Phoenix RPC12 Administrator's Guide

P/N 83-00004289-11
Revision B
August 2008

Contents

Preface

This guide describes how to use the web-browser interface (RAIDar) to configure and manage a Phoenix™ RPC12 storage system, and applies to the following enclosures:

- 69501 FC Controller Enclosure
- 69503 iSCSI Controller Enclosure
- SAS Expansion Enclosure

This guide introduces the interface, describes navigation, and presents the functions you need to perform to configure the system for the first time and then manage and monitor it.

This guide is written for system administrators who are familiar with Fibre Channel (FC), Internet SCSI (iSCSI), and Serial Attached SCSI (SAS) configurations, network administration, and RAID technology.

Before You Read This Book

Before you begin to follow the procedures in this book, you must have already installed the system and learned of any late-breaking information related to system operation as described in the *Getting Started Guide* and *Release Notes*.

Typographic Conventions

Typeface[1]	Meaning	Examples
AaBbCc123	Book title, new term, or emphasized word	See the *Release Notes*. A virtual disk (*vdisk*) can .... You *must* be an advanced user to ....
`AaBbCc123`	Directory or file name, value, command, or on-screen output	The default file name is `store.logs`. The default IP address is `10.0.0.1`. Type `exit`.
`AaBbCc123`	Text you type, contrasted with on-screen output	`# set password` `Enter new password:`
`AaBbCc123`	Variable text you replace with an actual value	Use the format `http://`*`ip-address`*.

1 The fonts used in your viewer might differ.

Related Documentation

Application	Title	Part Number
Site planning information	*Phoenix Storage System Site Planning Guide*	83-00004283
Late-breaking information not included in the documentation set	*Phoenix 69501 Release Notes* *Phoenix 69503 Release Notes*	83-00004282 83-00005032
Installing and configuring hardware	*Phoenix 69501 Getting Started Guide* *Phoenix 69503 Getting Started Guide*	83-00004284 83-00005034
Using the command-line interface (CLI)	*Phoenix RPC12 CLI Reference Manual*	83-00004288
Troubleshooting	*Phoenix RPC12 Troubleshooting Guide*	83-00004287
Recommendations for maximizing reliability, accessibility, and serviceability	*Phoenix RPC12 Best Practices Guide*	83-00004286

CHAPTER 1

Introducing and Using RAIDar

This chapter introduces the web-browser interface for Phoenix storage systems by describing its key elements, including icons and navigation. It describes how to configure this interface by setting system preferences, configuring users, and managing licenses.

This chapter contains the following sections:

What is RAIDar?

Each controller module contains a RAIDar web server. RAIDar is the primary interface for monitoring and managing Phoenix storage systems from a management host (a workstation with direct or network connections to a storage system's management ports).

RAIDar enables you to configure and maintain the storage for data hosts (a host that reads/writes data to the storage system). You can manage the following physical and logical storage components:

- Controller enclosures and controller modules
- Expansion enclosures and expansion modules
- Power-and-cooling modules
- Drive modules
- Virtual disks (*vdisks*)
- Volumes
- Volume-to-host mappings, including logical unit number (LUN) assignments
- Master volumes, snap pools, and snapshots

RAIDar also includes monitoring and diagnostic features that enhance the reliability, availability, and serviceability (RAS) of your storage system. You can configure the transmission of event notifications (alerts), which can be sent to the screen or to email addresses, and Simple Network Management Protocol (SNMP) traps, which can be sent to an SNMP application. Events are also recorded in an event log on the storage system from which they can be viewed.

In a dual-controller system, you can access all RAIDar functions from either controller. If one controller becomes unavailable, you can continue to monitor and manage the storage system from the partner controller.

Note – You can also monitor and manage storage using the scriptable command-line interface (CLI), as described in the *CLI Reference Manual*.

Preparing to Use RAIDar

RAIDar supports the following browsers:

- Microsoft Internet Explorer 5.5 or later
- Mozilla Firefox 1.0.7 or later

Configure your browser as follows:

- RAIDar uses pop-up windows to display various statistics and progress messages. You must enable pop-up windows on your browser for proper operation.
- (Internet Explorer only) When web page caching is enabled in RAIDar, which is the default, you can optimize performance by setting the browser *never* to check for newer versions of stored pages. For information about the RAIDar caching mode, see "Setting the Web Page Caching Mode" on page 45.

Note – This setting is used for all sites you visit with the browser.

- To optimize the display, use a color monitor and set its color quality to the highest setting.
- To ensure that you can view animated status icons, set your browser to play animations.
- To ensure that you can navigate beyond the RAIDar login page, set your browser's local-intranet security option to medium or medium-low.

Logging In and Out of RAIDar

RAIDar distinguishes users by the IP addresses from which they log in. If you log in to RAIDar using multiple browser instances on the same management host, RAIDar considers all instances as a single user. Actions you take in one RAIDar instance are reflected in the other RAIDar instances on the same host. A controller permits only one browser instance for each management host IP address. Do not log in more than once from the same host.

Each RAIDar user has a Monitor or Manage access level, as described in "User Roles" on page 26. If a Manage user does not log out of RAIDar when finished using it, other Manage users cannot log in to the same controller, and the IP address stays logged in until the auto-logout timeout expires.

RAIDar permits one Manage user and up to five Monitor users to be logged into a controller at the same time.

To log in to RAIDar:

1. In a web browser's address field, type the IP address of one of the Ethernet management ports and press Enter.

   The RAIDar Login page is displayed.

   Note – If the page does not display, verify that you entered the correct IP address. In a dual configuration, if you still cannot access a controller, try entering the IP address of the partner controller's Ethernet port. If you still cannot access RAIDar, use the CLI `show network-parameters` command to verify the IP addresses.

   ---

2. On the login page, type the username and password.

   The default Manage username is `manage` and the default password is `!manage`.

3. Click Log In.

   The Status Summary page displays the overall status and health of the system.

   Note – If you cannot navigate past the login page, your browser's security setting might be too high for web pages on the local intranet. Set it lower and try to log in again.

To log out of RAIDar:

1. Click Log Off at the bottom of the menu.

   The Log Off page is displayed.

2. Click Log Off.

Understanding the Interface

The topics in this section describe elements of RAIDar pages and provide help for navigating pages:

Interface Elements

The following figure shows RAIDar as it would appear for a dual-controller system with one healthy virtual disk and one virtual disk being initialized.

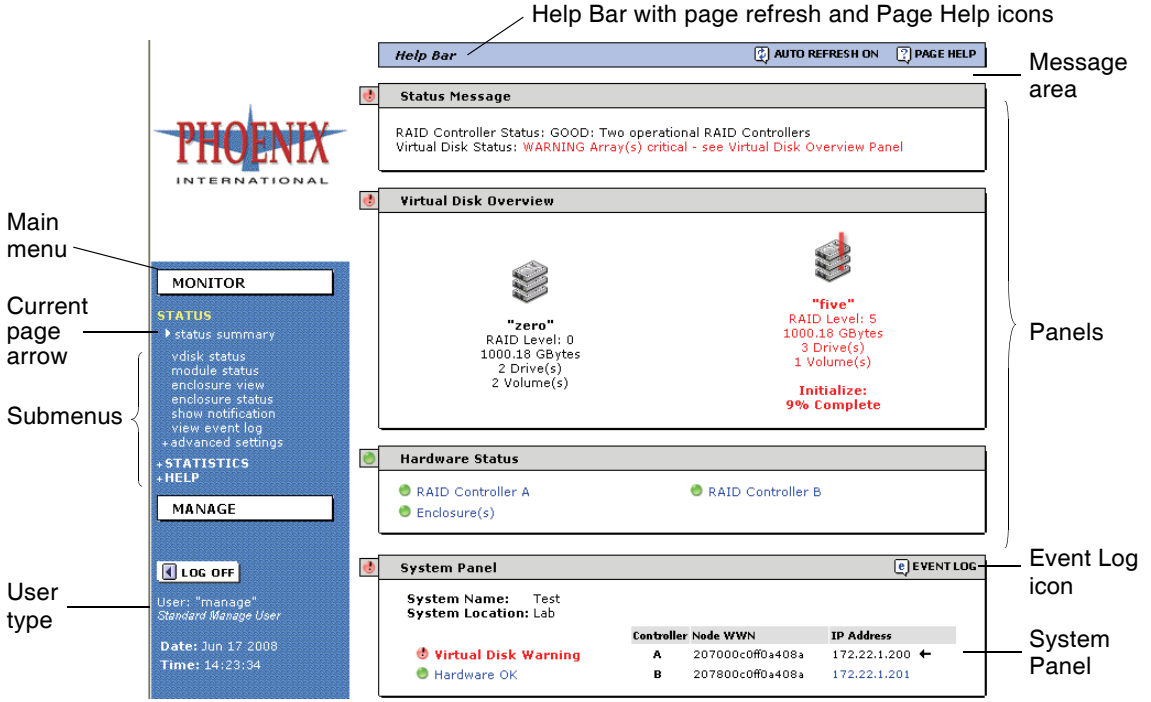

Figure 1-1 Key Elements of RAIDar Pages

The following table describes the key elements of RAIDar pages.

Table 1-1 Key Elements of RAIDar Pages

Element	Description
Menu area	This area on each page includes monitoring functions in the Monitor menu, management functions in the Manage menu, and a Log Off function. An arrow icon marks the menu item for the currently displayed page. The type of user that is logged in is displayed beneath the Log Off button. (User types are described in "Access Privileges" on page 26.)
Help Bar	Click the Page Help icon to display help for the current page (see "Help Bar Icons" on page 20).
Message area	After you make configuration changes a message in this area indicates whether the changes succeeded or failed.
Panels	Panels show information about your system configuration and available functions. The icon on the left side of a panel shows the overall status of items in the panel. In panels, blue text (or red text if there is a failure) is a link to additional information.
System Panel	Each page includes this panel, which shows the overall state of the two categories of system operation: virtual disk health and hardware health (see Figure 1-2). To view more detailed information, click a category name. Click the Event Log icon to display the event log page (see "Displaying the Event Log" on page 165).

Navigating RAIDar

The following table describes how to navigate RAIDar pages.

Table 1-2 RAIDar Navigation

Task	Navigation Action
Select a menu item	Click the menu item in the menu on the left side of each page. When you click some menu items, the menu changes to display different submenus. This book uses the following convention to indicate the steps in navigating to a function: Select Menu > Submenu > Function
View more information	Click a virtual disk or volume icon or click blue or red text.
View the most current status information on the current page	Click your browser's Refresh or Reload button.

While navigating RAIDar pages:

- Do not use the Tab key. Because RAIDar requires frequent visible and invisible data refreshes, using the Tab key can cause unpredictable behavior.
- When you use your browser's Back button, the last page viewed is displayed, but its content is not updated to show current data. If you use the Back button, manually refresh the page to get current data.
- Do not try to perform commands on multiple items (such as virtual disks and disk drives) by holding down the Shift key while clicking them with the mouse. In many cases, this will cause a new browser window and an error message to display.

Help Bar Icons

The Help Bar at the top of each page can include event notification, page refresh, and page help icons.

- VISUAL EVENT ALERT – An event occurred that is configured to display a visual alert. Click this icon to view the most recent events monitored by Event Notification. To control how you receive information about events, see "Configuring Event Notification" on page 47.
- AUTO REFRESH ON – The page refreshes automatically when its content changes status. The Page Refresh Rate preference controls how fast the page refreshes; see "Configuring Preferences" on page 24.
- CLICK TO REFRESH – The page does not refresh automatically to ensure that settings or values you are changing are not lost during the refresh process. To refresh the page manually, click this icon. Any unsaved changes are cleared.
- No refresh icon – The page has static content that does not need to be refreshed.
- PAGE HELP – Click this icon to show help for the current page.

Virtual Disk Icons

RAIDar has many status pages that enable you to monitor the status of your system, virtual disks, and disk drives. The top panel on many status pages includes an icon for each virtual disk with information about the selected virtual disk below it. The following table describes the virtual disk status icons.

Table 1-3 Virtual Disk Status Icons

Icon	Description
	Virtual disk is online with all drives working.
	Virtual disk is online in a critical state. The virtual disk can perform I/O functions for data hosts but is not fault tolerant. It is normal for a virtual disk to be critical while it is initializing online, or is reconstructing after a drive failure; in both cases, the utility name and percent complete are shown. If a virtual disk is critical for any reason other than online initialization or reconstruction, review the status information and take the appropriate action, such as replacing a disk drive. You can use a virtual disk in this state but resolve the problem as soon as possible. Refer to the *Troubleshooting Guide* for more information.

Table 1-3 Virtual Disk Status Icons *(Continued)*

Icon	Description
	RAID-6 virtual disk is online in a degraded state. The virtual disk can perform I/O functions for data hosts and is fault tolerant, but has degraded performance due to one missing drive. This might indicate that a disk drive has failed in the virtual disk or that the virtual disk is reconstructing. You can use a virtual disk in this state but resolve the problem as soon as possible. Refer to the *Troubleshooting Guide* for more information.
	A Verify or Expand utility is running. The utility and percent complete also appear. You can use a virtual disk in this state.
	Virtual disk is initializing offline or is offline for another reason. When a virtual disk is initializing offline the status indicates that the virtual disk is initializing and specifies the percent complete. You cannot use an offline virtual disk. If the virtual disk is offline and is not initializing then you must begin your disaster recovery process. Refer to the *Troubleshooting Guide* for more information.
	Virtual disk is quarantined. After a restart or rescan, one or more drives that are part of a formerly fault-tolerant virtual disk were missing. This virtual disk has been frozen until the drives are added back into the system or until the virtual disk is manually dequarantined using Virtual Disk Quarantine. A virtual disk can become quarantined if disk drives are removed, their enclosures are not powered on, or their enclosures are slow to power on. For more information, see "Dequarantining a Virtual Disk" on page 71.

Note – The Critical, Offline, and Quarantined icons are animated. To ensure that they display correctly, verify that animation is enabled in your browser.

System Panel

The System Panel at the bottom of each page includes system information, the overall status of system components, controller information, and the Event Log icon.

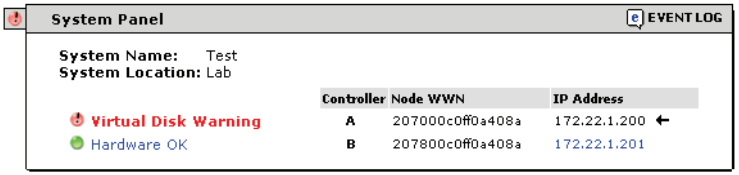

Figure 1-2 System Panel

The following information is shown:

- System information – The system's name and location.
- Overall status – The Virtual Disk category shows the status of virtual disks in the system. The Hardware category shows the status of I/O modules and enclosure components. To see more detailed information, click the health value. An icon shows the current status for each category:
  - – A green icon indicates that all virtual disks or hardware components are operating normally.
  - – A red icon with an exclamation point indicates that at least one virtual disk or hardware component is operating in a degraded state or is offline.
- Controller information – The node world wide name (WWN) and Ethernet management port IP address of each controller. A black arrow icon ← identifies the controller you are accessing. If the Ethernet ports of both controller modules are connected to the same network, clicking the IP address link for the other controller opens a new browser window for login.
- EVENT LOG – Click this icon to display the Event Log page. See "Displaying the Event Log" on page 165 for more information about the event log.

Help Menu

The Help submenu in the Monitor menu provides the following options for getting online help:

- Getting Started – Shows information about configuring your browser to use RAIDar and shows tips for using RAIDar.
- Subject Index – Provides an alphabetically ordered list of actions you can perform in RAIDar. If you have the proper role to perform an action, a link to the associated page is displayed; otherwise, the name of the associated page and the role required to access it are displayed. This index provides an alternative way to find where you can view system information or configure system settings.
- Support Information – Optionally displayed in a customized interface to describe how to get technical support and product documentation.

Size Representations in RAIDar

Data capacity and I/O statistics are calculated in decimal (base 10). Memory size is calculated in binary (base 2) using the memory industry standard.

Table 1-4 Size Representations in RAIDar

Unit	Data Capacity and I/O Statistics	Memory
Kbyte (KB)	1000 bytes	1024 bytes
Mbyte (MB)	1000 Kbyte 1 million bytes	1024 Kbyte 1,048,576 bytes
Gbyte (GB)	1000 Mbyte 1 billion bytes	1024 Mbyte 1,073,741,824 bytes
Tbyte (TB)	1000 Gbyte 1 trillion bytes	1024 Gbyte 1,099,511,627,776 bytes

Configuring RAIDar

The topics in this section describe how to configure RAIDar for your needs:

Configuring Preferences

You can configure RAIDar preferences to meet your needs. When you change preferences, the change takes effect immediately for the current RAIDar session but does not affect other active sessions logged in to either controller. RAIDar sessions started after the preferences are changed use the new preferences.

To configure preferences:

1. Select Manage > General Config > System Preferences.
2. Set the following options:

Preference	Description
Page Refresh Rate	Select how often you want RAIDar pages to refresh based on the speed of your computer and Ethernet connection. • Fast – Use for fast computers with a fast Ethernet connection. For example, Pentium III 500 MHz or higher with a T1 connection. The default is Fast. • Medium – Use for slower computers with a slower Ethernet connection. For example, Pentium III 400 MHz and slower with a cable modem or DSL connection. • Slow – Use for the slowest computers with a slow Ethernet connection. For example, Pentium II 200 MHz and slower with a dial-up modem of 33.5-5 Kbit/sec connection.
Auto-Logout Timeout	Type the number of minutes that a user's RAIDar session can be idle before being automatically logged off. The allowed values are 0–255 minutes, where 0 means no timeout. The default is 30 minutes.
Temperature Display Mode	Select Fahrenheit or Celsius for all temperature status indications. The default is Celsius.

Preference	Description
On Manage Login	Select which page to display when a Manage user logs in. • Go To Main Monitor Status Screen – The Status Summary page is displayed after login. This is the default. • Go Directly To Manage Screens – The Vdisk Status page is displayed after login.

3. Click Change Preferences.

Configuring User Access

By default, the system provides three users that can access the system. In addition to these users, which you can modify, you can add 10 other users (13 maximum). The user configuration function enables you to define user roles by setting specific access privileges. For each user you can set a password and enable or disable access to the following system interfaces: WBI (RAIDar), CLI (command-line interface), and FTP.

The default users are configured with the usernames, access levels, user types, and default passwords shown in the following table.

Table 1-5 Default User Configuration

Username	Access Level	User Type	Password
`monitor`	Monitor	Standard	`!monitor`
`manage`	Manage	Advanced	`!manage`
`ftp`	Manage	Advanced	`flash`

User Roles

Each user role is defined by an access level of either Monitor or Manage:

- Monitor – Enables access to functions on the Monitor menu.
- Manage – Enables access to functions on the Monitor and Manage menus.

Up to five Monitor users and only one Manage user can be logged in to each controller. RAIDar distinguishes users by their IP addresses. If you log in to RAIDar using multiple browser instances on the same management host, RAIDar considers all instances as a single user; actions you take in one instance are reflected in the other instances on the same host.

Note – If you are a Monitor user and you attempt to change a Manage user setting, RAIDar prompts you to log in as a Manage user. If you enter a correct Manage username and password, you are logged out of your Monitor session.

Access Privileges

User access privileges are based on the following user types:

- Standard – Enables access to most functions.
- Advanced – In addition to enabling Standard functions, enables access to infrequently used administrative functions.
- Diagnostic – In addition to enabling Standard and Advanced functions, enables access to troubleshooting functions.

  This guide describes Standard and Advanced functions only. Diagnostic functions are listed in Table D-3 in "RAIDar Menu Reference" on page 231. However, because they are for troubleshooting purposes, describing them further is outside the scope of this guide. Refer to the *Troubleshooting Guide* for information related to Diagnostic user functions.

How User Configuration Affects the RAIDar Menu

User configuration enables you to control which functions a user can access based on the user's role (assigned user type and access level). For example:

- In the Monitor > Status > Advanced Settings menu, Advanced Monitor users can view temperature and power status information which Standard Monitor users cannot.
- In the Manage > Utilities menu, Advanced Manage users can access a Host Utilities submenu which Standard Manage users cannot.

The current user's name and role are displayed at the bottom of the RAIDar menu.

Note – Appendix D lists all RAIDar functions and the roles required to use them.

Modifying Users

To modify a user:

1. Select Manage > General Config > User Configuration > Modify Users.

   The System User List displays the current list of configured users.

2. Select a user from the Username drop-down list and click Modify User.

   The Modify Selected User panel is displayed.

3. Change the username.

   The name is case-sensitive and can include 19 characters. Allowed characters include letters, numbers, hyphens, underscores, and spaces.

   Note – For security reasons, create different usernames unique to your site. If you keep the default ones, change their default passwords.

   ---

4. Change the user's password.

   The password is case-sensitive and can include 19 characters. Allowed characters include letters, numbers, hyphens, underscores, and spaces.

5. Change the user's access level:
   - Monitor enables access to all functions on the Monitor menu.
   - Manage enables access to all functions on the Monitor and Manage menus.

Note – In a list the current setting is marked with two asterisks (**).

---

6. Change the user type:
   - Standard enables access to most functions.
   - Advanced additionally enables access to infrequently used administrative functions.
   - Diagnostic additionally enables access to troubleshooting functions for use by service technicians.
7. Enable or disable user access to system interfaces:
   - WBI – The web-browser interface, RAIDar
   - CLI – The command-line interface
   - FTP – The file transfer protocol interface

Note – A system interface can be used only if the corresponding network management service is enabled on the Manage > General Config > Services Security page.

---

8. Click Save Changes.

   The System User List is updated.

Adding Users

RAIDar allows a maximum of 13 users, including the three default users shown in Table 1-5.

To add a user:

1. Select Manage > General Config > User Configuration > Add Users.

   The Add System User panel displays the current list of configured users.
2. Type a new username.

   The name is case-sensitive and can include 19 characters. Allowed characters include letters, numbers, hyphens, underscores, and spaces.

3. Type a password.

   The password is case-sensitive and can include # characters. Allowed characters include letters, numbers, hyphens, underscores, and spaces.

4. Select an access level:
   - Monitor enables access to all functions on the Monitor menu.
   - Manage enables access to all functions on the Monitor and Manage menus.

5. Select a user type:
   - Standard enables access to most functions.
   - Advanced additionally enables access to infrequently used administrative functions.
   - Diagnostic additionally enables access to troubleshooting functions for use by service technicians.

6. Enable or disable the user's access to system interfaces:
   - WBI – The web-browser interface, RAIDar
   - CLI – The command-line interface
   - FTP – The file transfer protocol interface

Note – A system interface can be used only if the corresponding network management service is enabled on the Manage > General Config > Services Security page.

---

7. Click Add User.

   The user is added to the Add System User panel.

Deleting Users

You can delete any user from the system, including the default users.

Note – The deletion of a system user cannot be undone.

To delete a user:

1. Select Manage > General Config > User Configuration > Delete Users.

   The System User List panel displays the current list of configured users.

2. Select a user from the Username drop-down list and click Delete User.

   A confirmation prompt is displayed.

3. Click OK to confirm the operation or Cancel to stop it.

   If you clicked OK, a message indicates whether the operation succeeded. If it succeeded, the user is removed from the System User List panel.

Managing Licenses

You can purchase a license to expand baseline functionality. Depending on the license options you purchase, you can:

- Increase the number of snapshots that can be taken
- Enable the ability to copy volumes
- Enable use of the host-based Volume Shadow Copy Service (VSS)
- Enable use of the host-based Virtual Disk Service (VDS)

You must obtain and install the license certificate file that enables the purchased options. A valid license certificate file meets the following requirements:

- The file is being installed on the controller enclosure for which the license file was generated.
- The file is a text file with a `.txt` extension.
- The file cannot be edited in any way.

To obtain and install a license, refer to *Obtaining and Installing the License Certificate File* at:
`http://crc.dothill.com`

Viewing Installed Licenses

To view installed licenses:

- Select Manage > General Config > License Management > Installed Licenses.

  The Licensed Features Installed panel shows whether a license certificate file is installed and the status of licensed features. For a licensed feature that has a quantity limit, the panel shows the maximum quantity available with the license and the baseline maximum quantity available without a license.

  - Snapshot – Shows whether snapshot services are enabled or disabled.
  - Snapshots Available – The maximum number of snapshots permitted, followed by the default number permitted.
  - Snapshots In Use – The number of snapshots that exist on the system.
  - Volume Copy – Shows whether volume-copy services are enabled or disabled.
  - VSS – Shows whether use of VSS is enabled or disabled. Volume Shadow Copy Service (VSS) is an API that enables snapshots to be managed by third-party applications.
  - VDS – Shows whether use of VDS is enabled or disabled. Virtual Disk Service (VDS) is an API that enables virtual disks and volumes to be managed by third-party applications.
  - License File Signature – License value from the installed license certificate file.

  For example, the following figure shows that the installed license allows more than the default (baseline) number of snapshots, and two snapshots are in use.

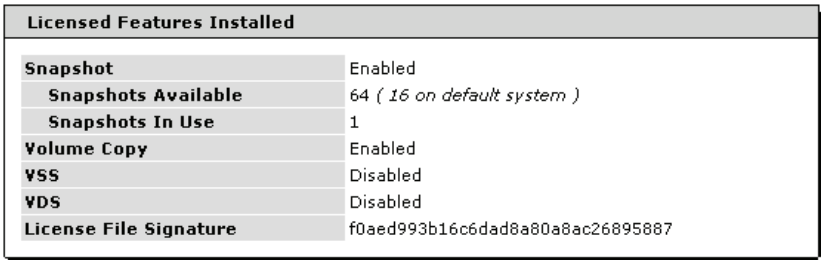

Figure 1-3 Licensed Features Installed Panel

Installing a License

To install a license certificate file that has been generated for this system:

1. Ensure that the license file has been saved to a location on your network that this system can access.
2. Select Manage > General Config > License Management > Install A License.

   The Load License File panel is displayed.
3. Click Browse to navigate to the location of the file and click Open.
4. Click Load License File.

   The license file is installed and a message is displayed informing you that it was installed successfully. The changes produced by the license file take effect immediately.

CHAPTER 2

Configuring Your System for the First Time

This chapter describes how to use RAIDar to configure your system for the first time. It contains the following sections:

Configuring User Access

By default, the system provides three users that can access the system. In addition to these users, which you can modify, you can add 10 other users (13 maximum). The user configuration function enables you to define user roles by granting specific access privileges.

For detailed information about adding, modifying, and deleting users, see "Configuring User Access" on page 25.

Setting System Information

You can specify information about the system to enable you to identify it. The system name and location are displayed in the System Panel.

To set system information:

1. Select Manage > General Config > System Information.

   The System Information panel is displayed.

2. Type information in each field.

   Each value can include 74 characters. Allowed characters include letters, numbers, hyphens, underscores, and spaces.

   - System Name – Name of the system as seen by other systems on the network. The default is Uninitialized Name.
   - System Contact – Name of a contact person responsible for the system. The default is Uninitialized Contact.
   - System Location – Location of the system. The default is Uninitialized Location.
   - System Information – Other information you want to specify, such as the system's purpose or type. The default is Uninitialized Info.

3. Click Save Changes.

Setting Date and Time

You can set the system's date and time, which are displayed at the bottom of the menu area. It is important to set the date and time so that entries in system logs have correct time stamps.

To set the system date and time:

1. Select Manage > General Config > Set Date/Time.

   The Current System Date & Time panel shows the current date and time.

2. In the Set System Date panel, select the current month, day, and year.

3. In the Set System Time panel, type time values using a 24-hour clock (where hour 8 represents 8 a.m. and hour 20 represents 8 p.m.) and select the proper time zone.

   The time zone setting affects the date stamp on email messages sent by the Event Notification function.

4. Click Change Date/Time.

Configuring Host Ports

This section describes how to configure host ports on Fibre Channel (FC) controller modules or on iSCSI controller modules.

Configuring FC Host Ports

On the Host Port Configuration page you can view the location, link speed, and topology of each FC host port in each controller module.

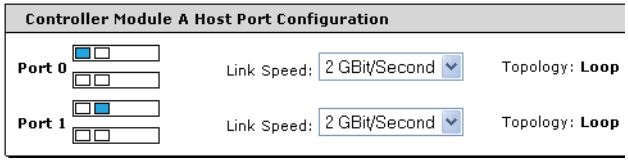

Figure 2-1 Configuration Settings for Host Ports on Controller Module A

This page shows the following information:

- Port location – The shaded box represents the port whose settings are shown to the right.
- Link Speed – 2 GBit/Second or 4 GBit/Second. A host port's link speed must match the speed of the HBA or switch to which the port is connected.
- Topology – Either Loop or Point to Point. “Not Available IOM Down” appears if the controller is down and topology information is unavailable.

On this page you can set FC host port link speeds. From the Advanced Options panel you can view and set the following:

- FC host port loop ID
- FC host port interconnect status (dual-controller system only)
- FC host port topology

Setting FC Host Port Link Speed

A host port's link speed must match the speed of the host (HBA or switch) to which the port is connected. In a dual-controller system, setting the speed of host port 0 on one controller also sets the speed of host port 1 on the other controller.

A speed mismatch with the host prevents the host from accessing the storage system.

Note – Model 0 host interface modules do not support 4-Gbit/second link speed with host port interconnects enabled. To determine which model your controllers use, see "Controller Versions" on page 154.

To set host port link speed:

1. Select Manage > General Config > Host Port Configuration.
2. For each port that is connected to a host, set the appropriate speed.

   The default is 2 Gbit/second.
3. Click Update Host Port Configuration.

Setting FC Host Port Loop IDs

A loop ID identifies a controller to a data host. A loop ID is only a requested value. The controller requests the specified ID when it arbitrates on the FC loop but the actual loop IDs assigned to each port during FC loop initialization might differ. This page shows the requested and current loop ID for each controller.

Generally, you only need to change this setting if you want a controller to be at a specific address; your system checks addresses in reverse order (lowest address first); or, an application requires that specific IDs be assigned to recognize the system.

To set host port loop IDs:

1. Select Manage > General Config > Host Port Configuration.
2. In the Advanced Options panel, click Change FC Loop ID.

   The Requested Loop ID for Host Ports panel displays the currently requested loop ID and the current loop ID for each controller's host ports.
3. Select a requested loop ID for each controller:
   - Soft – Select this software addressing setting if it doesn't matter whether the controller's loop ID changes after you power down and power up or after a loop initialization process (LIP). This setting enables the FC loop initialization process to determine the loop ID.
   - 0–125 – Select a specific number if you want the loop ID to stay the same after you power down and power up. RAIDar cannot determine which loop IDs are available. If the controller cannot get the specified loop ID during the loop initialization process, it tries to get a soft address.
4. Click Save And Continue.

   A controller restart is required for the change to take effect.
5. Click OK to restart the controller.

   When processing is complete, the main Host Port Configuration page is displayed.
6. To verify that the loop ID you wanted was assigned, click Change FC Loop ID and check the current loop ID values.

Configuring FC Host Port Interconnects

In an FC storage system, the host port interconnects act as an internal switch to provide data-path redundancy.

When the host port interconnects are enabled, port 0 on each controller is cross-connected to port 1 on the other controller. This provides redundancy in the event of failover by making volumes owned by either controller accessible from either controller.

When the host port interconnects are disabled, volumes owned by a controller are accessible from its host ports only. This is the default.

For a single-controller FC system, host port interconnects are always disabled.

For a dual-controller FC system in a direct attach configuration, host port interconnects are typically enabled — except in configurations where fault-tolerance is not required and the highest performance is required.

For a dual-controller FC system in a switch attach configuration, host port interconnects are typically disabled — except for applications where fault tolerance is required and highest performance is not, or when not enough switch ports are available.

You cannot enable host port interconnects if any host port is set to point-to-point topology.

Note – Model 0 host interface modules do not support 4-Gbit/second link speed with host port interconnects enabled. To determine which model your controllers use, see "Controller Versions" on page 154.

To change the host port interconnect setting:

1. Select Manage > General Config > Host Port Configuration.
2. In the Advanced Options panel, click Change FC Port Interconnect Settings.

   The Host Port Configuration panel displays the current interconnect setting.
3. Set Internal Host Port Interconnect to Interconnected (enabled) or Straight-through (disabled).

   The default is Straight-through.

   This setting affects all host ports on both controllers.

4. Click Save And Continue.

   The main Host Port Configuration page is displayed.

Setting FC Host Port Topology

For Phoenix storage systems, *topology* means the path that data travels between devices: either through a series of connected devices (loop) or directly from one device to another (point-to-point). In a switch-attach configuration, either topology is supported but loop is preferred. In a direct-attach configuration, only loop is supported.

In a dual-controller FC storage system, the topology to set for host ports depends on the system's host port interconnect setting (see "Configuring FC Host Port Interconnects" on page 38), and affects host access to volumes during failover, when their owning controller's host ports are inaccessible. This relationship is described in the following paragraphs.

Volumes can be mapped with access privileges through specific host ports to data hosts, with a LUN that identifies each mapping. Host access to volumes during a controller failover is determined by the storage system's host port interconnect and topology settings. For example, assume volumes are mapped through controller A's host ports and controller A fails over to controller B. The host can access controller A's volumes through controller B's host ports as follows:

- If all host ports are set to loop topology, both controllers' volumes are presented on controller B's host ports.
- If one or more host ports are set to point-to-point topology, controller B presents its volumes on half of its host ports and presents controller A's volumes on the remaining host ports.

If host port interconnects are enabled, the paired ports are connected in a loop and must be set to use loop topology. Changing the topology setting for one host port automatically changes the setting for the paired port on the partner controller.

If host port interconnects are disabled, you can change the topology setting for each host port individually.

Note – In a switch-attach configuration, if you change from loop to point-to-point after already establishing a public loop connection, the switch might ignore subsequent attempts to perform point-to-point initialization.

To set host port topology:

1. Select Manage > General Config > Host Port Configuration.
2. In the Advanced Options panel, click Change Host Port Topology.
3. For each port that is connected to a host, set the appropriate topology.

   The default is Loop (Fibre Channel Arbitrated Loop).
4. Click Save And Continue.

   The main Host Port Configuration page is displayed.

Note – For a system using loop topology, you might need to reset a host link to fix a host connection or configuration problem. See “Resetting Host Channels” on page 182 for steps to reset a host link.

Configuring iSCSI Host Ports

You can configure the following network parameters for each iSCSI port on each controller module. The 69503 supports only static addressing using IPv4 format.

- IP Address – IP address for a specific port. The system uses port 0 of each controller as one failover pair, and port 1 of each controller as a second failover pair. Therefore, port 0 of each controller must in the same subnet, and port 1 of each controller should be in a second subnet. For example:
  - Controller A port 0: 10.10.10.100
  - Controller A port 1: 10.11.10.120
  - Controller B port 0: 10.10.10.110
  - Controller B port 1: 10.11.10.130
- IP Mask – IP subnet mask for a specific port. The default is 255.255.255.0.
- Gateway – Gateway IP address for a specific port. The default is 0.0.0.0.

Caution – Changing IP settings can cause data hosts to lose access to the storage system.

To configure host ports:

1. Select Manage > General Config > Host Port Configuration.
2. For each port that is connected to a host, set the appropriate values.
3. Click Update Host Port Configuration.

Configuring Ethernet Management Ports

You can configure addressing parameters for each controller's Ethernet management port and the timeout value for Telnet sessions. You can also view and configure the SNMP event filter and the web page caching mode.

If you accessed RAIDar for the first time using the default IP address, you should set the IP address for each controller. You can also the change the IP settings as needed.

Caution – Changing IP settings can cause management hosts to lose access to the storage system.

Using DHCP to Obtain IP Settings

In DHCP mode, Ethernet management 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.

To use DHCP to obtain IP values for Ethernet management ports:

1. Select Manage > General Config > LAN Configuration.
2. In the IP Address Assignment panel, set Source For IP Address to DHCP.

   Settings in the RAID Controller IP Configuration panels are ignored.
3. Click Change LAN Configuration.

   The controllers try to obtain IP values from the DHCP server. The new IP values are displayed in a pop-up window. Record the new addresses.

   After 15 seconds you will be logged out and the browser will try to reconnect to RAIDar using the new IP address.

Using Static IP Settings

To set IP values for Ethernet management ports:

1. Select Manage > General Config > LAN Configuration.
2. In the IP Address Assignment panel, set Source For IP Address to Manual.
3. In the IP Configuration panel for each controller, set appropriate values for your network. Use dotted decimal notation.
   - The default IP address is 10.0.0.2 for controller A and 10.0.0.3 for controller B.
   - The default IP subnet mask is 255.255.255.0.
   - The default gateway IP address is 10.0.0.1.

   You must set a unique IP address for each Ethernet port. Record the IP values you assign.
4. Click Change LAN Configuration.

   After 15 seconds you will be logged out and the browser will try to reconnect to RAIDar using the new IP address.

Setting the Telnet Timeout

You can set the number of idle minutes before a Telnet connection to the storage system is automatically terminated. To set the Telnet timeout:

1. Select Manage > General Config > LAN Configuration.
2. In the Telnet Configuration panel, set the Timeout value.

   The allowed values are 0–255 minutes, where 0 means no timeout. The default is 60 minutes.
3. Click Change LAN Configuration.

Setting the SNMP Event Table Filter

Your storage system supports the following management information base (MIB) types for use with Simple Network Management Protocol (SNMP):

- MIB-II (RFC 1213), which reports basic system and TCP/IP statistics for the network stack
- FA2.2 (Fibre Alliance), which reports firmware and hardware revisions, sensor data, host port data, and events

You can filter the types of events that are included in the FA2.2 event table. The filter is applied as events are put into the table. Changing the filter does not affect events already recorded in the table; therefore, old events are reported even though they might not meet the current filter criteria.

Note – The event table is held in memory and is not an externally accessible file. For information on viewing the event log, see "Displaying the Event Log" on page 165.

For more information about using SNMP, see Appendix A.

To set the SNMP event table filter:

1. Select Manage > General Config > LAN Configuration.
2. In the Advanced LAN Options panel, click Advanced Options.

   The SNMP Event Table Configuration panel is displayed.
3. Set Event Table Filter to one of the following options:
   - Informational – Puts all events into the table. This is the default.
   - Warning – Puts warning and error events into the table.
   - Error – Puts only error events into the table. Error events are the most severe.
4. Click Change SNMP Event Table Configuration.

Setting the Web Page Caching Mode

The web page caching mode controls how RAIDar handles web page names. The names interact with your browser's caching operations to determine which pages and image files are retrieved.

To set the web page caching mode:

1. Select Manage > General Config > LAN Configuration.
2. In the Advanced LAN Options panel, click Advanced Options.
3. In the Change Web Page Caching Mode panel, set Web Page Caching Mode to Enabled or Disabled:
   - Enabled – Causes RAIDar to generate unique page names for all main web page accesses. This setting forces the web browser to always retrieve a new page from the system when needed. This mode is essential if your network has any kind of a proxy server that might be caching web requests from the system, which is undesirable as it could cause old pages or data to be displayed. This mode also prevents your web browser from caching pages that it shouldn't. This mode is the default.
   - Disabled – Web page names requested from RAIDar are not unique so you must assure that your browser and network are set up correctly to always retrieve a new page from the system when requested. You must also perform a top-level browser refresh (or close a browser and open a new one) to make the change take effect.
4. Click Change Web Page Caching Mode.

Configuring Network Management Services

You can configure network management services and in-band management services to limit the ways in which users and host-based management applications can access the system. If a service is disabled, it continues to run but cannot be accessed.

For information about permitting users to use enabled WBI, CLI, or FTP services, see “Configuring User Access” on page 25.

For information about in-band management services, see “Configuring In-band Management Services” on page 127.

To configure network management services:

1. Select Manage > General Config > Services Security.
2. In the Network Management Services panel, set these options:
   - Web Browser Interface (WBI) – RAIDar, the primary interface for managing the system. You can enable use of HTTP, of HTTPS for increased security, or both. The default is Enabled with HTTP and HTTPS.
   - Command Line Interface (CLI) – An advanced user interface for managing the system. You can enable use of Telnet, of SSH (secure shell) for increased security, or both. The default is Enabled with Telnet and SSH.
   - Storage Management Initiative Specification (SMIS) – Used for remote management of the system through your network. The default is Enabled.
   - File Transfer Protocol (FTP) – Used as an alternative to the WBI for upgrading system software. The default is Disabled.
   - Simple Network Management Protocol (SNMP) – Used for remote monitoring of the system through your network. The default is Enabled.
3. Click Update Network Management Services.

Configuring Event Notification

You can configure how and under what conditions the system alerts you when specific events occur. The system generates events having three severity levels:

- **Critical** – Something related to the system or to a virtual disk has failed and requires immediate attention.
- **Warning** – Something related to the system or to a virtual disk has a problem. Correct the problem as soon as possible.
- **Informational** – A problem occurred that the system corrected, or a system change has been made. These events are purely informational; no action required.

You can:

- Choose to be notified of all events, categories of events, or individual events.
- Enable or disable different notification methods for different events. Methods include visual alerts, email alerts, and SNMP traps.
- Configure options for each notification method.

To view the current notification settings:

- Select Manage > Event Notification > Notification Summary.

Event notification is controlled by three levels of settings. The settings are listed in order of precedence, meaning that the first settings override subsequent settings.

- **Notification Enabled** – This is the highest level of control. Enable allows the notification selected by the lower levels. Disable prevents any event notification.
- **Event Categories Selected** – If an event category is selected, and any events of that type occur, notification occurs.
- **Individual Events Selected** – Individual events can be selected for notification.

Note – All events are logged to the event log whether notification is enabled or not. See “Displaying the Event Log” on page 165 for more information.

Enabling or Disabling Event Notification

You can enable or disable the following notification methods for selected event categories or individual events:

- **Visual Alerts** – RAIDar shows a visual alert indicator that a notification event has occurred. To see this, RAIDar must be operating on a management host.
- **Email Alerts** – The system sends an email containing the events that have occurred to the designated users.
- **SNMP Traps** – The system sends an SNMP trap to the designated trap host.

For each notification method you enable, configure its options and select event categories or specific events to monitor.

You can combine the event selections in any way that meets your needs. When one of these events occurs in the system, RAIDar notifies you based on your event notification settings.

Note – Selecting entire event categories can result in the system sending numerous event notifications. Select the categories that are most important to you.

Selecting Event Categories to Monitor

To optimally configure the remote event notification feature, you must first understand the following event category options in the Event Notification Summary panel:

- All Critical Events – Serious events that might indicate system failure and require intervention. For example, a virtual disk is down.
- All Warning Events – Events that might require intervention although the system is still operating. For example, a virtual disk is critical.
- All Informational Events – Events that you expect to occur. For example, a virtual disk verification has completed.

Typically, you will want to select All Critical Events and All Warning Events when you are using email notification because it prevents unwanted email and paging from being sent to an administrator or other designated person. Warning and Error messages typically require some form of action whereas informational events are used to track specific behaviors when troubleshooting.

To select event categories for notification:

1. On the Event Notification Summary page, for each category you want to be notified of, select a notification method.

   For example, to receive email for all critical events, in the All Critical Events row select only the Email Alerts check box. To receive no notification of informational events, clear all check boxes in the All Informational Events row.

2. Click Change Notification Settings.

Selecting Individual Events to Monitor

In addition to selecting event categories, a Diagnostic Manage user can select individual events to be notified of. For information on selecting individual events, refer to the *Troubleshooting Guide*.

Configuring Visual Alerts

You can set the following options for visual notification of events:

- How you access the event listing, in a pop-up window or from the Help Bar
- The maximum number of events that are displayed and can be acknowledged at one time
- Whether visual alerts are enabled or disabled

To configure visual notification:

1. Select Manage > Event Notification > Visual Configuration.
2. Set Visual Alerts Method to one of the following:
   - Page Notification – Shows a visual alert icon  in the Help Bar when a visual alert event occurs. This icon is a link to the Show Notification page, which lists the events that have occurred. The notification only occurs automatically if the page is an auto-refresh page. On a non-auto-refresh page, the notification is not displayed until you refresh the page or go to another page.
   - Popup Notification – Causes a pop-up window to show the visual alert events when they occur. This window is displayed for all pages, remains on top of all other windows, and remains until you acknowledge the events by clicking an Acknowledge button. This method is the default.

3. Select a value for Maximum Events to Display at One Time.

   RAIDar can display a maximum of 100 events at a time; the default is 10.

   For example, if 10 events can display at a time and 15 are pending then the pop-up window shows the first 10 events and clicking the Acknowledge button will show the remaining events and new events that might have occurred.

   If more than 100 are pending, the oldest ones are dropped.

4. Enable or disable visual notification.

   The default is Disable.

5. Click Change Visual Alerts Configuration.

Note – Special events prompt you to take a specific action based on the event. The prompt is displayed in place of the normal Acknowledge button. For example, when the "Unwritable cache data exists for virtual disk" event occurs, you are prompted to either keep or discard the cache data.

Note – To view events that have been acknowledged or that don't cause notifications, click EVENT LOG in the System Panel.

Configuring Email Alerts

You can configure the following options for email notification of events:

- Email addresses to send notification messages to
- A comment to include in each message
- The mail server IP address
- The sender name and domain name
- Whether email alerts are enabled or disabled

You can also test the email configuration.

To configure and test email notification:

1. Select Manage > Event Notification > Email Configuration.
2. Type values in the following fields:
   - Email Address 1–4 – Email addresses that the system should send notifications to. Email addresses must use the format *user-name@domain-name*.
   - Email Comment – Text to send with email messages. For example, you might want to identify the location, name, or use of the system.
   - Mail Server – 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 configured on the General Config > LAN Configuration page.
   - Domain Name – The domain name that, with the sender name, forms the "from" address for remote notification. Because this name is used as part of an email address, do not include spaces. If no domain name is set, a default name is created. If the domain name is not valid, some email servers will not process the mail.
   - Sender Name – The sender name that, with the domain name, forms the "from" address for remote notification. Because this name is used as part of an email address, do not include spaces. If no sender name is set, a default name is created.
3. Enable or disable email notification.

   The default is Disable.
4. Click Change Email Alerts Configuration.

5. Click Send Test Email.

   Each configured email address should receive the test message.

   If the test fails, check the following:

   - The configured email addresses are correct.
   - The gateway is properly configured to enable email to be sent across subnets, and the Mail Server value is the IP address of the subnet's router. For information about setting up IP addresses, see "Configuring Ethernet Management Ports" on page 42.
   - The domain name and sender name do not include spaces.
   - Some mail servers are set up to reject mail if the mail does not pass a mail filter. Verify that mail can be sent and received from the configured domain and sender.

Configuring SNMP Traps

You can configure the following options for SNMP notification of events:

- Read and write community strings
- IP addresses of hosts that are configured to receive SNMP traps

To configure SNMP traps:

1. Select Manage > Event Notification > SNMP Configuration.
2. Type values in the following fields:
   - SNMP Read Community – The SNMP read password for your network. The value is case-sensitive and can include 15 characters. The default is `public`.
   - SNMP Write Community – The SNMP write password for your network. The value is case-sensitive and can include 15 characters. The default is `private`.
   - SNMP Trap Host IP Address 1–3 – The IP addresses of host systems that are set up to receive SNMP traps.
3. Enable or disable SNMP traps.

   The default is No (disable).
4. Click Change SNMP Traps Configuration.

Changing the Cache Mirroring Mode

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

The advantage of ICPM is that the two controllers can achieve very high write bandwidth and still use write-back caching. User data is still safely stored in nonvolatile RAM, with backup power provided by super-capacitors should a power failure occur. This feature is useful for high-performance applications that do not require a fault-tolerant environment for operation; that is, where speed is more important than the possibility of data loss due to a drive fault prior to a write completion.

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

Data loss does not automatically occur if a controller experiences a software exception, or if a controller module is removed from the enclosure. If a controller should experience a software exception, the controller module goes offline; no data is lost, and it is written to disks when you restart the controller. However, if a controller is damaged in a nonrecoverable way then you might lose data in ICPM.

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

Note – Independent cache performance mode disables partner firmware upgrade. Controllers must be upgraded manually.

To change the cache mirroring mode:

1. Select Manage > General Config > System Configuration.
2. Click Advanced Options.
3. In the Change Independent Cache Performance Mode panel, select Enabled or Disabled. The default is Disabled.
4. Click Enable/Disable Independent Cache Performance Mode.

   The system automatically restarts both controllers, which takes several minutes to complete.

Saving the Configuration to a File

As an Advanced Manage user, you can save the storage system's configuration settings to a file. This enables you to make a backup of your settings in case a subsequent configuration change causes a problem, or if you want to apply one system's settings to another system. For information on restoring configuration data, see "Restoring a Saved Configuration File" on page 184.

The configuration file contains all system configuration data, including:

- LAN configuration settings
- Host port configuration settings
- Enclosure management settings
- Disk configuration settings
- Services security settings
- System information settings
- System preferences settings
- System configuration settings
- Event notification settings

The configuration file does not include configuration data for virtual disks and volumes. You do not need to save this data before replacing a controller or expansion module because the data is saved as metadata in the first sectors of associated disk drives.

To save system configuration data to a file on the management host or network:

1. Select Manage > Utilities > Configuration Utilities > Save Config File.
2. Click Save Configuration File.
3. If prompted to open or save the file, click Save.
4. If prompted to specify the file location and name, do so using a `.config` extension.

   The default file name is `saved_config.config`.

Note – If you are using Firefox and have a download directory set, the file is automatically saved there.

Restarting and Shutting Down a Controller

You can restart or shut down controllers when a controller is not working properly or when the system will be serviced or moved.

Restarting a Controller

You can restart one or both controllers when:

- RAIDar informs you that you have changed a configuration setting that requires restarting
- A controller does not seem to be working properly

When you restart a controller, its Management Controller and Storage Controller processors are shut down and then restarted and any data in write-back cache is written to disk.

Caution – If you restart the controller you are connected to, you lose access to the controller and you must reconnect to it and log back in to access RAIDar. If you restart both controllers, you lose access to RAIDar and users lose access to data until the controllers have restarted. You must then log back in to access RAIDar.

To restart a controller:

1. Select Manage > Restart System > Shut Down/Restart.
2. In the Restart Controller panel, select a controller option.
3. Click Restart.

   A confirmation prompt is displayed.
4. Click OK to confirm the operation or Cancel to stop it.

   If you clicked OK, a message indicates whether the operation succeeded.

Note – If the storage system is connected to a Microsoft Windows host, the following event is recorded in the Windows event log: A connection to the target was lost, but Initiator successfully reconnected to the target.

Shutting Down a Controller

Shut down a controller module before you remove it from an enclosure, or before you power off its enclosure for maintenance, repair, or a move. Shutting down a controller module halts I/O to that module, ensures that any data in the write cache is written to disk, and initiates failover to the partner controller, if it is active.

Caution – If you shut down both controller modules, you lose access to RAIDar and users lose access to data. To restart the controllers, turn off both power-and-cooling modules and then turn them back on.

To shut down a controller:

1. Select Manage > Restart System > Shut Down/Restart.
2. In the Shut Down panel, select a controller option.
3. Click Shut Down.

   A warning might appear that data access redundancy will be lost until the selected controller is restarted. This is an informational message that requires no action.
4. Confirm the operation by clicking OK.

Note – If the storage system is connected to a Microsoft Windows host, the following event is recorded in the Windows event log: Initiator failed to connect to the target.

CHAPTER 3

Managing Storage

This chapter describes how to use RAIDar to configure and manage virtual disks, spare disks, volumes, volume-to-host mappings, and to use volume snapshot features. It contains the following sections:

Creating Virtual Disks and Volumes

You can create a virtual disk when you have enough available disk drives of the same type for the RAID level you want to use. A maximum of 16 virtual disks per controller can exist. The controller safeguards against improperly combining SAS and SATA disk drives in a virtual disk. The system displays an error message if you choose drives that are not of the same type.

Each virtual disk is owned by only one of the controllers. For most purposes, it does not matter which controller owns a virtual disk because RAIDar automatically selects the owner and balances the number of virtual disks each controller owns. Alternatively, you can select the owner yourself.

In a dual controller configuration, when a controller fails, the partner controller assumes temporary operation of the failed controller's virtual disks and assigned spares.

The following table specifies the number of disk drives that each RAID level supports. For more information about RAID levels, see Appendix B.

Table 3-1 Number of Disk Drives Required for Each RAID Level

RAID Level	Minimum and Maximum Number of Disk Drives
Non-RAID	1
0	2–16
1	2 (To create a mirror with more than two drives, use RAID 10.)
3	3–16
5	3–16
6	4–16
10	4–16
50	6–32

Note – RAID 50 must have the same number of drives in each sub-vdisk, so the total number of drives is a multiple of the number of drives in each sub-vdisk. Each sub-vdisk can have 3 to 16 drives. There are many possible configurations of RAID 50 virtual disks, depending on the number of drives in each sub-vdisk and the number of stripes.

When you create a virtual disk you can also create volumes within it. A volume is a logical subdivision of a virtual disk, and can be mapped to host ports for access by data hosts. This type of volume is not the same as a volume you create with your operating system or with third-party tools.

You can create a virtual disk automatically or manually:

- Automatic Virtual Disk Creation (Policy-based) creates a virtual disk based on minimal information. See “Creating a Virtual Disk Automatically” on page 61.
- Manual Virtual Disk Creation (Detail-based) creates a virtual disk based on parameters you select, which provides greater control over the configuration than Automatic Virtual Disk Creation. See “Creating a Virtual Disk Manually” on page 63.

Creating a Virtual Disk Automatically

If your system has only one type of disk drive inserted (SAS or SATA), you can create a virtual disk “automatically” by using the Automatic Virtual Disk Creation option. This option creates a virtual disk based on minimal information.

To create a virtual disk automatically:

1. Select Manage > Virtual Disk Config > Create A Vdisk.
2. Select Automatic Virtual Disk Creation.
3. Type a name for the virtual disk.

   The name is case-sensitive and can include 17 characters. Allowed characters include letters, numbers, hyphens, underscores, and spaces.
4. Set Fault Tolerance Level to one of the following options:
   - None – None creates a non-RAID virtual disk if a single disk drive is required or a RAID 0 virtual disk if multiple drives are required; either will stop working if a drive fails.
   - Medium – Creates a RAID 5 virtual disk that can tolerate and recover from a failed disk. This is the default.
   - High – Creates a RAID 50 virtual disk that can tolerate and recover from a failed disk.

   Use of fault tolerance consumes some of the data capacity. To the right of this field, Largest Possible Virtual Disk estimates the largest virtual disk size that can be created for the selected fault-tolerance level, and depends on the number and size of available drives in the system.
5. Set Minimum Size Of Virtual disk to the amount of available space to use for all volumes on the new virtual disk.

   This value is rounded to the nearest Gbyte and is shown to the right of this field as Targeted Virtual Disk Size. Because RAIDar allocates entire drives to virtual disks, the resulting virtual disk is typically larger than the requested size; capacity beyond that allocated to volumes is designated as free space. For example, if you set this value to 600 Gbyte and your system has 500-Gbyte drives, the resulting virtual disk will be approximately 1000 Gbyte, including approximately 400 Gbyte of free space.

6. Set Number Of Volumes to the number of individual volumes the virtual disk is to be divided into.

   You can create a virtual disk that has one volume or multiple volumes. Single-volume virtual disks work well in environments that need one large, fault-tolerant storage space for data on one server. A large database accessed by users on a single server that is used only for that application is an example. Multiple-volume virtual disks work well when you have very large disk drives and you want to make the most efficient use of disk space for fault tolerance (parity and spares).

   The Size Of Each Volume field shows the Targeted Virtual Disk Size divided by the number of volumes. Volumes created are approximately the size requested; they may be a few percent larger than the requested size. Unused capacity in the virtual disk is designated as free space.

   Note – If you have disk drives of different sizes, the calculations are based on the smaller drives. This can result in virtual disks that have larger actual size than you requested and space on the larger drives that will be unused. To avoid these problems, you can use Manual Virtual Disk Creation to select disk drives to include in a virtual disk.

   ---

7. Click Create New Virtual Disk.

   A new page shows the progress of virtual disk initialization. See “Virtual Disk Initialization” on page 66.

Creating a Virtual Disk Manually

To create a virtual disk manually:

1. Select Manage > Virtual Disk Config > Create A Vdisk.
2. Select Manual Virtual Disk Creation.
3. Type a name for the virtual disk.

   The name is case-sensitive and can include 17 characters. Allowed characters include letters, numbers, hyphens, underscores, and spaces.
4. Set Virtual Disk RAID Level to one of the following options:
   - RAID 0 - Disk Striping
   - RAID 1 - Disk Mirroring, 2 Disks only
   - RAID 3 - Parity RAID, 1 Parity Disk
   - RAID 5 - Parity RAID, Parity Distributed
   - RAID 6 - Double-parity RAID, Parity Distributed
   - Non-RAID
   - RAID 10 - Data Striped Over Mirrors
   - RAID 50 - Data Striped Across RAID 5
5. (Optional) As an Advanced user, you can set the initialization type:

   a. Click Advanced Virtual Disk Creation Options.

   b. Set Initialization Type to one of the following options:
      - Online – Enables you to use the virtual disk immediately after creating it while it is initializing. Because Online uses the verify method to create the virtual disk, it takes longer to complete initializing than Offline. This option is the default.
      - Offline – You must wait for the virtual disk initialization process to finish before using the virtual disk; however, Offline takes less time to complete initializing than Online. You can only create a virtual disk with one volume using this option. You can add volumes when the virtual disk initialization is complete.
6. Click Create New Virtual Disk.
7. Select the drives to use in the virtual disk.

   Only available drives are selectable. Available drives are neither in a virtual disk nor assigned as a spare.

   The minimum and maximum number of drives that you can select when creating a virtual disk are shown in Table 3-1.

In a multi-enclosure system, for certain RAID levels you can select drives in a way that provides some protection against enclosure failure:

- RAID 1, 3, 5, or 6 – Select each drive from a different enclosure.
- RAID 10 – Select the first half of the drives from one enclosure and the second half from another enclosure. The first set is assigned to one mirror group and the second to other mirror group, which limits the effect of an enclosure failure to one mirror.
- RAID 50 – Drives that you select consecutively are assigned to different sub-vdisks in the virtual disk. Therefore, you can force the drives in each sub-vdisk to be selected from different enclosures, improving the protection of each sub-vdisk from an enclosure failure.

8. (Optional) Calculate whether the formatted virtual disk will have the capacity you want:

   a. Click Calculate Virtual Disk Size.

      The results of the calculation are displayed.

   b. Click OK.

      If the capacity is insufficient for your application, change your drive selections and repeat this step.

9. (Optional) For RAID 1, 3, 5, 6, 10, or 50, select spare drives for this virtual disk:

   a. Set the add dedicated spares option to Yes.

      The default is No.

Note – If you want to designate spares that can be used by any virtual disk, create the virtual disk without adding spares, and then create global spares.

---

   b. Click Continue.

      The Select Spare Drives page is displayed.

   c. Select the check box of each drive to use as a spare in the virtual disk.

      A drive has a check box if the drive is available and is no smaller than the smallest disk drive in the virtual disk.

      You can add four spares to a virtual disk.

   d. Click Continue.

      The Configure Volumes For Virtual Disk page is displayed and summarizes your selections.

10. (Optional) Set How Many Volumes to the number of standard volumes you want in your virtual disk.

    You can create a virtual disk that has no volumes (the default), one volume or multiple volumes. One volume works well in environments that need one large, fault-tolerant storage space for data for one host, such as a large database accessed by users on a single host. Multiple volumes work well when you have very large disk drives and you want to make the most efficient use of disk space for fault tolerance (parity and spares). If you choose to create no volumes, you can later add standard volumes or volumes of other types.

11. If you specified to create one or more volumes, set the following volume options:
    - Create Volumes Of Equal Size? – (Online initialization only) The default is Yes. If you select No, you can type the size of each volume on the next page.
    - Expose Volumes To All Hosts? – The default is No, which sets the volumes' LUN to None so hosts cannot access the volumes until you manually map them. If you select Yes, the volumes are automatically mapped to all connected hosts with read-write access on all controller host ports, and the Automatically Assign LUNs option is enabled.
    - Automatically Assign LUNs? – The default is Yes. If you select No, you can type the LUN for each volume on the next page.
    - Would You Like To Name Your Volumes? – (Online initialization only) The default is No. If you select Yes, you can type a name for each volume on the next page.

12. (Optional) Set the following advanced options for the virtual disk:
    - Virtual Disk Chunk Size – The chunk size is the amount of contiguous data that is written to a virtual disk member before moving to the next member of the virtual disk. The allowed values are 16K, 32K, or 64K (Kbyte). The default is 64K. If you are using the virtual disk for a database with very small records, you might want to use a smaller chunk size. Check the data chunk size that your application is sending to the virtual disk, then set the virtual disk chunk size to best match that of your application.
    - Preferred Virtual Disk Owner – Select the controller that should own the virtual disk. If you do not make a selection here, RAIDar automatically selects the owner and balances the number of virtual disks each controller owns.

    Click Continue to return to the previous window and continue the virtual disk creation process.

13. Click Create Virtual Disk.

    The system creates the virtual disk and shows the next page in the process.

    - If you accepted the default volume options, the final page shows the progress of virtual disk initialization. Proceed to “Virtual Disk Initialization” on page 66.
    - If you changed any of the volume options, an Add Volumes To Virtual Disk page is displayed showing information based on your selections from the previous page. Continue with Step 14.

14. If any of the following options are displayed, set them appropriately:
    - Volume Size – Shows the volume size based on evenly-sized volumes. Type the size you want for each volume in Mbyte. To calculate the total size of the volumes based on the values you typed, click Calculate The Total Size.
    - Volume LUN ID – Shows the default LUN. Select the LUN you prefer.
    - Volume Name – Shows the default name. Type a new name. The name is case-sensitive and can include 20 characters. Allowed characters include letters, numbers, hyphens, underscores, and spaces.

15. Click Add Volumes.

    The final page shows the progress of virtual disk initialization.

Virtual Disk Initialization

The initialization process takes from several minutes to more than an hour depending on the RAID level (RAID 0 and RAID 1 are the fastest), virtual disk size, drive speed, and other processes running on the system.

If the virtual disk is initializing online, you can start using it immediately. If the virtual disk is initializing offline, you must wait for initialization to complete before using the virtual disk.

If you must change the virtual disk's configuration or use of disk drives before initialization is complete, you can stop initialization.

Managing Virtual Disks

RAIDar enables you to manage virtual disks in a variety of ways. You can:

- View the status of virtual disks and disk drives
- Expand virtual disk capacity
- Dequarantine a virtual disk
- Verify a virtual disk
- Change a virtual disk's owner
- Change a virtual disk's name
- Delete a virtual disk

For information about reconstructing a failed virtual disk, see the *Troubleshooting Guide*.

Viewing Virtual Disk and Disk Drive Status Information

You can view status information for a virtual disk and for disk drives associated with a virtual disk.

Virtual Disk Status

To view information about a virtual disk:

1. Select Manage > Virtual Disk Config > Vdisk Configuration > Vdisk Status.

   For each virtual disk, the virtual disk panel shows a status icon; the name, RAID level, size, number of disk drives, and number of volumes; and utility status, if any.

2. Select a virtual disk.

   The Virtual Disk Status panel shows the following information:

   - RAID Level – Either RAID 0, 1, 3, 5, 6, 10, 50, or Non-RAID.
   - Virtual Disk Size – Virtual disk size in Gbyte.
   - Virtual Disk Status – Either Online, Offline, Critical, or Fault Tolerant.
   - Number Of Drives – Number of drives in the virtual disk when fault tolerant. For example, if a three-drive RAID 5 virtual disk loses a drive, this value remains 3.
   - Spare Drives – Number of spares assigned to this virtual disk.
   - Number Of Volumes – Number of volumes in the virtual disk.
   - Virtual Disk Serial Number – Unique number assigned by the owning controller.

- Virtual Disk Owner – Controller that owns the virtual disk.
- Chunk Size – Amount of contiguous data in Kbyte that is written to a virtual disk member before moving to the next member of the virtual disk.
- Date Created – Date when the virtual disk was created.
- Utility – Name of any utility running on the virtual disk, or None. The utility status is shown in the virtual disk panel.

The Enclosure View panel shows a graphical representation of disk drives by enclosure. On this page only, drives in any virtual disk except the selected one are gray. For more information, see "Disk Drives by Enclosure" on page 150. If this panel does not display, see "Enclosure View is Unavailable" on page 151.

Disk Drive Status

To view information about the drives in a virtual disk:

1. Select Manage > Virtual Disk Config > Vdisk Configuration > Disk Drive Status.

   For each virtual disk, the virtual disk panel shows a status icon; the name, RAID level, size, number of disk drives, and number of volumes; and utility status, if any.

2. Select a virtual disk.

   The Virtual Disk Drive List panel shows the following information about each drive in the virtual disk:

   - Status – Up if operational or Down if failed
   - Size – Drive size in Gbyte
   - Manufacturer – Drive manufacturer
   - Model – Drive model number
   - Revision – Drive firmware revision number
   - Node WWN – Drive node World Wide Name
   - Serial Number – Drive serial number
   - Enclosure:Slot – Enclosure number and slot number containing the drive
   - Enclosure – Name of the enclosure containing the drive

   The Dedicated Spares For Selected Virtual Disk panel shows the same information about each spare assigned to the virtual disk. If no spares are assigned, the panel is not displayed.

Expanding Virtual Disk Capacity

You can expand the capacity of a virtual disk by adding drives to it. Because virtual disk expansion does not require I/O to be quiesced, the virtual disk can continue to be used while the Expand utility runs. Expanding a virtual disk adds free space after the space used by existing volumes. You can then create or expand a volume to use the free space. You can expand only one virtual disk at a time.

The RAID level determines how the virtual disk can be expanded and the maximum number of drives the virtual disk can have, as shown in the following table.

Table 3-2 Virtual Disk Expansion by RAID Level

RAID Level	Expansion Capability	Maximum Drives
Non-RAID	Cannot expand.	1
0, 3, 5, 6	You can add 1–4 drives at a time.	16
1	Cannot expand.	2
10	You can add 2 or 4 drives at a time.	16
50	You can expand the virtual disk, one sub-vdisk at a time. The added sub-vdisk must contain the same number of drives as each of the existing sub-vdisks.	32

Note – Expansion can take hours or days to complete, depending on the virtual disk RAID level and size, drive speed, utility priority, and other processes running on the storage system. You can stop an expansion only by deleting the virtual disk. Before starting the expansion, make sure to back up the data so that if you need to delete the virtual disk, you can move the data into a new, larger virtual disk.

To expand a virtual disk:

1. Select Manage > Virtual Disk Config > Vdisk Configuration > Expand Virtual Disk.

   For each virtual disk, the virtual disk panel shows a status icon; the name, RAID level, size, number of disk drives, and number of volumes; and utility status, if any.

2. Select the virtual disk to expand.
3. Select available drives to add to the virtual disk.

4. Click Expand Virtual Disk.

   Expansion begins and the percentage completed is shown. You can perform other functions during the expansion. You can view the status of the expansion on the Vdisk Utility Progress page or on any page that shows virtual disk icons.

Checking the Progress of a Utility

To check the status of any running virtual disk utilities:

- Select Manage > Virtual Disk Config > Vdisk Utility Progress.

  For each virtual disk, the virtual disk panel shows a status icon; the name, RAID level, size, number of disk drives, and number of volumes; and utility status, if any.

  For each virtual disk where a utility is running, a Utility Running For Virtual Disk panel specifies its status.

Note – To stop the Initialize or Verify utility, go to the Abort A Vdisk Utility page. To stop background scrub of virtual disks, go to the General Config > System Configuration page. You cannot stop the Expand or Reconstruct utility unless you delete the virtual disk.

Dequarantining a Virtual Disk

The quarantine icon  indicates that a previously fault-tolerant virtual disk is quarantined because not all of its drives were detected after a restart or rescan. Quarantine isolates the virtual disk from host access, and prevents the storage system from making the virtual disk critical and starting reconstruction when drives are "missing" for these reasons:

- Slow to spin up after system power-up
- Not properly seated in their slots
- In an powered-off enclosure
- Inserted from a different system and retain old metadata

The virtual disk can be fully recovered if the missing drives can be restored. Make sure that no drives have been inadvertently removed and that no cables have been unplugged. Sometimes not all drives in the virtual disk power up. Check that all enclosures have rebooted after a power failure. If these problems are found and then fixed, the virtual disk recovers and no data is lost.

The quarantined virtual disk's drives are "write locked," and the virtual disk is not available to hosts until the virtual disk is dequarantined. The system waits indefinitely for the missing drives. If the drives are found, the system automatically dequarantines the virtual disk. If the drives are never found because they have been removed or have failed, you must dequarantine the virtual disk manually.

If the missing drives cannot be restored (for example, a failed drive), you can use dequarantine to restore operation in some cases. If you dequarantine a fault-tolerant virtual disk that is not missing too many drives, its status changes to critical. Then, if spares of the appropriate size are available, reconstruction begins.

Caution – If the virtual disk does not have enough drives to continue operation, when a dequarantine is done, the virtual disk goes offline and its data cannot be recovered.

To dequarantine a virtual disk:

1. Select Manage > Utilities > Recovery Utilities > Vdisk Quarantine.

   For each virtual disk, the virtual disk panel shows a status icon; the name, RAID level, size, number of disk drives, and number of volumes; and utility status, if any.

2. Select the virtual disk to dequarantine.

3. Click Dequarantine Selected Virtual Disk.

Verifying a Virtual Disk

When you suspect that a redundant virtual disk has a problem, you can verify its data. For example, if the system was operating outside the normal temperature range for any length of time. The RAID level determines the Verify utility's behavior:

- For RAID 3, 5, 6, and 50, the Verify utility verifies all parity blocks in the virtual disk and corrects any bad parity.
- For RAID 1 and 10, the Verify utility compares the primary and secondary drives. If a mismatch occurs, the primary is copied to the secondary.

The verification process ensures that the redundancy data in the virtual disk is consistent with the user data in the virtual disk. If an inconsistency is found, the redundancy data is updated to reflect the current state of the user data. The number of inconsistencies found is noted in the "Virtual disk verification complete" event in the event log.

The number of virtual disk verifications you can initiate is determined by the current load on your controllers. If an error is displayed when you try to verify a virtual disk and multiple utilities running, wait until those utilities have completed and try again.

Starting Virtual Disk Verification

To verify a virtual disk:

1. Select Manage > Virtual Disk Config > Vdisk Configuration > Verify Virtual Disk.

   For each virtual disk, the virtual disk panel shows a status icon; the name, RAID level, size, number of disk drives, and number of volumes; and utility status, if any.

2. Select the virtual disk to verify.

3. Click Verify & Update Virtual Disk Parity.

   Verification begins and the percentage of verification completed is displayed. You can continue to use the virtual disk during verification. To check the progress of the verification, select Manage > Virtual Disk Config > Vdisk Utility Progress.

Stopping Virtual Disk Verification

You can stop the virtual disk verification process at any time. If you stop verification, you cannot resume; you must restart the verification from the beginning.

1. Select Manage > Virtual Disk Config > Abort A Vdisk Utility.

   For each virtual disk, the virtual disk panel shows a status icon; the name, RAID level, size, number of disk drives, and number of volumes; and utility status, if any.

2. Select the virtual disk to stop verifying.

3. Click Abort Verify.

   When processing is complete, the virtual disk icon changes to show that no utility is running.

Changing Virtual Disk Ownership

Each virtual disk is associated with one controller. RAIDar balances the number of virtual disks each controller owns.

When a controller fails, the partner controller assumes temporary ownership of the failed controller's virtual disks and resources. If the system uses a fault-tolerant cabling configuration, both controller's LUNs will be accessible through the partner.

Typically, you should not need to change virtual disk ownership.

Caution – When you change the ownership of a virtual disk whose volumes are mapped to hosts, the assigned LUNs become invalid and hosts lose access to the volumes. After changing ownership, you must reassign the LUNs and, depending on the host operating system, either rescan or restart to detect the LUN changes.

To change the owner of a virtual disk:

1. Select Manage > Virtual Disk Config > Vdisk Configuration > Change Vdisk Owner.

   For each virtual disk, the virtual disk panel shows a status icon; the name, RAID level, size, number of disk drives, and number of volumes; and utility status, if any.

2. Select the virtual disk to change ownership of.

   The Change Virtual Disk Owner panel shows the current owner.

3. Click Change Virtual Disk Owner To RAID Controller *X*, where *X* is whichever controller does not currently own the virtual disk.

4. Assign a new LUN to each volume (see "Managing Volume Mappings" on page 90).

Changing a Virtual Disk Name

To change the name of a virtual disk:

1. Select Manage > Virtual Disk Config > Vdisk Configuration > Change Vdisk Name.

   For each virtual disk, the virtual disk panel shows a status icon; the name, RAID level, size, number of disk drives, and number of volumes; and utility status, if any.

2. In the Change Virtual Disk Name field, type a new name.

   The name is case-sensitive and can include 17 characters. Allowed characters include letters, numbers, hyphens, underscores, and spaces.

3. Click Change Virtual Disk Name.

Deleting a Virtual Disk

You can delete a virtual disk when you no longer need the virtual disk or you need its disks for another use. You do not need to stop any utilities running on the virtual disk.

Caution – Deleting a virtual disk deletes all volumes and data contained in the virtual disk.

To delete a virtual disk:

1. Select Manage > Virtual Disk Config > Delete A Vdisk.

   For each virtual disk, the virtual disk panel shows a status icon; the name, RAID level, size, number of disk drives, and number of volumes; and utility status, if any.

2. Select the virtual disk to delete.

3. Click Delete This Virtual Disk.

Managing Spares

Controllers in your system automatically reconstruct redundant (fault-tolerant) virtual disks (RAID 1, 3, 5, 6, 10, and 50) if a virtual disk becomes critical and a properly sized spare disk is available. A virtual disk becomes critical when one or more of its disks fails.

There are three types of spares:

- A *vdisk spare* is an available drive that is assigned to a specific virtual disk.
- A *global spare* is an available drive that can act as a spare for any failed drive in any redundant virtual disk. Global spares are available to any redundant virtual disk in the system. If a drive in a virtual disk fails, the controller can use a global spare to reconstruct the critical virtual disk.
- A *dynamic spare* is a properly sized available drive that is automatically assigned by the system.

When a disk fails, the system looks for a vdisk spare first. If it does not find a properly sized vdisk spare, it looks for a global spare. If it does not find a properly sized global spare and the dynamic spares option is enabled, it takes any properly sized available drive. If no properly sized spares are available, reconstruction must be started manually.

For more information, see “Managing Dynamic Spares” on page 77, “Managing Vdisk Spares and Global Spares” on page 78, or the topic about reconstructing a virtual disk in the *Troubleshooting Guide*.

Managing Dynamic Spares

The dynamic spares feature lets you use all of your disk drives in redundant virtual disks without designating one as a spare. With dynamic spares enabled, if a drive fails and you replace it with a properly sized drive, the storage system rescans the bus, finds the new drive, automatically designates it a spare, and starts reconstructing the virtual disk. A properly sized drive is one whose capacity is equal to or greater than the smallest drive in the virtual disk.

If a vdisk spare, global spare, or properly sized available drive is already present, the dynamic spares feature uses that drive to start the reconstruction and the replacement drive can be used for another purpose.

To configure dynamic spares:

1. Select Manage > General Config > System Configuration.
2. Set Dynamic Spare Configuration to Enabled.
3. Click Change System Configuration.

   When Dynamic Spare Configuration is enabled, the Dynamic Spare Rescan Rate option is displayed. Use the default rescan rate.
4. Click Change System Configuration.

Managing Vdisk Spares and Global Spares

This section describes how to designate available drives as spares for use by one virtual disk or by any virtual disk. It also describes how to return spares to the pool of available drives.

Adding Vdisk Spares

You can add a maximum of four available drives to a redundant virtual disk (RAID 1, 3, 5, 6, 10, and 50) for use as spares. If a drive in the virtual disk fails, one of these *vdisk spares* is automatically used to reconstruct the virtual disk. You cannot add a spare with insufficient capacity to replace the smallest drive in the virtual disk. Vdisk spares are also called *dedicated spares*.

The controller automatically uses the vdisk spare for reconstruction of the critical virtual disk to which it belongs. The virtual disk remains in Critical status until the parity or mirror data is completely written to the spare, at which time the virtual disk returns to Fault Tolerant status. For RAID 50 virtual disks, if more than one sub-vdisk becomes critical, reconstruction and use of vdisk spares occur in the order sub-vdisks are numbered.

Although using a vdisk spare is the most secure way to provide spares for your virtual disks, it is also expensive to keep a spare assigned to each virtual disk. An alternative method is to enable dynamic spares or to assign one or more unused drives as global spares.

To add spares to a virtual disk:

1. Select Manage > Virtual Disk Config > Vdisk Configuration > Add Vdisk Spares.

   For each virtual disk, the virtual disk panel shows a status icon; the name, RAID level, size, number of disk drives, and number of volumes; and utility status, if any.

2. Select a virtual disk.

3. In the Select Drives To Be Vdisk Spares panel, select drives to be spares for the selected virtual disk. Only appropriate drives are selectable.

4. Click Add Vdisk Spares.

   A processing message is displayed.

Deleting Vdisk Spares

You can delete vdisk spares from a virtual disk at any time. To delete vdisk spares:

1. Select Manage > Virtual Disk Config > Vdisk Configuration > Delete Vdisk Spares.

   For each virtual disk, the virtual disk panel shows a status icon; the name, RAID level, size, number of disk drives, and number of volumes; and utility status, if any.

2. Select a virtual disk.

3. In the Select Spare Drives To Delete panel, select the spares to delete.

   Only spares in the selected virtual disk are selectable.

4. Click Delete Vdisk Spares.

   When processing is complete, enclosure view shows the drives as available.

Adding Global Spares

You can designate a maximum of eight *global spares* for the system. If a disk in any redundant virtual disk (RAID 1, 3, 5, 6, 10, and 50) fails, a global spare is automatically used to reconstruct the virtual disk. At least one virtual disk must exist before you can add a global spare. You cannot add a spare that has insufficient capacity to replace the smallest drive in an existing virtual disk.

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

To add global spares:

1. Select Manage > Virtual Disk Config > Global Spare Menu > Add Global Spares.

2. Select drives to designate as global spares.

   Only appropriate drives are selectable.

3. Click Add Global Spares.

   When processing is complete, the drive's icon changes to gray with a "G" in the enclosure view.

Deleting Global Spares

You can delete global spares at any time. To delete global spares:

1. Select Manage > Virtual Disk Config > Global Spare Menu > Delete Global Spares.
2. Select the global spares to delete.
   Only global spares are selectable.
3. Click Delete Global Spares.
   When processing is complete, enclosure view shows the drives as available.

Displaying Global Spares

To display global spares:

- Select Manage > Virtual Disk Config > Global Spare Menu > Show Global Spares.

Drives whose icons are gray with a “G” are global spares.

Managing Volumes

RAIDar lets you manage volumes in a variety of ways. You can:

- Add a volume
- Expand a volume
- View volume status information
- Change a volume name
- Control access to volumes by mapping volumes to hosts
- Change a volume’s read-ahead cache settings
- Enable or disable a volume’s write-back cache
- Delete a volume

For information about master volumes, snap-pool volumes, and snapshots, see “Using Snapshot Services” on page 98. For information about copying volumes, see “Using Volume Copy Services” on page 117.

Understanding Volumes

A volume is a logical subdivision of a virtual disk. Using RAIDar you can add, expand, rename, delete volumes, and map them to data hosts. This type of volume provides the storage for a file system partition you create with your operating system or third-party tools. A dual-controller system supports a maximum of 256 volumes.

A virtual disk can have one or more volumes. Using multiple volumes lets you create one very large virtual disk making efficient use of your disk drives. For example, you could create one very large RAID 5 virtual disk and assign one vdisk spare to the virtual disk. This minimizes the amount of disk space allocated to parity and spares compared to the space required if you created five or six smaller RAID 5 virtual disks.

You can give each volume a name. Assign names that indicate how the volumes are to be used. For example, if the first volume will be used to store your customer database, give it a name such as: `cust_database`.

When you create a virtual disk, you can specify the number of volumes you want and their sizes. If the total size of the volumes equals the size of the virtual disk, you will not have any free space, as shown in Figure 3-1. In this example, the volumes in VirtualDisk-1 are equal in size and use all of the virtual disk's space.

You can also create fewer volumes that do not equal the virtual disk's size. This leaves free space in which you can add or expand volumes later as shown by VirtualDisk-2 in the following figure.

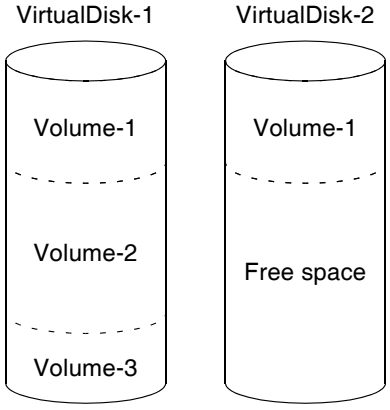

Figure 3-1 A Virtual Disk With Multiple Volumes and a Virtual Disk With One Volume and Free Space

After expanding a virtual disk, you can either add a volume or expand a volume to use the new free space. You can also delete one or more volumes and expand a volume into the space.

For information about mapping volumes and assigning LUNs (logical unit numbers), see “Understanding Volume Mapping” on page 85.

Adding a Volume

You must have free space in a virtual disk before you can add a volume. You can create free space by deleting a volume (see “Deleting a Volume” on page 97) or by expanding the virtual disk (see “Expanding Virtual Disk Capacity” on page 69). You can add volumes to a virtual disk until you use all of the free space.

To add a volume:

1. Select Manage > Volume Management > Volume Menu > Add Volume.

   For each virtual disk, the virtual disk panel shows a status icon; the name, RAID level, size, number of disk drives, and number of volumes; and utility status, if any.

2. Select a virtual disk.

   The selected virtual disk’s volume names, sizes, and types are displayed, and the amount of free space.

3. Type a size in increments of 1 Mbyte for the new volume.

4. (Optional) Change the name for the new volume.

   The default is *vdisk-name*_V*number*. For example, MyVdisk_V1.

   The name is case-sensitive and can include 20 characters. Allowed characters include letters, numbers, hyphens, underscores, and spaces.

5. (Optional) Change the LUN setting.
   - NONE – The volume is not accessible by connected hosts. This setting is the default. You can map the volume to hosts later; see “Managing Volume Mappings” on page 90.
   - 0–127 – The volume is accessible with this LUN by all connected hosts.

6. Click Add Volume.

   When processing is complete, the new volume is displayed in the Volume Menu panel.

Expanding a Volume

You can expand a standard volume or a snap pool if the virtual disk has free space and sufficient resources. Because volume expansion does not require I/O to be quiesced, the volume can continue to be used while it is expanded.

To expand a volume:

1. Select Manage > Volume Management > Volume Menu > Expand Volume.

   For each virtual disk, the virtual disk panel shows a status icon; the name, RAID level, size, number of disk drives, and number of volumes; and utility status, if any.

2. Select a virtual disk.

   The selected virtual disk's volume names, sizes, and types are displayed, and the amount of free space.

3. Select a standard volume or snap pool to expand.

4. Enter the amount of free space in increments of 1 Mbyte to add to the volume.

5. Click Expand Volume.

   When processing is complete, the new size is displayed in the Volume Menu panel.

Viewing Volume Status Information

Volume status information is available from the following pages in RAIDar:

- Monitor > Status > Vdisk Status. Includes volume information for the selected virtual disk. See “Virtual Disk Status” on page 144.
- Manage > Volume Management > Volume Menu > Volume Status. Includes more detailed volume information for the selected virtual disk.

On Virtual Disk Config and Volume Management pages, the virtual disk panel shows an icon for each virtual disk with information about the virtual disk below it. See “Displaying Status Information” on page 143 for a description of the virtual disk icons. Click a virtual disk icon to display panels with additional information. The information that is available varies, depending on the page.

On Volume Management pages, the Volume Menu panel shows a color-coded “map” of the space used by each volume in the selected virtual disk. The color codes are:

- Gray – Free space
- Green – Standard volume
- Blue – Snap pool
- Orange – Master volume
- Yellow – Snapshot

The panel also shows a table with information about each volume and the amount of free space. The information that is shown varies, depending on the page.

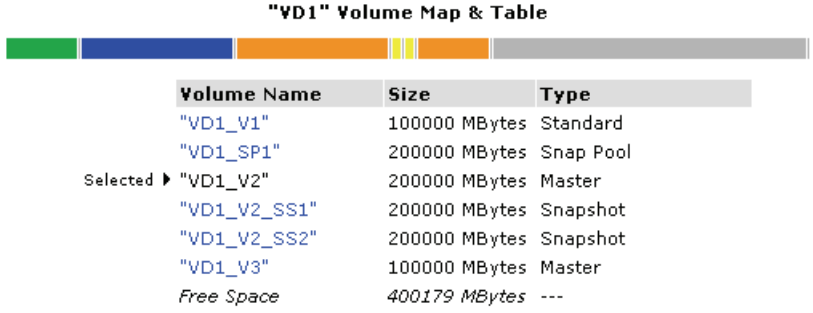

Figure 3-2 Volume Map and Table Example

Note – For an explanation of sizes represented by various units, see “Size Representations in RAIDar” on page 23.

Changing a Volume Name

You can change the name of a volume. This does not affect the target ID or LUN values of the volume.

To change a volume name:

1. Select Manage > Volume Management > Volume Menu > Change Volume Name.

   For each virtual disk, the virtual disk panel shows a status icon; the name, RAID level, size, number of disk drives, and number of volumes; and utility status, if any.

2. Select a virtual disk.

   The selected virtual disk's volume names, sizes, and types are displayed, and the amount of free space.

3. Select the volume to rename.

4. In the Change Volume Name field, type a new name.

   The name is case-sensitive and can include 20 characters. Allowed characters include letters, numbers, hyphens, underscores, and spaces.

5. Click Change Volume Name.

   When processing is complete, the new name is displayed in the Volume Menu panel.

Understanding Volume Mapping

You can enable data-host access to specific volumes and control the type of access each data host has to each volume. You do this by mapping a volume to a data-host port through one or more controller host ports, setting access privileges for each controller host port, and assigning a LUN (logical unit number) to the mapping.

A LUN identifies a volume to a host through a specific access path. This enables a volume to be accessed with a different LUN by different hosts, or a specific LUN to be used by different hosts to access different volumes.

For a given volume you can configure the same access for all hosts, or configure different access for specific hosts than for all other hosts. The access privilege for a controller host port can be set to read-write, read-only, or none (no access).

For the 69501 when host port interconnects are enabled between controller A and B, the access settings for controller A port 0 also apply to controller B port 1, and access settings for controller A port 1 also apply to controller B port 0.

The following figure shows an example of using volume mapping to control access to volumes.

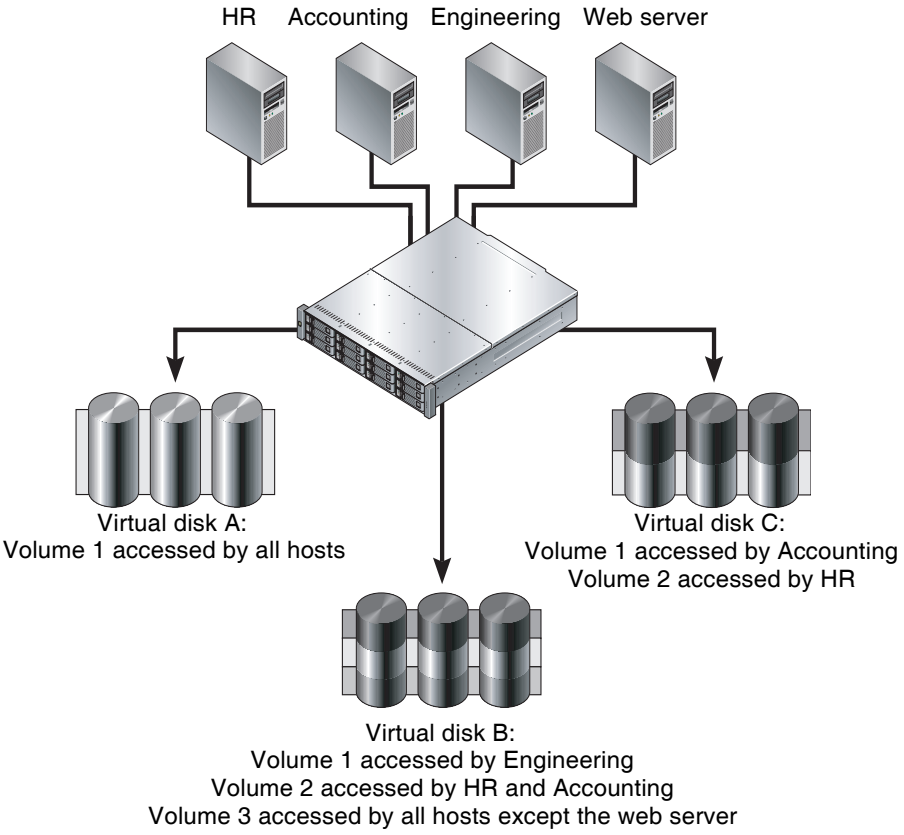

Figure 3-3 Volume Mapping Example

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

In a dual-controller system, if one controller fails the partner controller takes temporary ownership of the failed controller's volumes and other resources. For information about how controllers present mapped volumes in different configurations during active-active operation and failover, see Appendix C.

For information about managing the list of data-host ports that can be used for volume mapping, see "Managing the Global Host Port List" on page 87.

For the procedures to manage volume mappings, see "Managing Volume Mappings" on page 90.

Managing the Global Host Port List

The *global host port list* is a list of data-host ports that can be used for volume mapping. You can assign a nickname to a port to make it easily recognizable.

The list is automatically populated with the port WWNs (FC) or IP addresses (iSCSI) of hosts that have sent an `inquiry` or a `report luns` command to the system. Hosts will typically do this when they boot up or rescan for devices. When the command from the host occurs, the system saves the host port information; however, the information is not retained after you restart the system unless you have assigned the port a nickname. You can also add ports manually to the list. On an FC system, 63 nicknames can be assigned; on an iSCSI system, 56 can be assigned.

Note – Before you can manually add a port to the list you must know the port's WWN (FC) or IP address (iSCSI).

Managing the Global Host Port List on an FC System

On the Manage Host List page you can:

- Display the list of known host ports
- Set or change nicknames for host ports
- Add host ports to the list
- Delete host ports and nicknames from the list

The following figure shows the host port list with two example entries:

Figure 3-4 FC Host Port List With Two Example Entries

To display the global host port list:

- Select Manage > Volume Management > Volume Mapping > Manage Host List.

  The Current Global Host Port List panel shows the port WWN, controller ID and port number, manufacturer, and nickname (if any) for each data-host port. The port associated with the host that most recently scanned for devices is first in the list.

To set or change a port nickname:

1. Type a new nickname in the port's Nickname field.

   The name is case-sensitive and can include 15 characters. Allowed characters include letters, numbers, hyphens, underscores, and spaces.

2. Click Update.

To add a port:

1. In the Add Port To Global Host Port List panel, type the port WWN and a nickname.

   The name is case-sensitive and can include 15 characters. Allowed characters include letters, numbers, hyphens, underscores, and spaces.

2. Click Add New Port.

   If you add a WWN and nickname and the WWN is already in the list, the new nickname replaces the old. If you add a WWN and nickname and the nickname is the same as an existing one, the attempt is rejected.

To delete either a manually added port or the nickname of an automatically added port:

- In the port's row, click Delete.

  If the host had scanned for devices since you last restarted the system, you need to restart the system to complete the deletion.

Managing the Global Host Port List on an iSCSI System

On the Manage Host List page you can:

- Display the list of known host ports
- Set or change nicknames for host ports
- Add host ports to the list
- Delete host ports and nicknames from the list

The following figure shows the host port list with two example entries:

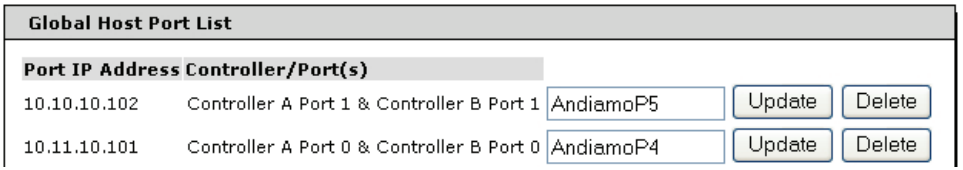

Figure 3-5 iSCSI Host Port List With Two Example Entries

To display the global host port list:

- Select Manage > Volume Management > Volume Mapping > Manage Host List.

  The Current Global Host Port List panel shows the port IP address, controller ID and port number, and nickname (if any) for each data-host port. The port associated with the host that most recently scanned for devices is first in the list.

To set or change a port nickname:

1. Type a new nickname in the port's Nickname field.

   The name is case-sensitive and can include 15 characters. Allowed characters include letters, numbers, hyphens, underscores, and spaces.

2. Click Update.

To add a port:

1. In the Add Port To Global Host Port List panel, type the port IP address and a nickname.

   The name is case-sensitive and can include 15 characters. Allowed characters include letters, numbers, hyphens, underscores, and spaces.

2. Click Add New Port.

   If you add an IP address and nickname and the IP address is already in the list, the new nickname replaces the old. If you add an IP address and nickname and the nickname is the same as an existing one, the attempt is rejected.

   To delete either a manually added port or the nickname of an automatically added port:

- In the port's row, click Delete.

  If the host had scanned for devices since you last restarted the system, you need to restart the system to complete the deletion.

Managing Volume Mappings

In the Map Hosts To Volume page you can add, change, or delete host access to the selected volume.

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 a mapped volume from a host system before changing the mapping's LUN.

To manage host-to-volume mappings:

1. Select Manage > Volume Management > Volume Mapping > Map Hosts To Volume.

   For each virtual disk, the virtual disk panel shows a status icon; the name, RAID level, size, number of disk drives, and number of volumes; and utility status, if any.

2. Select a virtual disk.

   The selected virtual disk's volume names, sizes, default LUNs, and types are displayed, and the amount of free space.

3. Select a volume.

   The Current Host-Volume Relationships panel shows which data-host ports have access to the selected volume. For the selected volume you might see the following mappings:

   - All Hosts – Shows the settings used by all data-host ports to access the volume. This entry is displayed only if no specific ports are mapped. If a specific port is mapped, All Hosts changes to All Other Hosts.
   - WWN value (FC) or IP address (iSCSI) – Shows the settings used by a data-host port to access the volume.

- All Other Hosts – Shows the access settings used by all data-host ports except by specifically mapped ports. This entry is displayed only if specific ports are mapped. If no specific port is mapped, All Other Hosts changes to All Hosts.

For each entry, the port identifier, the assigned LUN, and each controller host port's access privilege are shown. The access privilege for a controller host port can be read-write, read-only, or none (no access). A mapping cannot include both read-write and read-only access.

4. To add or change a mapping:

   a. In the Assign Host Access Privileges panel, select a host port identifier or All Other Hosts.

   b. Specify a LUN and port-access settings.

   c. Click Map It.

      When processing is complete, the page shows the new mapping.

5. To delete a mapping:

   a. In the Assign Host Access Privileges panel, select a host port identifier.

   b. Click Unmap It.

      When processing is complete, the mapping is removed from the page.

Changing a Volume's Read-Ahead Cache Settings

As an Advanced Manage user, you can optimize a volume for sequential reads or streaming data by changing its read-ahead cache settings. The read-ahead cache settings enable you to change the amount of data read in advance after two back-to-back reads are made. Read ahead is triggered by two back-to-back accesses to consecutive logical block address (LBA) ranges. Read ahead can be forward (that is, increasing LBAs) or reverse (that is, decreasing LBAs). Increasing the read-ahead cache size can greatly improve performance for multiple sequential read streams. However, increasing read-ahead size will likely decrease random read performance.

The default read-ahead size, which sets one chunk for the first access in a sequential read and one stripe for all subsequent accesses, works well for most users in most applications. The controllers treat volumes and mirrored virtual disks (RAID 1) internally as if they have a stripe size of 64 Kbyte, even though they are not striped.

Caution – Only change the read-ahead cache settings if you fully understand how your operating system, application, and FC HBA or iSCSI Ethernet adapter move data so that you can adjust the settings accordingly. Be prepared to monitor system performance using the virtual disk statistics and adjust read-ahead size until you find the optimal size for your application.

To change a volume's read-ahead cache settings:

1. Select Manage > Volume Management > Volume Menu > Read Ahead Cache.

   For each virtual disk, the virtual disk panel shows a status icon; the name, RAID level, size, number of disk drives, and number of volumes; and utility status, if any.

2. Select a virtual disk.

   The selected virtual disk's volume names, sizes, and read-ahead cache sizes are displayed, and the amount of free space.

3. Select the standard, snap-pool, or master volume whose cache settings you want to change.

4. Set Read Ahead Size to one of the following options:
   - Default – 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 virtual disk (the default is 64 KB). Non-RAID and RAID 1 virtual disks are considered to have a stripe size of 64 KB.
   - Disabled – 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 use the volume statistics read histogram to determine what size accesses the host is doing.
   - 64, 128, 256, or 512 KB; 1, 2, 4, 8, 16, or 32 MB – Sets the amount of data to read first, and the same amount is read for all read-ahead accesses.
   - Maximum – 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.

Note – Only use Maximum when disk drive latencies must be absorbed by cache.

---

5. Set Cache Optimization to one of the following options:
   - Standard – Works well for typical applications where accesses are a combination of sequential and random. This method is the default.
   - Super-Sequential – Slightly modifies the controller's standard read-ahead caching algorithm by enabling the controller to discard cache contents that have been accessed by the host, making more room for read-ahead data. This setting is not optimal if random accesses occur; use it only if your application is strictly sequential and requires extremely low latency.
6. Click Set Read Ahead Cache Options.

   When processing is complete, the new setting is displayed in the Volume Menu panel.

Changing a Volume's Write-Back Cache Setting

As an Advanced Manage user, 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 disk, 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 drive. 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 disk 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 switch from write-back caching to write-through caching as described in "Changing Auto-Write-Through Triggers and Behaviors" on page 96.

In both caching strategies, active-active failover of the controllers is enabled.

You can enable and disable the write-back cache for each volume. By default, volume write-back cache is enabled. Because controller cache is backed by super-capacitor technology, if the system loses power, data is not lost. For most applications, this is the correct setting. But because back-end bandwidth is used to mirror cache and because this mirroring uses back-end bandwidth, if you are writing large chunks of sequential data (as would be done in video editing, telemetry acquisition, or data logging), write-through cache has much better performance. Therefore, you might want to experiment with disabling the write-back cache. You might see large performance gains (as much as 70 percent) if you are writing data under the following circumstances:

- Sequential writes
- Large I/Os in relation to the chunk size
- Deep queue depth

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

Caution – Only disable write-back cache if you fully understand how your operating system, application, and FC HBA or iSCSI Ethernet adapter move data. You might hinder your storage system's performance if used incorrectly.

To change a volume's write-back cache setting:

1. Select Manage > Volume Management > Volume Menu > Write Back Cache.

   For each virtual disk, the virtual disk panel shows a status icon; the name, RAID level, size, number of disk drives, and number of volumes; and utility status, if any.

2. Select a virtual disk.

   The selected virtual disk's volume names, sizes, and write-back cache settings are displayed, and the amount of free space.

3. Select the standard, snap-pool, or master volume whose cache setting you want to change.

4. Depending on the current setting, click Enable Write Back Cache or Disable Write Back Cache.

   When processing is complete, the new value is displayed in the Volume Menu panel.

Changing Auto-Write-Through Triggers and Behaviors

You can set conditions that cause ("trigger") a controller to change the cache mode from write-back to write-through. You can also specify actions for the system to take when write-through caching is triggered.

For an explanation of cache modes, see "Changing a Volume's Write-Back Cache Setting" on page 94.

To change auto-write-through triggers and behaviors:

1. Select Manage > General Config > System Configuration.
2. Click Advanced Options.
3. In the SCSI Configuration Options panel, under Auto-Write Through Triggers, select the options to enable:
   - Controller Failure Trigger – Switches to write-through if a controller fails. The default is Disabled.
   - Cache Power Trigger – Switches to write-through if cache backup power is not fully charged or fails. The default is Enabled.
   - A/C Power Trigger – Switches to write-through if A/C power fails. The default is Disabled.
   - Power Supply Failure Trigger – Switches to write-through if a power supply unit fails. The default is Disabled.
   - Fan Failure Trigger – Switches to write-through if a cooling fan fails. The default is Disabled.
   - Overtemp Failure Trigger – Forces a controller shutdown if a temperature is detected that exceeds system threshold limits. The default is Disabled.
4. Under Auto-Write Through Behaviors, select the options to enable:
   - Revert when Trigger Condition Clears – Switches back to write-back caching after the trigger condition is cleared. The default is Enabled.
   - Notify Other Controller – In a dual-controller configuration, the partner controller is notified that the trigger condition is met. The default is Disabled.
5. Click Change SCSI Configuration Options.

Deleting a Volume

You can delete a volume when you no longer need it and you want to use the space for another purpose.

Caution – Deleting a volume deletes all data contained in the volume.

To delete a volume:

1. Select Manage > Volume Management > Volume Menu > Delete Volume.

   For each virtual disk, the virtual disk panel shows a status icon; the name, RAID level, size, number of disk drives, and number of volumes; and utility status, if any.

2. Select a virtual disk.

   The selected virtual disk's volume names, sizes, and types are displayed, and the amount of free space.

3. Select the volume to delete.

4. Click Delete Volume.

   A confirmation prompt is displayed.

5. Click OK to confirm the operation or Cancel to stop it.

   If you clicked OK, a message indicates whether the operation succeeded. If it succeeded, the volume is removed from the Volume Menu panel.

Using Snapshot Services

Snapshot services provide data protection by enabling you to create and save snapshots of a volume, where each snapshot preserves the volume's data state at the point in time when the snapshot was created.

Snapshots can be taken of master volumes only. A master volume is a volume that has been enabled for snapshots. You can either create a master volume directly or convert a standard volume to a master volume.

Master volumes are associated with a snap pool, which contains pre-allocated reserve space for the snapshot data. A snap pool represents the storage area that is to hold the copy of the data or pointers to the data created by the snapshot. A snap pool can have 16 associated master volumes. A master volume and its associated snap pool must be owned by the same controller. Threshold levels and associated policies specify the action that the storage system takes when the threshold value of the snap pool is reached.

A snapshot is a virtual volume. While really a set of pointers to a portion of the snap pool, a snapshot behaves like a volume in that it can be mapped to data hosts and the mapping can be assigned a LUN and be made accessible as read-only or read-write, depending on the purpose of the snapshot.

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

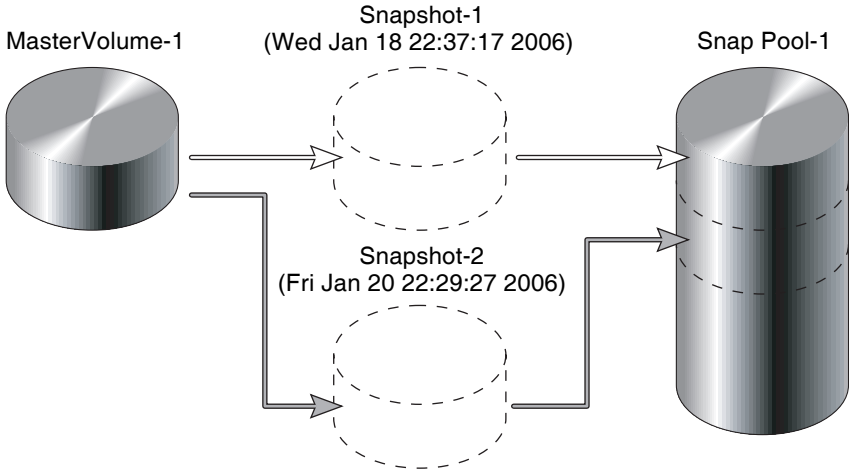

Figure 3-6 Relationship Between a Master Volume and its Snapshots and Snap Pool

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

Maximum Number of Snapshots

Each storage system permits a maximum number of snapshots to be retained. For example, if the maximum number of snapshots allowed on your system is four, when the fifth snapshot is taken, an error message informs you that you have exceeded the maximum number of snapshots allowed on your system. You can delete an existing snapshot and take another snapshot if the size of the snapshot is within the limits of the snap pool threshold. The maximum number of snapshots can be increased by adding a license. See “Managing Licenses” on page 30 for more information.

Determining the Snap Pool Size

Before you can create a master volume you must create a snap pool. A snap pool is the storage area that will hold the copy of the data or pointers to the data created by snapshot of the master volume. When you create a snap pool, RAIDar prompts you to enter the size for the snap pool. To help you calculate the size, first you must know or supply the following information:

- **Snap pool reserve space.** Each snap pool requires a reserve space of 750 Mbyte for internal use.
- **What is the master volume size?** The size of the master volume is specified at the time the master volume is created.
- **How many snapshots is the system going to retain?** This number is dependent upon the configuration limits for your system.
- **What is the average percent of change to the master volume?** If the master volume is going to be updated frequently, the snap pool size will be greater than if the master volume is not updated frequently.

- **What is the number of snapshots that will be written to? If the snapshots are being written to, what is the average amount of data that will be written to a snapshot?** If a snapshot has been made accessible as read-write, you can write to it.
- **Safety margin.** Add a 25% safety margin to the snap pool size.

Snap Pool Sizing Formula: One Master Volume Per Snap Pool

Based on the information from above, use the following formula to calculate the snap pool size for a configuration of one master volume per snap pool:

```
(reserve-space + (volume-size x avg-change x snapshots-retained)
+ (snapshots-modified x avg-write-data) x (1 + safety-margin) =
snap-pool-size
```

Example

The following example information illustrates calculating the snap pool size for a 10-Gbyte volume.

Snap Pool Sizing Parameters	Volume Data
Overhead (Mbyte)	750
Volume size (Mbyte)	10,000
Average percent of change (%)	0.05
Number of snapshots retained	4
Number of modified snapshots	4
Average write data (Mbyte)	1,000
Snap pool size (Mbyte)	6,750
Safety margin (%)	0.25
Snap pool size (Mbyte)	8,438

If you substitute the values from the above example into the snap pool sizing formula, the snap pool size is as follows:

750 + (10,000 x 0.05 x 4) + (4 x 1,000) = 6,750 x 1.25 = 8,438 Mbyte

---

Note – For an explanation of sizes represented by various units, see “Size Representations in RAIDar” on page 23 for more information.

Reverting to Original Data

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

- Deleting only modified data on a snapshot
- Rolling back the data in a master volume

For snapshots that have been made accessible as read-write, you can delete just the modified (write) data that was written directly to a snapshot. When the modified data is deleted, the snapshot data reverts to the original data that was snapped. This feature is useful for application test, for example. You might want to test some code, which writes data to the snapshot. Rather than having to take another snapshot, you can just delete any write data and start again.

The roll back feature enables you to revert the data in a master volume to the data that existed when a specified snapshot was created (preserved data). You also have the option of rolling back to include the modified (write) data on the snapshot since the snapshot was taken. For example, you might want to take a snapshot, mount that snapshot for read/write, and then install new software on that snapshot for test purposes. If the software installation is successful, you can roll back the master volume to the contents of the modified snapshot (preserved data plus the write data).

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

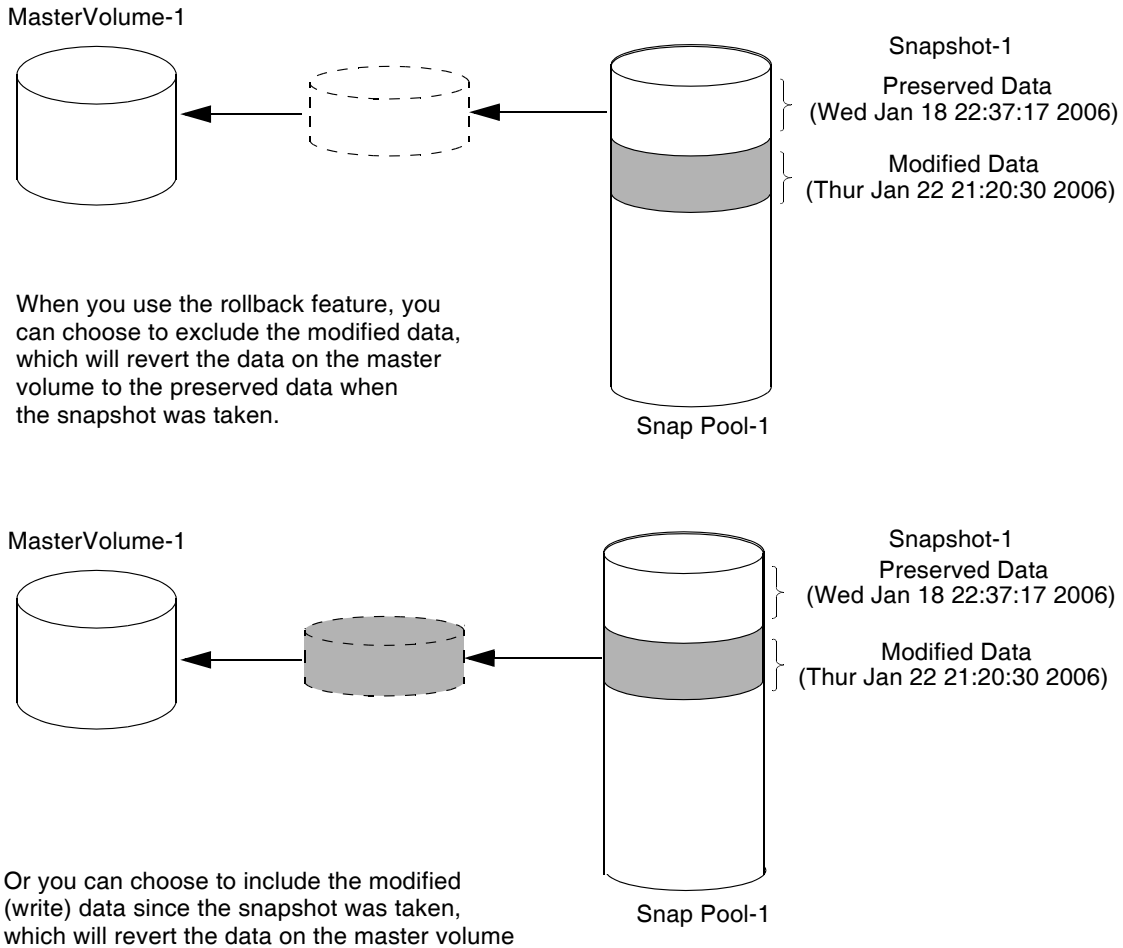

Figure 3-7 Rolling Back the Master Volume

Creating a Snap Pool

Before you can convert a standard volume to a master volume or create a master volume for snapshots, a snap pool must exist. A snap pool and its associated master volumes can be in different virtual disks, but must be owned by the same controller. You can create a maximum of 16 snap pools.

To create a snap-pool volume:

1. Select Manage > Volume Management > Snapshot Services > Create Snap-Pool.

   For each virtual disk, the virtual disk panel shows a status icon; the name, RAID level, size, number of disk drives, and number of volumes; and utility status, if any.

2. Select a virtual disk.

   The selected virtual disk's volume names, sizes, and types are displayed, and the amount of free space.

3. Type a size in Mbyte for the snap pool.

   The size should be between 1,000 Mbyte (1 Gbyte) and the size shown by Free Space Available.

   For information about calculating the snap pool size, see "Determining the Snap Pool Size" on page 99.

4. (Optional) Change the name for the new snap pool.

   The default name is *vdisk-name*_SP*number*. For example, MyVdisk_SP1.

   In a later step, you will associate a master volume with this snap pool. Name the snap pool in such a way that it can be easily identified with the correct virtual disk. The name is case-sensitive and can include 20 characters. Allowed characters include letters, numbers, hyphens, underscores, and spaces.

5. Click Create Snap Pool.

   When processing is complete, a message indicates whether the operation succeeded. If the operation succeeded, the new snap pool is displayed in the Volume Menu panel.

   You should now configure notification thresholds and policies for the snap pool; see "Setting Snap Pool Policies and Thresholds" on page 104.

Setting Snap Pool Policies and Thresholds

Each snap pool has three policy levels that notify you when the snap pool is reaching decreasing capacity. Each policy level has an associated policy that specifies system behavior when the threshold is reached. The following table summarizes the default thresholds and policies. You can set the Warning and Error thresholds and the Error and Critical policies.

Policy Level	Threshold	Policy
Warning	75%	Notify Only.
Error	90%	Delete Oldest Snapshots.
Critical	99%	Delete Snapshots.

Note – RAIDar notifies you of events based on your event notification settings. See "Configuring Event Notification" on page 47 for more information about how and under what conditions the system alerts you when specific events occur.

Policy Trigger Behavior

A policy might be triggered before it appears that a specified threshold has been reached. This is because a threshold percentage is based on the size of the snap pool, less a fixed amount of 750 Mbyte for internal use. This fixed amount guarantees that there is enough reserve space to store pending data for which the controller has space. The following example demonstrates policy trigger behavior:

Snap pool size = *10,000 Mbyte* (10 Gbyte)

Snap pool reserve = *750 Mbyte*

Space available = *9,250 Mbyte*

Policy trigger set at default error level of 90% = *9,250 Mbyte* x 0.9 = *8,325 Mbyte*

In the above example, the 90% trigger occurs at 8.325 Gbyte. If you did not take into account the amount of reserve space, you might expect the trigger to occur at 9 Gbyte (10 Gbyte x 0.9). For larger snap pools, the impact of this reserve space is less noticeable.

To set a snap pool's policies and thresholds:

1. Select Manage > Volume Management > Snapshot Services > Set Snap-Pool Policy.

   For each virtual disk, the virtual disk panel shows a status icon; the name, RAID level, size, number of disk drives, and number of volumes; and utility status, if any.

2. Select a virtual disk.

   The selected virtual disk's volume names, sizes, and types are displayed, and the amount of free space.

3. Select the snap pool to configure.

4. Set the Warning Policy threshold.

   When the snap pool reaches the specified percent of capacity, the system generates an event that the threshold has been reached. Notify Only is the only Warning Policy. The default is 75%. The Warning threshold must be less than the Error threshold.

5. Select an Error Policy and set the threshold to a value less than 99%.

   When the snap pool reaches the specified percent of capacity, the system generates an event that the threshold has been reached. The default is 90%. The Error threshold must be less than the Critical threshold of 99%.

   The system takes further action, depending on the Error Policy, as follows:

   - Delete Snapshots – Automatically deletes *all* snapshots associated with the snap pool.
   - Auto Expand – Expands the snap pool by the value specified in the Size To Expand (MBytes) field. The amount of space specified must exist as free space on the virtual disk on which the snap pool resides. If there isn't enough space on the virtual disk, the auto-expand operation fails, and an insufficient virtual disk free space error is logged.
   - Halt Writes – Halts all writes to the master volume (each write returns an error). Snapshot data is preserved.
   - Delete Oldest Snapshots – Deletes the oldest snapshots until the amount of data in the snap pool is below the threshold. This option is the default.
   - Notify Only – No further action.

6. Select a Critical Policy.

   When the snap pool reaches 99% capacity, the system takes action depending on the Critical Policy, as follows:

   - Delete Snapshots – Automatically deletes *all* snapshots associated with the snap pool. This option is the default.
   - Halt Writes – Halts all writes to the master volume (each write returns an error). Snapshot data is preserved.
   - Delete Oldest Snapshots – Deletes the oldest snapshots until the amount of data in the snap pool is below the threshold.

7. Click Set Policy & Threshold.

   The changes take effect immediately.

Creating a Master Volume

You can take snapshots of a snapshot-enabled volume (a master volume). A maximum of 16 master volumes can exist and they all can be associated with a single snap pool. A master volume and its snap pool can be in different virtual disks, but must be owned by the same controller.

You can either:

- Create a master volume, as described below
- Convert a standard volume to a master volume; see "Converting a Standard Volume to a Master Volume" on page 108

Creating a New Volume as a Master Volume

To create a master volume:

1. Select Manage > Volume Management > Snapshot Services > Create Master Volume.

   For each virtual disk, the virtual disk panel shows a status icon; the name, RAID level, size, number of disk drives, and number of volumes; and utility status, if any.

2. Select the virtual disk where you want to create the volume.

   The selected virtual disk's volume names, sizes, and types are displayed, and the amount of free space.

3. Select a snap pool.

   Only snap pools owned by the same controller as the selected virtual disk are listed.

4. Type a size in increments of 1 Mbyte for the new volume.

5. (Optional) Change the name for the new volume.

   The default name is *vdisk-name*_V*number*. For example, MyVdisk_V1.

   The name is case-sensitive and can include 20 characters. Allowed characters include letters, numbers, hyphens, underscores, and spaces.

6. (Optional) Change the LUN setting.
   - NONE – The volume is not accessible by connected hosts. This setting is the default. You can map the volume to hosts later; see "Managing Volume Mappings" on page 90.
   - 0–127 – The default LUN to use when mapping this volume to hosts.

7. Click Create Master Volume.

   When processing is complete, the new volume is displayed in the Volume Menu panel.

Converting a Standard Volume to a Master Volume

To convert a standard volume to a master volume:

1. Select Manage > Volume Management > Snapshot Services > Snapshot-Enable Volume.

   For each virtual disk, the virtual disk panel shows a status icon; the name, RAID level, size, number of disk drives, and number of volumes; and utility status, if any.

2. Select a virtual disk.

   The selected virtual disk's volume names, sizes, and types are displayed, and the amount of free space.

3. Select the standard volume to convert.

4. Select a snap pool.

   Only snap pools owned by the same controller as the selected virtual disk are listed.

5. Click Convert To Master Volume.

   When processing is complete, the volume type is updated in the Volume Menu panel.

Taking a Snapshot

You can take a snapshot of the data state of a selected master volume. The snapshot data is stored in the snap pool associated with the master volume.

To take a snapshot:

1. Select Manage > Volume Management > Snapshot Services > Take Snapshot.

   For each virtual disk, the virtual disk panel shows a status icon; the name, RAID level, size, number of disk drives, and number of volumes; and utility status, if any.

2. Select a virtual disk.

   The selected virtual disk's volume names, sizes, and types are displayed, and the amount of free space.

3. Select the master volume to take a snapshot of.

4. (Optional) Change the name for the snapshot.

   The default name is *vdisk-name_volume-name_*SS*number*. For example, MyVdisk_V2_SS1.

   Name the snap pool in such a way it can be easily identified with the correct virtual disk. The name is case-sensitive and can include 20 characters. Allowed characters include letters, numbers, hyphens, underscores, and spaces.

5. Click Take Snapshot.

   When processing is complete, the new snapshot is displayed in the Volume Menu panel.

Updating a Snapshot by Resetting

You can reset a snapshot to replace its content with the current data state of the associated master volume. The selected snapshot is replaced with a current snapshot having the same characteristics, such as name and LUN. The snapshot data is stored in the snap pool associated with the master volume. Before being reset, a snapshot must be unmounted from hosts.

Caution – Before resetting a snapshot you must unmount it from data hosts to avoid data corruption.

To reset a snapshot:

1. Unmount the snapshot from hosts.
2. Select Manage > Volume Management > Snapshot Services > Reset Snapshot.

   For each virtual disk, the virtual disk panel shows a status icon; the name, RAID level, size, number of disk drives, and number of volumes; and utility status, if any.
3. Select a virtual disk.

   The selected virtual disk's volume names, sizes, and types are displayed, and the amount of free space.
4. Select the snapshot to reset.
5. Click Reset Snapshot.

   When processing is complete, a message indicates whether the operation succeeded.

Deleting Modified Data

If a snapshot has been made accessible as read-write, you have the option of deleting only the modified (write) data that has been written to it. (See “Managing Volume Mappings” on page 90 for information about setting access privileges.) The amount of data that has been written to a snapshot is shown in the Unique Data field on the Snapshot Overview page. (See the information provided for snapshots in “Viewing Information About All Snap Pools, Master Volumes, and Snapshots” on page 114.) You must unmount the snapshot from hosts before deleting modified data.

Caution – Before deleting modified data you must unmount the snapshot from data hosts to avoid data corruption.

To delete the modified (write) data from a snapshot:

1. Unmount the snapshot from hosts.
2. Select Manage > Volume Management > Snapshot Services > Delete Modified Data.

   For each virtual disk, the virtual disk panel shows a status icon; the name, RAID level, size, number of disk drives, and number of volumes; and utility status, if any.
3. Select a virtual disk.

   The selected virtual disk’s volume names, sizes, and types are displayed, and the amount of free space.
4. Select the snapshot from which you want to delete modified data.
5. Click Delete Modified Data.

   When processing is complete, on the Snapshot Overview page the snapshot’s Unique Data field shows zero.

Rolling Back a Master Volume

You can roll back (revert) the data in a master volume to the data that existed when a specified snapshot was created. You also have the option of rolling back to the modified (write) data on the snapshot.

Caution – Before rolling back a master volume you must unmount it from data hosts to avoid data corruption. If you want to include snapshot modified data in the rollback, you must also unmount the snapshot. You can remount the master volume after the rollback has started. You can remount the snapshot when the rollback is complete.

Caution – Whenever you perform a rollback, the data that existed on the master volume is replaced by the data on the snapshot; that is, all data on the master volume written since the snapshot was taken is lost. As a precaution, take a snapshot of the master volume before starting a rollback.

Only one rollback is allowed on the same master volume at one time. Multiple rollbacks on subsequent volumes on the same snap pool are performed sequentially; that is, additional rollbacks are queued until the current rollback is complete. However, after the rollback is requested, the master volume is available for use as if the rollback has already completed.

To rollback a master volume:

1. Unmount the master volume from hosts.
2. If the rollback will include snapshot modified data, unmount the snapshot from hosts.
3. Select Manage > Volume Management > Snapshot Services > Rollback Volume.

   For each virtual disk, the virtual disk panel shows a status icon; the name, RAID level, size, number of disk drives, and number of volumes; and utility status, if any.
4. Select a virtual disk.

   The selected virtual disk's volume names, sizes, and types are displayed, and the amount of free space.
5. Select the master volume to rollback.
6. Select the snapshot to rollback the master volume to.

7. Select whether to include or exclude data modified in the snapshot since it was taken.

   The default is Exclude, which means that the master volume will contain only the data that existed when the snapshot was taken.

8. Click Rollback Master Volume.

   When processing is complete, a message indicates whether the operation succeeded.

Deleting a Snapshot

You can delete snapshots at any time, including when:

- The associated snap pool is reaching capacity and you want to free some space
- The maximum number of snapshots is reached and you want to delete older snapshots
- You no longer need the data associated with the snapshot

To delete a snapshot:

1. Unmount the snapshot from hosts.
2. Select Manage > Volume Management > Snapshot Services > Delete Snapshot.

   For each virtual disk, the virtual disk panel shows a status icon; the name, RAID level, size, number of disk drives, and number of volumes; and utility status, if any.

3. Select a virtual disk.

   The selected virtual disk's volume names, sizes, and types are displayed, and the amount of free space.

4. Select the snapshot to delete.
5. Click Delete Snapshot.

   When processing is complete, the snapshot is removed from the Volume Menu panel. All data uniquely associated with the snapshot is deleted and associated space in the snap pool is freed for use.

Viewing Information About All Snap Pools, Master Volumes, and Snapshots

To view information about all snap pools, master volumes, and snapshots:

1. Select Manage > Volume Management > Snapshot Services > Snapshot Overview.

   For each virtual disk, the virtual disk panel shows a status icon; the name, RAID level, size, number of disk drives, and number of volumes; and utility status, if any.

2. Select a virtual disk.

   The selected virtual disk's volume names, sizes, and types are displayed, and the amount of free space.

3. Select a volume.

   The Volume Status panel shows the following information, depending on the type of volume selected.

Volume Type	Field	Description
(All)	Volume Type (not shown for standard volumes)	One of the following volume types: Master – A standard volume that is enabled for snapshots and is associated with a snap pool. Snap Pool – A virtual volume in which snapshots of the associated master volume are stored. Snapshot – A volume that preserves the data state of a master volume at the point in time when the snapshot was created.
	Volume Name	Name assigned to the volume.
	Belongs to Virtual Disk	Name of the virtual disk the volume is part of.
	Volume Presented Globally?	Specifies whether the volume is visible to all connected hosts. If the value is Yes, the LUN is also shown.
Standard	Volume Size	Volume size in Mbyte. To change the volume size, go to Volume Menu > Expand Volume.
	Percent of Total Virtual Disk	The percentage of the total virtual disk that this volume occupies.

Volume Type	Field	Description
Snap Pool	Free Space	The amount of free space in the snap pool.
	Master Volumes	The number of master volumes using the snap pool.
	Snapshots	The number of snapshots in the snap pool.
	Thresholds	For each threshold level, the action configured to occur when snap-pool usage exceeds the specified percentage. For information about thresholds, see "Setting Snap Pool Policies and Thresholds" on page 104.
	Master Volumes on this Snap Pool	If the snap pool contains snapshots, the name of each associated master volume, the number of snapshots taken, and their total size.
Master	Associated Snap Pool	The name of the associated snap pool.
	Number of Snapshots of Volume	The number of existing snapshots.
	Snapshot Data Size	Total size of stored snapshots.
	Rollback Percentage	The approximate percentage complete, if the master volume is being rolled back.
	Snapshots of this Master Volume	If any snapshots have been taken, each snapshot's name, the date and time created, and the size.
Snapshot	Date Created	The date and time when the snapshot was created.
	Committed?	One of the following snapshot status values: Valid – The snapshot is available to be mounted and used. The snapshot volume has been created and associated with a master volume. Pending – The snapshot is not yet valid for use. During this brief period, the snapshot volume is allocated but not yet activated; that is, data is not yet being collected for this snapshot. Offline – The snapshot volume is offline. The cause could be a rollback or volume-copy including modified data. Snap-Pool Offline – The snap pool is not available for use. The cause could be a problem with the snap pool itself or its virtual disk.
	Master Volume	The name of the master volume that the snapshot was taken of.

Volume Type	Field	Description
	Snap Pool	The name of the snap pool that the snapshot data is stored in.
	Data	Specifies the following amounts of data associated with the snapshot: Snap Data – The total amount of data associated with the specific snapshot (data copied from a master volume to a snapshot and data written directly to a snapshot). Unique Data – The amount of data that has been written to the snapshot since the last snapshot was taken. If the snapshot has not been written or is deleted, this field will show a value of 0. Shared Data – The amount of data that is potentially shared with other snapshots and associated amount of space that is guaranteed to be freed if the snapshot is deleted. This represents the amount of data written directly to the snapshot. It also includes data copied from the master volume to the snap pool for the oldest snapshot, since that snapshot does not share data with any other snapshot. For a snapshot that is not the oldest, if the modified data is deleted or if it had never been written to, this field will show a value of 0.

Using Volume Copy Services

Volume copy services enable you to copy a source volume, view the status of an in-progress volume copy, and cancel an in-progress volume copy. The source volume can be either a master volume or a snapshot.

Copying a Volume

You can copy a master volume or a snapshot to a new standard volume. The volume-copy operation takes a snapshot of all data in the source volume and creates a destination volume that you specify. The destination volume must be in a virtual disk owned by the same controller as the source volume.

While the operation is in progress, the destination volume type is shown as Standard*. When the operation is complete, the volume type becomes Standard and the destination volume can be mapped for use.

Caution – Before copying a master volume or a snapshot volume and its modified data, you must unmount the volume from data hosts to avoid data corruption. You can remount the master volume after the copy has started. You can remount the snapshot when the copy is complete.

To copy a volume:

1. If copying a master volume or a snapshot volume and its modified data, unmount the volume from hosts.
2. Select Manage > Volume Management > Volume-Copy Services > Volume-Copy.

   For each virtual disk, the virtual disk panel shows a status icon; the name, RAID level, size, number of disk drives, and number of volumes; and utility status, if any.
3. Select the master volume or snapshot that you want to copy.

   The selected virtual disk's volume names, sizes, and types are displayed, and the amount of free space.
4. Select a destination virtual disk.
5. Type a name for the destination volume.

   The name is case-sensitive and can include 20 characters. Allowed characters include letters, numbers, hyphens, underscores, and spaces.

6. If the source volume is a snapshot, select whether the copy should include or exclude data modified in the snapshot since it was taken.

   The default is Exclude, which means that the copy will contain only the data that existed when the snapshot was taken.

7. Click Volume Copy.

   When processing is complete, the new volume is displayed in the Volume Menu panel.

Viewing the Status of a Volume Copy

You can view the status of a destination volume being created by a volume-copy operation. If the operation has completed or if you select a different type of volume, the page shows that there is no status information.

To view the status of an in-progress volume copy:

1. Select Manage > Volume Management > Volume-Copy Services > Volume-Copy Status.

   For each virtual disk, the virtual disk panel shows a status icon; the name, RAID level, size, number of disk drives, and number of volumes; and utility status, if any.

2. Select the destination volume's virtual disk.

   The selected virtual disk's volume names, sizes, and types are displayed, and the amount of free space.

3. Select a destination volume whose type is Standard*.

   The Volume-Copy Status panel shows the following information:

   - Volume Type – Standard*.
   - Volume Name – Name assigned to the volume.
   - Belongs to Virtual Disk – Name of the virtual disk the volume is part of.
   - Volume Presented Globally? – Specifies whether the volume is visible to all connected hosts. If the value is Yes, the LUN is also shown.
   - Volume Serial Number – Serial number of the volume being created.
   - Source Volume Name – Name of the volume being copied.
   - Percent Complete – Percent complete of the volume copy.

Canceling a Volume Copy

You can cancel an in-progress volume-copy operation. When the cancellation is complete, the destination volume is deleted.

To cancel an in-progress volume copy:

1. Select Manage > Volume Management > Volume-Copy Services > Abort Volume-Copy.

   For each virtual disk, the virtual disk panel shows a status icon; the name, RAID level, size, number of disk drives, and number of volumes; and utility status, if any.

2. Select either the source volume's virtual disk or the destination volume's virtual disk.

   The selected virtual disk's volume names, sizes, and types are displayed, and the amount of free space. The destination volume's type is Standard*.

3. Select either the source volume or the destination volume.

4. Click Abort Volume Copy.

   A confirmation prompt is displayed.

5. Click OK to confirm the operation or Cancel to stop it.

   If you clicked OK, a message indicates whether the operation succeeded. If it succeeded, the destination volume is removed from the Volume Menu panel.

Using the Scheduler

You can use the Scheduler feature to create tasks and define schedules at which the system will automatically perform those tasks.

Actions you can perform on the Scheduler page are:

- Create tasks to take a snapshot, reset a snapshot, or copy a volume
- View task information
- Delete tasks
- Schedule tasks
- View schedule information
- Delete schedules

Panels on this page have these icons:

- ▶ – Click to show the panel's content.
- ▼ – Click to hide the panel's content.
- ☒ – Click to cancel creating a task or schedule.

While you are managing tasks and schedules, running tasks or use of other storage system interfaces can cause displayed data to become outdated. The following update notice and button are displayed in the message area so you can update the page when you are ready:

Creating a Take Snapshot Task

You can create a task to take a snapshot of a master volume, if at least one master volume exists. When you create the task you a prefix to identify snapshots taken by that task, and the number of snapshots with that prefix to retain (known as the retention count). When the task runs, the Scheduler compares the number of snapshots that exist with the retention count:

- If the retention count has not been reached, the snapshot is taken.
- If the retention count has been reached, the oldest snapshot with that prefix is unmapped, reset, and renamed to the next name in the sequence.

To create a task to take a snapshot of a master volume:

1. Select Manage > Scheduler > Manage Scheduler.
2. In the Tasks panel click Add New Task.

   The Create Task panel is displayed.
3. Select Take Snapshot.
4. Select a master volume to take snapshots of.
5. Specify a prefix to identify snapshots created by this task.

   The prefix is case-sensitive and can include 14 characters. Allowed characters include letters, numbers, hyphens, underscores, and spaces.

   Automatically created snapshots are named *prefix*_S*xxxx*, where *xxxx* increments from 0001 to 9999 before rolling over.
6. Specify the number of snapshots with this prefix to retain.

   The default and minimum value is 1. Your license determines the maximum value.
7. Specify a name for the task.

   The name is case-sensitive and can include 32 characters. Allowed characters include letters, numbers, hyphens, underscores, and spaces.
8. Click Create Task.

   The Tasks panel is updated and task information is displayed in the Task Details panel.

Creating a Reset Snapshot Task

You can create a task to reset a snapshot, which deletes the data in the snapshot and resets it to the current data in the associated master volume. The snapshot's name and other volume characteristics are not changed.

Caution – Before scheduling a reset snapshot task, consider that if the snapshot is mounted to a host operating system, the snapshot must be unmounted before the reset is performed; leaving it mounted can cause data corruption.

To create a task to reset a snapshot:

1. Select Manage > Scheduler > Manage Scheduler.
2. In the Tasks panel click Add New Task.

   The Create Task panel is displayed.
3. Select Reset Snapshot.
4. Select a snapshot volume to reset.
5. Specify a name for the task.

   The name is case-sensitive and can include 32 characters. Allowed characters include letters, numbers, hyphens, underscores, and spaces.
6. Click Create Task.

   The Tasks panel is updated and task information is displayed in the Task Details panel.

Creating a Volume Copy Task

If an installed license enables this feature, you can copy a snapshot or a master volume to a new standard volume. The destination volume must be in a virtual disk owned by the same controller as the source volume.

To create a task to copy a volume:

1. Select Manage > Scheduler > Manage Scheduler.
2. In the Tasks panel click Add New Task.

   The Create Task panel is displayed.
3. Select Volume Copy.
4. Select a snapshot or master volume to copy.
5. Select a destination virtual disk for the copy.
6. Specify a prefix to identify volumes created by this task.

   The prefix is case-sensitive and can include 14 characters. Allowed characters include letters, numbers, hyphens, underscores, and spaces.

   Automatically created volumes are named *prefix*_V*xxxx*, where *xxxx* increments from 0001 to 9999 before rolling over.
7. Select whether to include or exclude modified write data from the snapshot in the copy.

   The default is Exclude.
8. Specify a name for the task.

   The name is case-sensitive and can include 32 characters. Allowed characters include letters, numbers, hyphens, underscores, and spaces.
9. Click Create Task.

   The Tasks panel is updated and task information is displayed in the Task Details panel.

Viewing Task Information

To view information about existing tasks:

1. Select Manage > Scheduler > Manage Scheduler.

   The Tasks panel shows the name, type, and status of existing tasks.

   If a task fails, an error icon is displayed and the task type and status are shown in red. The task remains in the current state until an associated schedule initializes the task to run again. An error message in the Task Details panel specifies the failure reason.

2. For more information about a task, click a task name.

   For a Take Snapshot task, the Task Details panel shows:

   - Task name and type
   - Task status (Ready or Active)
   - Task state (Init, Vol Verified, License Checked, Name Created, Snap Created, or Snap Verified)
   - Master volume name and serial number
   - Snapshot prefix
   - Retention count
   - Last snapshot created, if the task has run
   - Retained snapshots, if any
   - Error message, if any

   For a Reset Snapshot task, the Task Details panel shows:

   - Task name and type
   - Task status (Ready or Active)
   - Task state (Init or Snap Verified)
   - Snapshot name and serial number
   - Error message, if any

   For a Volume Copy task, the Task Details panel shows:

   - Task name and type
   - Task status (Ready or Active)
   - Task state (Init, Vol Verified, Name Created, or Vol Created)
   - Source volume name and serial number
   - Destination virtual disk name and serial number
   - Destination volume prefix
   - Include modified data
   - Last copy created, if the task has run
   - Error message, if any

Deleting a Task

You can delete an unscheduled task. If the task is scheduled, you must delete the schedule first.

To delete a task:

1. Select Manage > Scheduler > Manage Scheduler.
2. In the Tasks panel click a task name.
3. In the Task Details panel, click Delete Task.
4. Click OK to confirm the operation or Cancel to stop it.

Creating a Schedule

To schedule a task:

1. Select Manage > Scheduler > Manage Scheduler.
2. In the Schedules panel click Add New Schedule.

   The Create Schedule panel is displayed.
3. Specify the date when the schedule should start running. Either:
   - Type a date using the format mm/dd/yyyy
   - Click  to display a calendar window in which you can select the date

   The default is the current date.
4. Specify the time when the schedule should start running.

   The default is the current time.
5. Enable and configure recurrence and constraint rules:
   - Every – Specifies how often the task should run.
   - Between – Specifies a time range within which the task should run.
   - Only On – Specifies days when the task should run.

     You can select a combination of: any day or a day by number; a day by type or name; and all months or a month by name.

     For the day number, the Specific option uses a number you type in an adjacent field.

- Repeat – Specifies the number of times the task should run, including the first time.
- Expires On – Specifies the date and time when the task should stop running.

6. Select a task to schedule.
7. Specify a name for the schedule.

   The name is case-sensitive and can include 32 characters. Allowed characters include letters, numbers, hyphens, underscores, and spaces.

8. Click Create Schedule.

   The task list in the Schedules panel is updated and schedule information is displayed in the Schedule Details panel.

Viewing Schedule Information

To view information about existing schedules:

1. Select Manage > Scheduler > Manage Scheduler.

   The Schedules panel shows the name, associated task, and the next time the task will run.

2. For more information about a schedule, click a schedule name.

   The Schedule Details panel shows:

   - Schedule name
   - Schedule specifications (recurrence and constraint settings)
   - Schedule status (Ready or Expired)
   - Next time the scheduled task will run
   - Task to run

Deleting a Schedule

To delete a schedule:

1. Select Manage > Scheduler > Manage Scheduler.
2. In the Schedules panel click a schedule name.
3. In the Schedule Details panel, click Delete Schedule.
4. Click OK to confirm the operation or Cancel to stop it.

Configuring In-band Management Services

You can manage the storage system in-band with custom applications written using the Configuration API (CAPI). If you are not using CAPI-based applications, you can disable in-band management. You can also monitor system status in-band based on SCSI Enclosure Services (SES) data.

To configure in-band management services:

1. Select Manage > General Config > Services Security.
2. In the Inband Management Services panel, set these options:
   - Inband CAPI Capability – Used for in-band management of the system from host-based management applications, such as a VDS or VSS client. If this option is disabled, the applications will lose access to the system. The default is Enabled.
   - Inband SES Capability – Used for in-band monitoring of system status based on SCSI Enclosure Services data. The default is Enabled.
3. Click Update Inband Management Services.

CHAPTER 4

Managing Disk Drives and Enclosures

This chapter describes how to use RAIDar to manage a system's disk drives and enclosures. Topics covered in this chapter are:

Managing Disk Drives

RAIDar provides a variety of functions related to disk drives.

Viewing Disk Drive Information

You can view two types of information about disk drives:

- A list of all disk drives connected to the system
- The status of all disk drives in a virtual disk

Viewing All Disk Drives

To view information about all disk drives connected to the system:

- Select Monitor > Status > Advanced Settings > Disk Drive List.

  For a description of the information contained on the Disk Drive List page, see "Disk Drive List" on page 149.

  Disk drives that are not members of any virtual disk are listed as Available. Drives that contain leftover metadata from a previous virtual disk are listed as Leftover. Leftover drives occur when drives are removed and reinserted, or fail temporarily and are not operating again.

  To clear leftover metadata, use the Clear Metadata utility. See "Clearing Metadata From a Disk Drive" on page 130.

Viewing Disk Drive Status

To view the status of the drives in a selected virtual disk:

- Select Manage > Virtual Disk Config > Vdisk Configuration > Disk Drive Status.

  For a description of the information contained on this page, see “Disk Drive Status” on page 68.

Note – If a disk drive has failed or malfunctioned, it might not be listed.

Clearing Metadata From a Disk Drive

All of the member disk drives in a virtual disk contain metadata in the first sectors. The system uses the metadata to identify virtual disk members after restarting or replacing enclosures.

Clear the metadata if you have a disk drive that was previously a member of a virtual disk. Disk drives in this state are identified as Leftover. After you clear the metadata, you can use the disk drive in a virtual disk or as a spare.

To clear metadata from drives:

1. Select Manage > Utilities > Disk Drive Utilities > Clear Metadata.

   An enclosure view is displayed in which only Leftover and Available drives are selectable. Available drives are considered to have had their metadata cleared, but are selectable in case a drive with partial metadata has been inserted into the system.

2. Select the drives whose metadata you want to clear.
3. Click Clear Metadata For Selected Disk Drives.

Enabling or Disabling SMART Changes

As an Advanced Manage user, you can configure the ability to change the Self-Monitoring Analysis and Reporting Technology (SMART) settings for all disk drives in the storage system. When SMART is enabled, SMART events are recorded in the event log and are counted on the Disk Error Stats page. This information enables you to monitor your disk drives or analyze why a disk drive failed.

For more information about the event log, see “Displaying the Event Log” on page 165. For more information about disk error statistics, see “Disk Drive Error Statistics” on page 173.

To configure SMART:

1. Select Manage > General Config > Disk Configuration.
2. Set SMART to one of the following options:
   - Don’t Modify – Allows current drives to retain their individual SMART settings and does not change the setting for new drives added to the system.
   - Enable – Enables SMART for all current drives after the next rescan and will automatically enable SMART for new drives added to the system. This option is the default.
   - Disable – Disables SMART for all current drives after the next rescan and will automatically disable SMART for new drives added to the system.
3. (Optional) Select or deselect the automatic rescan option.

   If you want to perform a manual rescan before the drive settings take effect, clear the automatic rescan check box. You can perform a manual rescan on the Manage > Utilities > Disk Drive Utilities > Rescan page (see “Rescanning for Drive Changes” on page 182).
4. Click Change Disk Option Configuration.

Viewing Disk Drive Read-Cache Status

To view a disk drive's read-cache status:

1. Select Manage > Utilities > Disk Drive Utilities > Display Disk Cache.

   The page shows the enclosure view with a drive selected, and shows the drive's node WWN and read-cache status. The first drive in the enclosure is selected by default.

2. Select a drive.

3. Click Show Disk Drive Cache Status.

   The selected drive's node WWN and read-cache status are displayed.

Illuminating a Drive Module LED

You can blink a drive module's Power/Activity/Fault LED to help you visually locate the drive in its enclosure. For information about locating the enclosure that contains the drive, see "Illuminating Enclosure LEDs" on page 138. For information about locating a faulty disk drive, see the *Troubleshooting Guide*.

To locate disk drives:

1. Select Manage > Utilities > Disk Drive Utilities > Locate Disk Drive.

   The page shows the enclosure view.

2. Select the drives to locate.

3. Click Update LED Illumination.

   The lower LED on each selected drive starts blinking yellow.

To stop blinking a drive's LED:

1. Clear the drive's check box.

2. Click Update LED Illumination.

Viewing and Updating Disk Drive Firmware Versions

You can view the firmware version and type of each disk drive in each enclosure connected to the system. If your drives support it, you can also update the disk drive firmware using RAIDar.

Viewing Disk Drive Types and Firmware Versions

To view the firmware version (revision) and type of each disk drive in each enclosure connected to the system, do either of the following:

- Select Manage > Update Software > Disk Drive Firmware > Show Disk Drives.

  The page shows similar information to the Disk Drive List page; see "Disk Drive List" on page 149.

- Select Manage > Update Software > Disk Drive Firmware > Show Disk Drive Types.

  The following information is displayed for each drive:

  - Vendor – Drive manufacturer.
  - Model – Drive model.
  - Firmware Revision – Revision code for the firmware currently in the drive.
  - Drive Size – Drive size in Gbyte.
  - Total Number of this Type – The number of drives that have the same vendor, model, and firmware revision. For example, two identical drives with different firmware revisions are considered to be different types.

Updating Disk Drive Firmware

You can update disk drive firmware by loading a firmware update file obtained from the disk drive manufacturer or your reseller.

Caution – Updating the firmware of disk drives in a virtual disk risks the loss of data and causes the drives to be temporarily inaccessible. Before performing a firmware update, back up the virtual disk data and stop I/O to the virtual disk.

To update disk drive firmware:

1. Select Manage > Update Software > Disk Drive Firmware > Update Firmware.
2. Select the type of disk drives to update.

   Drives that have the same manufacturer, model, and firmware revision are considered the same type. For example, two identical disk drives with different firmware revisions are considered to be different types. If firmware update is not supported for a disk drive type, the Select column shows "Not Supported" for that type and you cannot continue the firmware update process.
3. Click Select Type And Continue.

   Disk drives of the type you selected are listed and the following information is displayed for each disk drive:
   - Device WWN – The disk drive's node WWN.
   - Location Encl:Slot – Enclosure number and slot number containing the drive.
   - Size – The size of the disk drive in Gbyte.
   - Manufacturer – The disk drive manufacturer.
   - Model – The disk drive model number.
   - Rev – The four-digit firmware revision code for the firmware currently on the disk drive.
   - Serial Number – The disk drive's vendor-specific serial number.
   - Virtual Disk Member – Specifies whether this disk drive is part of a virtual disk.

   If more than two drives are listed, a Select All check box is displayed.
4. Select the disk drives to update.
5. Click Continue.
6. Click Browse to select the firmware update file.

7. Click Load Device Firmware File.
8. To start the firmware update, click Start Firmware Update.

   To cancel the firmware update, click Cancel.

   The file is transferred to the controller where it is temporarily stored prior to download to the disk drives. Once the firmware update process has started, the Drive Firmware Loading Progress page provides the update progress of each disk drive, including when the firmware update completes successfully.

   This operation can take many minutes or hours to complete. During the update, the following operations are blocked so that they do not interfere with the update:

   - Updating controller software (buffer interference)
   - Saving logs to a file (buffer interference)
   - Displaying disk drive read-cache status (SCSI interference)
9. When processing is complete, verify that the proper firmware version, size, and speed are reported for each updated disk drive.

Stopping or Aborting a Disk Drive Firmware Update

The Stop Device Firmware Update button stops the update operation at the next point that will leave the disk drives in a clean state. If you click this button while the file is being downloaded to the controller, the download stops in a few seconds and no disk drives are updated. If you click this button after download to the disk drives has started, the process is stopped after the update to the current disk drive is complete. No updates that are already done or started are undone. It can take up to several minutes for a stop operation to complete.

Caution – The Abort Device Firmware Update button immediately stops the firmware update and leaves the disk drive in an unknown, possibly unusable state. If you choose this option, wait two minutes to enable the disk drive to possibly finish writing to its nonvolatile memory.

Managing Enclosures

Each controller module and expansion module contains an Expander Controller (EC). The storage system can query EC for information about enclosure environmental conditions such as temperature, power supply and fan status, and the presence or absence of disk drives. The system can also communicate information to the EC about RAID activities such as disk drive rebuilds and failed disk drives. The EC is also referred to as the enclosure management processor (EMP).

Displaying Enclosure Status

You can view enclosure status information from the following RAIDar pages:

- Monitor > Status > Enclosure Status. See “Enclosure Status” on page 156 for more information.
- Monitor > Status > Module Status. See “Module Status” on page 153 for more information.
- Manage > General Config > Enclosure Management. See “Using the Enclosure Management Page” on page 136.
- System Panel at the bottom of every page. See “System Panel” on page 22 for more information.

Using the Enclosure Management Page

For a multi-enclosure system, the enclosure view shows enclosures in order by enclosure ID, with enclosure ID zero at the top. If the physical arrangement of enclosures in your system differs, configuring enclosures' display order by setting rack and position values might help you to manage the system. The enclosure view shows enclosures in the following order:

1. By rack number, if set
2. By rack position, if set
3. By enclosure ID, if neither rack number nor position number is set

On the Enclosure Management page you can perform the following tasks:

- View enclosure details
- Enter information to identify an enclosure
- Illuminate the Unit Locator LED to locate an enclosure
- Change the EMP poll rate

Viewing Enclosure Details

To view enclosure details:

1. Select Manage > General Config > Enclosure Management.
2. Pause your cursor over an enclosure icon.

   A pop-up shows the enclosure status and other details.

Entering Enclosure Information

To enter the name, location, rack number, and rack position for an enclosure:

1. Select Manage > General Config > Enclosure Management.
2. If there is more than one enclosure, select the enclosure for which you want to enter information.
3. Set the following values:
   - Enclosure Name – Type a name to identify the enclosure.

     The name can include 19 characters. Allowed characters include letters, numbers, hyphens, underscores, and spaces.
   - Enclosure Location – Type a description of the enclosure's physical location.

     The location can include 19 characters. Allowed characters include letters, numbers, hyphens, underscores, and spaces.
   - Rack Number – Select a number from 1 to 16to identify the rack the enclosure is in.

     The default is Not Set.
   - Enclosure Position in Rack – Select a number from 1 to 16 to identify where the enclosure is positioned in the rack.

     The default is Not Set.

     By convention, 1 indicates top and 16 indicates bottom.

Note – Virtual Disk Enclosure View sorts the enclosures by rack number first, then rack position as explained in “Using the Enclosure Management Page” on page 136.

---

4. Click Update Enclosure Information.

   The Order of Enclosures panel shows the order in which enclosures will display in enclosure view.

Illuminating Enclosure LEDs

To locate an enclosure by using its Unit Locator LED:

1. Select Manage > General Config > Enclosure Management.
2. If there is more than one enclosure, select the enclosure to locate.
3. Click Illuminate Locator LED.

   The top LED on the enclosure's right ear blinks so you can determine the enclosure's physical location.
4. To stop blinking the LED, click Turn Off Locator LED.

Changing the Enclosure Polling Rate

You can change the interval at which the storage system polls the EC (EMP) for status changes. Typically, use the default rate.

- 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 enclosure polling rate:

1. Select Manage > General Config > Enclosure Management.
2. In the Advanced Enclosure Options panel, click Advanced Options.
3. Change the polling rate.

   The default is 5 seconds.
4. Click Change EMP Poll Rate.

Viewing Expansion Enclosure Versions

To view the firmware version (revision) and type of expansion modules in each expansion enclosure connected to the system, do either of the following:

- Select Manage > Update Software > Enclosure Firmware > Show Enclosures.

  Information is displayed for all expansion enclosures connected to the system:

  - Enclosure WWN – The expansion enclosure's node World Wide Name.
  - Address – The channel and loop ID of the expansion module.
  - Manufacturer – The expansion module manufacturer.
  - Model – The expansion module model number.
  - Revision – The revision code for the firmware currently in the expansion module.

- Select Manage > Update Software > Enclosure Firmware > Show Enclosure Types.

  Information is displayed for each type of expansion enclosure connected to the system:

  - Vendor/Model – The vendor and model of the expansion module.
  - Firmware Revision – The revision code for the firmware currently in the expansion module.
  - Total Number of this Type – The number of expansion modules that have the same vendor, model, and firmware revision. For example, two identical expansion modules with different firmware revisions are considered to be different types.

Updating Expansion Enclosure Firmware

You can update expansion enclosure firmware by loading a firmware update file obtained from the enclosure vendor.

Caution – Updating enclosure firmware causes all disk drives to be temporarily inaccessible. Stop all I/O to virtual disks before performing this operation.

To update expansion enclosure firmware:

1. Select Manage > Update Software > Enclosure Firmware > Update Firmware.
2. Select the type of expansion modules to update.

   Expansion modules that have the same manufacturer, model, and firmware revision are considered the same type. For example, two identical expansion modules with different firmware revisions are considered to be different types.
3. Click Select Type And Continue.

   Enclosure processors of the type you selected are listed and the following information is displayed for each enclosure:

   - Device WWN – The expansion enclosure's node World Wide Name.
   - Address – The channel and loop ID of the expansion module.
   - Manufacturer – The expansion module manufacturer.
   - Model – The expansion module model number.
   - Rev – The revision code for the firmware currently in the expansion module.

   If more than two enclosure modules are listed, a Select All check box is displayed.
4. Select the enclosure modules to update.
5. Click Continue.
6. Click Browse to select the firmware update file.
7. Click Load Device Firmware File.
8. To start the firmware update, click Start Firmware Update.

   To cancel the firmware update, click Cancel.

   The file is transferred to the RAID controller where it is temporarily stored prior to download to the enclosure. Once the firmware update process has started, a page shows the update progress of each enclosure, including when the firmware update has completed successfully.

This operation can take several minutes to complete. During the update, the following operations are blocked so that they do not interfere with the update:

- Updating controller software (buffer interference)
- Saving logs to a file (buffer interference)

9. When processing is complete, verify that the proper firmware revision is reported for each updated enclosure.

Stopping or Aborting an Expansion Enclosure Firmware Update

The Stop Device Firmware Update button stops the update operation at the next point that will leave all expansion modules in a clean state. If you click this button while the file is being downloaded to the controller, the download stops in a few seconds and no expansion modules are updated. If you click this button after download to the enclosures has started, the process is stopped after the update to the current expansion module is complete. No updates that are already done or started are undone. It can take several minutes for a stop operation to complete.

Caution – The Abort Device Firmware Update button immediately stops the firmware update and leaves the expansion module in an unknown, possibly unusable state. If you choose this option, wait two minutes to enable the expansion module to possibly finish writing to its nonvolatile memory.

CHAPTER 5

Monitoring System Status

This chapter describes how to use RAIDar to monitor your system to ensure that its components are working properly. Topics covered in this chapter are:

Displaying Status Information

RAIDar includes many status pages that enable you to monitor the status of your system, virtual disks, and disk drives. The top panel on many status pages includes an icon for each virtual disk with information about the selected virtual disk below it. For information about the virtual disk icons, see Table 1-3.

Status Summary

Unless you have changed the On Manage Login preference (see "Configuring Preferences" on page 24), you see the Status Summary page when you log in to RAIDar. This page includes:

- Status Message panel – Briefly describes the storage system's overall status. If a warning or error condition exists, a message specifies to see a certain RAIDar page for details.
- Virtual Disk Overview panel – Shows information about existing virtual disks. To see more detail about a virtual disk, click its icon. For a description of the virtual disk icons, see Table 1-3.
- Hardware Status panel – Shows the status of each controller module and the overall status of system enclosures. To see more detail, click a status link.

- System Panel – Shows system status and which RAID controller you are connected to. For a description of the status icons, see "System Panel" on page 22. To see more detail, click a status link.

To display the Status Summary page from another RAIDar page:

- Select Monitor > Status > Status Summary.

Virtual Disk Status

You can view detailed information about a virtual disk's status, including its disk drives and volumes. You can display virtual disk status information in two ways.

To view virtual disk (vdisk) status from the menu:

- Select Monitor > Status > Vdisk Status.

To view virtual disk status from another page's Virtual Disk Overview panel:

- Click the virtual disk's icon.

Using either method, the Virtual Disk Status page is displayed. For a description of the virtual disk icons, see Table 1-3. Details about the selected virtual disk are displayed in four panels.

The Virtual Disk Status Details panel shows the following information:

- RAID Level – Either RAID 0, 1, 3, 5, 6, 10, 50, or Non-RAID.
- Virtual Disk Size – Virtual disk size in Gbyte.
- Virtual Disk Status – Either Online, Offline, Critical, or Fault Tolerant.
- Number Of Drives – Number of drives in the virtual disk when fault tolerant. For example, if a three-drive RAID 5 virtual disk loses one drive, the number still shows 3.
- Number Of Spares – Number of spares assigned to the virtual disk.
- Number Of Volumes – Number of volumes in the virtual disk.
- Virtual Disk Name – Name assigned to the virtual disk.
- Virtual Disk Serial Number – Unique number assigned by the owning controller.
- Virtual Disk Owner – Controller that owns the virtual disk.
- Chunk Size – Amount of contiguous data that is written to a virtual disk member before moving to the next member of the virtual disk.

- Date Created – Date when the virtual disk was created.
- Utility – Name of any utility running on the virtual disk, or None. The utility status is shown in the virtual disk panel.

The Virtual Disk Drive List panel shows the following information:

- Status – Up if operational or Down if failed
- Size – Drive size in Gbyte
- Manufacturer – Drive manufacturer
- Model – Drive model number
- Revision – Drive firmware revision
- Node WWN – Drive node World Wide Name
- Serial Number – Drive serial number
- Encl.Slot – Enclosure number and slot number containing the drive
- Encl Name – Name of the enclosure containing the drive

The Dedicated Spares For Selected Virtual Disk panel is displayed only if spares are assigned to the virtual disk. This panel shows the following information:

- Status – Up if operational or Down if failed
- Size – Drive size in Gbyte
- Manufacturer – Drive manufacturer
- Model – Drive model number
- Revision – Drive firmware revision
- Node WWN – Drive node World Wide Name
- Serial Number – Drive serial number
- Encl.Slot – Enclosure number and slot number containing the drive
- Encl Name – Name of the enclosure containing the drive

The Volume Information Panel shows the following information.

- Name – Name assigned to the volume
- LUN – Logical unit number assigned to use for all connected hosts, or None
- Size – Volume size in Mbyte

Host Port Status

This section describes status information shown for host ports on Fibre Channel (FC) controller modules or on iSCSI controller modules.

FC Host Port Status

The Host Port Status page shows a graphical representation of the host ports on each controller, including a color-coded status for each port.

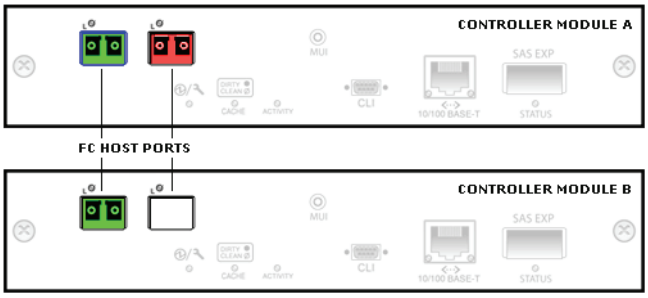

Figure 5-1 FC Host Port Status Example

To display host port status information:

1. Select Monitor > Status > Advanced Settings > Host Port Status.

   The status of each port is determined by the following color code:

   - Green – Host link is up
   - Red – Host link is down
   - White – Port is unused and does not contain an SFP

2. To view information about a specific port, click the port.

   The selected port is outlined in blue.

   Details in the lower part of the panel vary depending on the selected port's status. Some values are assigned only when the host link is up. An asterisk (*) indicates a value that will take effect after loop initialization.

   - Host Port Status Details – Selected controller and port number.
   - SFP Detect – SFP is present or not present. An SFP is used to connect the FC host port through an FC cable to another FC device.
   - Receive Signal – Signal is present or not present.
   - Link Status – Link is up (active) or down (inactive).
   - Signal Detect – Signal is detected or no signal.

- Topology – One of the following values:
  - Point-to-Point
  - Loop, if the loop is inactive
  - Private Loop, if the port is directly attached to a host
  - Public Loop, if the port is attached to a switch

  To change this setting, see "Setting FC Host Port Topology" on page 39.
- Speed – 2 Gbit/sec or 4 Gbit/sec. To change this setting, see "Setting FC Host Port Link Speed" on page 36.
- FC Address – 24-bit FC address, or Unavailable if the FC link is not active.
- Loop ID – (Loop topology only) Current and requested loop ID values. To change this setting, see "Setting FC Host Port Loop IDs" on page 37.
- Node WWN – Controller module node World Wide Name.
- Port WWN – Port World Wide Name.

iSCSI Host Port Status

The Host Port Status page shows a graphical representation of the host ports on each controller, including a color-coded status for each port.

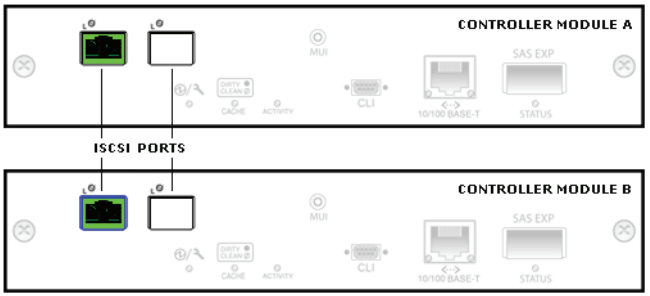

Figure 5-2 iSCSI Host Port Status Example

To display host port status information:

1. Select Monitor > Status > Advanced Settings > Host Port Status.

   The status of each port is determined by the following color code:

   - Green – Host link is up
   - White – Host link is down

2. To view information about a specific port, click the port.

   The selected port is outlined in blue.

   Details in the lower part of the panel vary depending on the selected port's status.

   - iSCSI Port Status Details – Selected controller and port number
   - Link Status – Link is up (active) or down (inactive)
   - Qualified Name – iSCSI qualified name (IQN)
   - Link Speed – 1 Gbit/sec
   - IP Version – IP addressing version; 4 for IPv4
   - IP Address – Port IP address
   - IP Mask – Port IP subnet mask
   - IP Gateway – Port gateway IP address
   - Service Port – iSCSI port number
   - Hardware Address – Port MAC address

Disk Drive List

To view information about all disk drives in the storage system:

- Select Monitor > Status > Advanced Settings > Disk Drive List.

This page shows the total number of drives that:

- Are installed
- Are available for use
- Contain leftover metadata from a previous virtual disk
- Are down (failed)

It also shows the following information about each drive:

- Status – Up if operational or Down if failed.
- Size – Drive size in Gbyte.
- Speed – Data transfer rate in Gbit per second.
- Manufacturer – Drive manufacturer.
- Model – Drive model number.
- Revision – Drive firmware revision.
- Node WWN – Drive node World Wide Name.
- Serial Number – Drive serial number.
- Encl.Slot – Enclosure number and slot number containing the drive.
- Belongs To Virtual Disk – Different information depending on the drive's status:
  - If used in a virtual disk, the virtual disk name.
  - If used as a spare, the type of spare.
  - If unused, Available.
  - If contains leftover metadata, Leftover. A Manage user can return leftover drives to available status; see "Clearing Metadata From a Disk Drive" on page 130.
- Enclosure Name – Name of the enclosure containing the drive.

Disk Drives by Enclosure

You can view a graphical representation of the disk drives and enclosures in the system. In this representation, drives in a virtual disk are the same color, which differs for each virtual disk. If the graphical representation isn't displayed, see "Enclosure View is Unavailable" on page 151.

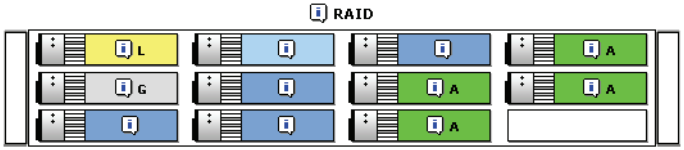

Figure 5-3 Enclosure View Example

Note – For a multi-enclosure system, you can configure the enclosure display order by setting rack positions; see "Entering Enclosure Information" on page 137.

To view disk drives by enclosure:

1. Select Monitor > Status > Enclosure View.

   Drives are shown with the following color and text codes:

   - White – Drive not installed
   - Yellow with 'L' – Leftover drive; not in use and contains old metadata
   - Green with 'A' – Available drive; not in use
   - Gray with 'G' – Global spare
   - Other color – Virtual disk member drives, including vdisk spares
   - 'SP-A' or 'SP-B' – Drive has a single-port connection to controller A or B

2. To see details for an enclosure or drive, pause the cursor over that device's information icon. If you click the icon, the information remains shown until the cursor passes over a similar icon.

   For an enclosure the following information is displayed:

   - Enclosure Status – Whether the enclosure is OK or has an error
   - Name – Name assigned to the enclosure
   - Mfr – Enclosure manufacturer
   - Model – Enclosure model number
   - Version – Expander Controller software version
   - WWN – Enclosure node World Wide Name

For a drive the following information is displayed.

- Drive Status – Up if operational or Down if failed.
- Encl – Number of the enclosure containing the drive.
- Slot – Number of the drive slot in the enclosure.
- Mfr – Drive manufacturer.
- Model – Drive model number.
- Size – Drive size in Gbyte.
- Type – Drive architecture (SATA or SAS).
- WWN – Drive node WWN.
- Serial # – Drive serial number.
- Part Of Vdisk – Different information depending on the drive's status:
  - If used in a virtual disk, the virtual disk name.
  - If used as a spare, the type of spare.
  - If unused, Available.
  - If contains leftover metadata, Leftover. A Manage user can return leftover drives to available status; see "Clearing Metadata From a Disk Drive" on page 130.

Enclosure View is Unavailable

If the system is unable to communicate with a drive enclosure or its drives, a message notifies you that there is an invalid enclosure or drive configuration, or that the display of drives by enclosure is unavailable.

In this situation the drive information is displayed in list format without enclosure information instead of in graphical format. This can occur for the following reasons:

- A drive has been recently added or removed, and the data reported by the individual drives has not yet been fully synchronized with the configuration data reported by the EMP.
- All paths to the enclosure have been disconnected.
- Enclosure polling has been suspended (the rate is set to zero), which can be done only by a service technician.

In these situations, there might be a short period where the displayed information is unpredictable. However, the display corrects itself once the environment stabilizes.

LAN Information

You can view Ethernet and IP information for each controller, and information about the system. To change the LAN settings, see "Configuring Ethernet Management Ports" on page 42. To change the system information, see "Setting System Information" on page 34.

To view LAN information:

- Select Monitor > Status > Advanced Settings > LAN Information.

  The following information is displayed:

  - Ethernet Address – Each controller's unique Media Access Controller (MAC) hardware address, also known as the physical address.
  - Ethernet Link – Each controller's Ethernet link status: Active (operational) or Link Down (not operational).
  - IP Address – Each controller's Ethernet management port IP address. The default is 10.0.0.2 for controller A and 10.0.0.3 for controller B.
  - IP Subnet Mask – Each controller's IP subnet mask. The default is 255.255.255.0.
  - IP Gateway – Each controller's gateway IP address. The default is 10.0.0.1.
  - Source For IP Address – Manual or DHCP.
  - Telnet Timeout – The number of idle minutes before the Telnet session times out. The allowed values are 0–255 minutes, where 0 means no timeout. The default is 60 minutes.
  - System Name – Name of the system as seen by other systems on the network. The default is Uninitialized Name.
  - System Contact – Name of a contact person responsible for the system. The default is Uninitialized Contact.
  - System Location – Location of the system. The default is Uninitialized Location.
  - System Information – Additional information about the system. The default is Uninitialized Info.

Module Status

You can view summary status information for each controller module and all enclosures in the storage system. More detail is available on other Monitor pages.

To view module status:

- Select Monitor > Status > Module Status.

  The Rear Panel Chassis View shows the back of the controller enclosure and the current status of power-and-cooling modules and controller modules. Failed modules are displayed in red.

  The following information is displayed for each controller module:

  - Present? – Yes if installed or No if not installed. Click Yes to see software and hardware version details on the Controller Versions page.
  - Primary Status – One of the following:
    - Online – The module is present and operating correctly.
    - Offline – The module is either not installed or has been taken out of service by the system or by user request.
    - Failed – A hardware or system error has been detected and the module is not functioning correctly.
  - Secondary Status – Additional status information.
  - Serial Number – Controller module serial number.
  - HW Version – Controller module hardware version and CPLD version.

  Summary information is displayed about the status of system enclosures. Any abnormal conditions are displayed, indicating the enclosure where the problem is, the element within the enclosure that is reporting the condition, and the status of the element. The following categories are monitored:

  - Power Supplies – Indicates the presence of critical or warning conditions in any enclosure power supply.
  - Cooling – Indicates the presence of critical or warning conditions in any cooling fan.
  - Temperature Sensors – Indicates the presence of critical or warning conditions detected by an enclosure temperature sensor.
  - Voltage Sensors – Indicates the presence of critical or warning conditions detected by an enclosure voltage sensor.
  - Drives – Indicates the presence of enclosure-detected critical or warning conditions in disk drives. Does not include “not installed” conditions.

The status is OK when there are no critical or warning conditions for the element type. If no enclosure polling data is available, a message is displayed stating this. For information about a critical or warning condition, view the event log; see "Displaying the Event Log" on page 165.

Controller Versions

You can view the software, hardware, and other version information for each controller module. During normal operation, all software versions should be the same on both controller modules. Software versions might differ briefly while you are updating the firmware on each controller module.

To view version information:

- Select Monitor > Status > Advanced Settings > Controller Versions.

  Version information is displayed in the following four panels.

  The Storage Controller Code Versions panel shows the following information:
  - Code Version – Storage Controller software version
  - Memory Controller – Memory controller software version
  - Loader Version – Storage Controller loader software version

  The Management Controller Code Versions panel shows the following information:
  - Code Version – Management Controller software version
  - Loader Version – Management Controller loader software version

  The Enclosure Controller Code Versions panel shows the Expander Controller software version.

  The RAID Controller Hardware Versions panel shows the following information.
  - Hardware Version – Board version number
  - CPU Type – Type of RAID controller processor:
    - Celeron 566MHz, used in a standard controller module
    - Pentium III 700MHz, used in a "Turbo" controller module
  - CPLD Version – Version of the complex programmable logic device (CPLD)
  - Host Interface Module Model – Model number of the host interface module within a controller module
  - Host Interface Module Version – Version of the host interface module within a controller module
  - Cache Memory Size – Cache memory size in Mbyte

FRU Information

You can view information about field-replaceable units (FRUs) other than drive modules in an enclosure. For information about installed drive modules, see “Disk Drive List” on page 149.

To view FRU information:

1. Select Monitor > Status > Advanced Settings > FRU Information.

   The drive enclosure panel shows all enclosures in the system and the status of each enclosure.

2. Select an enclosure.

   The following panels are displayed:

   - Enclosure Midplane – Shows information about the chassis-and-midplane FRU
   - Enclosure Controller A – Shows information about the controller module or expansion (I/O) module FRU in the upper slot
   - Enclosure Controller B – Shows information about the controller module or expansion (I/O) module FRU in the lower slot
   - Enclosure Power Supply 1 – Shows information about the power-and-cooling module FRU in the left slot (with respect to the back of the enclosure)
   - Enclosure Power Supply 2 – Shows information about the right power-and-cooling module FRU in the right slot (as viewed from the back of the enclosure)

Enclosure Status

You can view status information about each enclosure component. To change the enclosure name, location, and rack position values, see “Entering Enclosure Information” on page 137.

To view enclosure status:

1. Select Monitor > Status > Enclosure Status.

   The drive enclosure panel shows all enclosures in the system and the status of each enclosure.

2. Select an enclosure.

   The Enclosure Details panel shows the following information about the selected enclosure:

   - Name – Name assigned to the enclosure.
   - Vendor – Enclosure manufacturer.
   - Location – Enclosure location, if set.
   - Status – Specifies whether the enclosure is OK or has an error.
   - Misc – Enclosure ID, which is 0 for a controller enclosure and increments from 1 for attached expansion enclosures.
   - World Wide Name – Enclosure node World Wide Name.
   - Model – Enclosure model number.
   - Rack:Position – Assigned rack number and position of the enclosure within the rack, or 0:0 if not set. Position 1 is the top and 16 is the bottom.
   - Firmware Version – Version of the Expander Controller, which performs SES functions.
   - CPLD Revision – Revision of the complex programmable logic device (CPLD).

   The Components of Enclosure panel shows the status and other available information about non-drive components in the enclosure:

   - Component – Component name and location (as viewed from the back of the enclosure)
   - Status – OK or error
   - Details – Status details such as current temperature and voltage

   The Enclosure Drive List panel shows the slot number, node WWN, and status of each installed disk drive. If the system is unable to communicate with an enclosure or its drives, messages appear as described in “Enclosure View is Unavailable” on page 151.

Temperature Status

As an Advanced user, you can view the current temperature status of each temperature sensor in each controller module. Each controller has six temperature sensors. To change the temperature display mode, see “Configuring Preferences” on page 24.

To view temperature status:

- Select Monitor > Status > Advanced Settings > Temperature Status.

  The following panels are displayed:

  - Temperature Status – Shows each sensor's temperature value, status, and normal and warning operating ranges. For CPU and FPGA sensors the critical shutdown range is also shown.
  - Common Temperature Sensors Status – Shows each power supply sensor's temperature value, status, and normal operating range.

  For information about what to do when temperature errors occur, see the *Troubleshooting Guide*.

Power Status

As an Advanced user, you can view the current status of power supplies and super-capacitor packs in each enclosure. Each power supply has a 12-volt, 5-volt, and 3.3-volt sensor. The super-capacitor pack in each controller module provides backup power for controller cache.

To view power status:

- Select Monitor > Status > Advanced Settings > Power Status.

  The following panels are displayed:

  - Power Sensors Status – Shows the capacitor pack's total voltage (the sum of its individual cell voltages), status, and normal operating range; the voltage, status, and normal operating range of each cell in the capacitor pack; and the charge level of the capacitor pack.
  - Common Power Sensors Status – Shows the voltage, status, and normal operating range for each power supply sensor.

  For information about what to do when power errors occur, see the *Troubleshooting Guide*.

LUN Information

To view information about all volumes in the system:

- Select Monitor > Status > Advanced Settings > LUN Information.

The Volume LUN Information panel shows the following information.

- IOM – Owning controller (shown in dual-controller mode)
- Node WWN – Owning controller's node World Wide Name
- LUN – Default LUN to use when mapping the volume to a host if no other LUN is specified
- Vdisk Name – Name of the virtual disk that the volume is part of
- Vol Num – Volume number in the virtual disk
- Volume Name – Name assigned to the volume
- Volume Size – Volume size in Gbyte
- Read-Ahead Cache Size – Specifies the read-ahead cache size setting: Default, Disabled, a specific size in Kbyte or Mbyte, or Maximum
- Write-Back Cache Enable – Specifies whether write-back cache is enabled or disabled

To change the virtual disk owner, see "Changing Virtual Disk Ownership" on page 74. To change the virtual disk name, see "Changing a Virtual Disk Name" on page 75.

To change volume information, see "Managing Volumes" on page 80.

To change the read-ahead cache size, see "Changing a Volume's Read-Ahead Cache Settings" on page 92. To change the write-back cache setting, see "Changing a Volume's Write-Back Cache Setting" on page 94.

Misc Configuration

As an Advanced user, you can view the following categories of configuration settings: general, RAID controller, EMP, security access to services, and user preferences.

To view miscellaneous configuration settings:

- Select Monitor > Status > Advanced Settings > Misc Configuration.

  The information is displayed in five panels.

  The General Configuration Status panel shows the following information:

  - Background Scrub – Shows whether virtual disks are automatically analyzed to find disk-drive defects. The default is Disabled. If Enabled, disk drives associated with virtual disks are continuously analyzed and information about disk-drive defects is reported and is stored in disk-drive metadata. If Disabled, virtual disks are not automatically scrubbed. For more information, see "Enabling and Disabling Background Scrub for Disks" on page 186.
  - Partner Firmware Upgrade – Shows whether the system automatically upgrades firmware on one controller when a newer version of firmware is loaded on the partner controller. The default is Enabled. If Enabled, the partner controller is automatically upgraded. If Disabled, the partner controller must be upgraded manually.

    If directed by a service technician, a Manage user can disable partner firmware upgrade on the Manage > General Config > System Configuration page.
  - Utility Priority – Shows the priority at which all system utilities run when there are active I/O operations competing for the controller's CPU. The setting can be High (default), Medium, or Low. The default is High. For more information, see "Changing the Utility Priority" on page 181.
  - Host Control Of Cache – Shows whether hosts are prevented from using SCSI `MODE SELECT` commands to change the system's write-back cache setting. Some operating systems disable write cache. The default is Enabled. If host control is Disabled, the host cannot modify the cache setting. For more information, see "Controlling Host Access to the System's Write-Back Cache Setting" on page 187.
  - Dynamic Spare – Shows whether the system can automatically take a properly sized available drive to reconstruct a virtual disk when no spares are designated. The default is Disabled. For more information, see "Managing Dynamic Spares" on page 77.

- SMART – Shows whether Self-Monitoring, Analysis, and Reporting Technology (SMART) settings for all drives in the system can be changed. The setting can be Enabled, Disabled, or Don't Modify. The default is Enabled. For more information, see“Enabling or Disabling SMART Changes” on page 131.

The RAID Controller Status panel shows the following information for each controller:

- Hardware Status – Shows one of the following statuses:
  - Redundant Operation – Current controller and partner controller are online
  - Only Operational Controller – Current controller is online and partner controller, if installed, is offline
  - Not Up – Current controller is offline
- Write-Back Cache Status – Shows status of the controller’s write-back cache based on whether cache backup power is operating properly. If the cache backup power is faulty, the write-back is disabled.

The EMP Status panel shows the poll rate, which is the interval in seconds at which the system polls each enclosure’s EC (EMP) for status changes. To change this setting, see “Changing the Enclosure Polling Rate” on page 138.

The Security Access To Services panel shows the following information. To change these settings, see “Configuring Network Management Services” on page 46.

- FTP – Shows whether `ftp` access is enabled, which provides an alternate way to to update system software. The default is Disabled.
- Telnet – Shows whether Telnet access is enabled, so the CLI can be used to manage the system. The default is Enabled.
- HTTP – Shows whether `http` access is enabled, so RAIDar can be used to manage the system. The default is Enabled.
- SNMP – Shows whether SNMP is enabled, so the system can be remotely monitoring through your network, is enabled. The default is Enabled.
- Internet Debug – (Advanced users) Shows whether this diagnostic option, which can be used for technical support, is enabled. The default is Disabled.

The User Preferences panel shows the following information. To change these settings, see “Configuring Preferences” on page 24.

- Page Refresh Rate – Shows how often RAIDar refreshes its pages based on the speed of your computer and Ethernet connection. The setting can be Fast, Medium, or Slow. The default is Fast.
- Auto-Logout Timeout – The number of idle minutes before RAIDar session times out and requires you to log back in, or “No timeout.” The default is 30 minutes.
- Temperature Display Mode – Fahrenheit or Celsius for all temperature status displays. The default is Celsius.

Expander Status

Each controller module and expansion module has an Expander Controller (EC) that manages the module's SAS expander. A SAS expander has 24 serial ports (PHYs) that are used for communication between the ECs and the disk drives.

The SAS expander uses four types of PHYs:

- Ingress PHYs (4) to communicate with upstream devices such as the enclosure's EC
- Disk PHYs (12) to communicate with the enclosure's disk drives
- Egress PHYs (4) to communicate with downstream devices such as attached expansion enclosures
- Inter-expander PHYs (4), in a controller module only, to communicate with the expander in the partner controller module

When the storage system's PHY isolation feature is enabled, PHYs are monitored for faults and a PHY is automatically disabled if it experiences too many faults.

For a selected enclosure you can view the status of PHYs managed by each EC. The status information can identify where faults have occurred in the communication path.

To view expander status information:

1. Select Monitor > Status > Advanced Settings > Expander Status.
2. Select an enclosure.

   The information is displayed in three panels.

   The Enclosure Details panel shows the following information about the selected enclosure:

   - Name – Name assigned to the enclosure.
   - Vendor – Enclosure manufacturer.
   - Location – Enclosure location, if set.
   - Status – Specifies whether the enclosure is OK or has an error.
   - Misc – Enclosure ID, which is 0 for a controller enclosure and increments from 1 for attached expansion enclosures.
   - World Wide Name – Enclosure node World Wide Name.
   - Model – Enclosure model number.
   - Rack:Position – Assigned rack number and position of the enclosure within the rack, or 0:0 if not set. Position 1 is the top and 16 is the bottom.
   - Firmware Version – Version of the EC, which performs SES functions.

   The Phy Isolation Details panel shows the following settings for each EC:

   - Phy Isolation – Shows whether all PHYs in the expander are monitored for faults and automatically isolated if too many faults are detected. The default is Enabled.
   - Auto Recover – Disabled.
   - Monitoring Period – Specifies how often the EC checks each PHY and determines whether it should be isolated. The default is 100 milliseconds.

   The Expander Controller Phy Detail panel shows the following information about each PHY in each EC:

   - Status – Specifies one of the following:
     - OK – The PHY is healthy.
     - Error – The PHY experienced an unrecoverable error condition or received an unsupported PHY status value.
     - Disabled – The PHY has been disabled by a Diagnostic Manage user or by the system.
     - Non-Critical – The PHY is not coming to a ready state or the PHY at the other end of the cable is disabled.
     - Not Used – The module is not installed.

- Type – Specifies one of the following:
  - Disk – Communicates between the expander and a disk drive.
  - Inter-Exp – (Controller module only) Communicates between the expander and the partner's expander.
  - Ingress – Communicates between the EC and the expander.
  - Egress – Communicates between the expander and an expansion port or SAS Out port.
- State – Specifies whether the PHY is enabled or disabled.
- ID – Identifies a PHY's logical location within a group based on the PHY type. Logical IDs are 0–11 for disk PHYs and 0–3 for inter-expander, egress, and ingress PHYs.
- Details – Pause the cursor over or click the information icon to view a popup with more information. If you click the icon, the information remains shown until the cursor passes over a similar icon.
  - Status – The same status value shown in the panel's Status field.
  - Physical Phy ID – Identifies a PHY's physical location in the expander.
  - Type – The same type value shown in the panel's Type field.
  - Phy Change Count – Specifies the number of times the PHY originated a BROADCAST (CHANGE). A BROADCAST (CHANGE) is sent if doubleword synchronization is lost or at the end of a Link Reset sequence.
  - Code Violation Count – Specifies the number of times the PHY received an unrecognized or unexpected signal.
  - Disparity Error Count – Specifies the number of doublewords containing running disparity errors that have been received by the PHY, not including those received during Link Reset sequences. A running disparity error occurs when positive and negative values in a signal don't alternate.
  - CRC Error Count – In a sequence of SAS transfers (frames), the data is protected by a cyclic redundancy check (CRC) value. This error count specifies the number of times the computed CRC does not match the CRC stored in the frame, which indicates that the frame might have been corrupted in transit.
  - Inter-Connect Error Count – Specifies the number of times the lane between two expanders experienced a communication error.
  - Lost Doubleword Count – Specifies the number of times the PHY has lost doubleword synchronization and restarted the Link Reset sequence.
  - Invalid Doubleword Count – Specifies the number of invalid doublewords that have been received by the PHY, not including those received during Link Reset sequences.

- Reset Error Count – Specifies the number of times the expander performed a reset.
- Phy Disabled – Specifies whether the PHY is enabled (True) or disabled (False).
- Fault Reason – A coded value that explains why the EC isolated the PHY. If the PHY is active, this value is 0x0.

For example, assume that a SAS cable connects Enclosure 0's "out" port to Enclosure 1's "in" port. If the connection has no faults then PHYs associated with each port have OK status, as shown in the following figure.

Enclosure 0

OK	Egress	Enabled	0	
OK	Egress	Enabled	1	
OK	Egress	Enabled	2	
OK	Egress	Enabled	3	

Enclosure 1

OK	Ingress	Enabled	0	
OK	Ingress	Enabled	1	
OK	Ingress	Enabled	2	
OK	Ingress	Enabled	3	

However, if there is a fault in the SAS cable or either of the SAS connectors then associated PHYs have Non-Critical status as shown in the following figure.

Enclosure 0

Non-critical	Egress	Enabled	0	
Non-critical	Egress	Enabled	1	
Non-critical	Egress	Enabled	2	
Non-critical	Egress	Enabled	3	

Enclosure 1

Non-critical	Ingress	Enabled	0	
Non-critical	Ingress	Enabled	1	
Non-critical	Ingress	Enabled	2	
Non-critical	Ingress	Enabled	3	

Displaying the Event Log

The system's event log contains important information about the status of the system, virtual disks, and disk drives. Check it regularly to monitor the status of your system. For more information about specific warning and error events and specific disk and port errors, refer to the *Troubleshooting Guide*.

Some of the key warning and error events included in the event log during operation include the following:

- Disk detected error
- Disk channel error
- Drive down
- Virtual disk critical
- Virtual disk offline
- Temperature warning
- Temperature failure (this leads to a shutdown which is also logged)
- Voltage warning
- Voltage failure (this leads to a shutdown which is also logged)

The event log stores the most recent events with a time stamp next to them with one-second granularity.

Note – If you are having a problem with the system or a virtual disk, check the event log before calling technical support. Event messages might enable you to resolve the problem.

You can save the event log to a file; see "Saving Log Information to a File" on page 166.

To view the event log:

1. Do one of the following:
   - Click the EVENT LOG icon in the System Panel.
   - Select Monitor > Status > View Event Log.

   The Event Log is displayed.

2. Click one of the following buttons in the Select Event Table To View panel to see the corresponding events:

Button	Description
Controller A & B Events	Shows all events for both controllers.
Controller A & B Critical/ Warning Events	Shows only critical and warning events for both controllers.
Controller A Events	Shows events logged by controller A.
Controller B Events	Shows events logged by controller B.

The Events panel shows up to 200 events for a single controller or 400 events for both controllers. The events display in reverse chronological order (the most recent first). The following information is displayed:

Field	Description
C/W	Either C for critical, W for warning, or blank for informational.
Date/Time	Month, day, and time the event occurred.
EC	Event code that assists service personnel when diagnosing problems.
ESN	Event Serial Number. The prefix (A or B) indicates which controller logged the event.
Message	Information about the event.

Saving Log Information to a File

You can save the following types of log information to a file:

- Device status summary, which includes basic status and configuration information for the system.
- Event logs from both controllers when in active-active mode.
- Debug logs from both controllers when in active-active mode.
- Boot logs, which show the startup sequence for each controller.
- Up to four critical error dumps from each controller. These will exist only if critical errors have occurred.
- Management Controller traces, which trace interface activity between the controllers' internal processors and activity on the management processor.

To save log information to a file:

1. Select Manage > Utilities > Debug Utilities > Save Logs To File.
2. Type contact information and comments to include in the log information file.

   Contact information provides the support representatives who are reviewing the file a means to identify who saved the log. Comments can explain why the logs are being saved and include pertinent information about system faults.
3. Under File Contents, select the logs to include in the file.

   By default, all logs are selected.

Note – Select logs judiciously. Gathering log data can be a lengthy operation, especially if the system is performing I/O.

---

4. Click Generate Log Information.

   When processing is complete, a summary page is displayed.
5. Review the summary of contact information, comments, and selected logs.
6. Click Download Selected Logs To File.
7. If prompted to open or save the file, click Save.
8. If prompted to specify the file location and name, do so using a `.logs` extension.

   The default file name is `store.logs`. If you intend to capture multiple event logs, be sure to name the files appropriately so that they can be identified later.

Note – If you are using Firefox and have a download directory set, the file is automatically saved there.

Setting Up the Debug Log

When instructed to do so by service personnel, as an Advanced Manage user you can configure the debug log. The debug log captures data that will help service personnel locate problems within the system logic.

After you configure the debug log as instructed, you will need to perform I/O to the system or re-create the situation that is causing the fault. This populates the debug log with information that service personnel can use to diagnose the system.

Note – The debug log only collects data after you configure it. It will not contain information about any problems that occurred before you configure it.

To configure the debug log:

1. Select Manage > Utilities > Debug Utilities > Debug Log Setup.

   The Debug Log Setup page is displayed.

2. Select the debug log setup you want.
   - Standard – Used for diagnosing general problems. With minimal impact on I/O performance, it collects a wide range of debug data.
   - Fibre Channel/Performance – Used for diagnosing Fibre Channel problems. Using this option, the debug log is dedicated to collecting Fibre Channel information, with minimal impact on I/O performance.
   - Device-Side – Used for diagnosing device-side problems. It collects device failure data as well as Fibre Channel information, with minimal impact on I/O performance.
   - Device Management – Collects very verbose information, including all Configuration API (CAPI) transactions. Because this option collects a lot of data, it has a substantial impact on performance and quickly fills up the debug trace.
   - Custom Debug Tracing – Shows that specific events are selected for inclusion in the log. This is the default. If no events are selected, this option is not displayed.
3. Click Change Debug Logging Setup.
4. If instructed by service personnel, click Advanced Debug Logging Setup Options and select one or more additional types of events.

   Under normal conditions, you should not select any of these options because they have a slight impact on read/write performance.

Viewing Statistics

Viewing statistics can help you interpret performance based on configuration of an individual element of your storage solution, such as FC HBA, iSCSI Ethernet adapter, driver, SAN, or host operating system. The statistical information is useful to profile applications and their usage of a virtual disk, which can be used to determine if additional virtual disks would increase performance and what RAID level fits your needs. You can analyze the performance of the same application using different RAID levels to determine which level has the best performance. See Appendix B for more details on RAID levels.

Note – The statistics are provided as general information for your use when analyzing system performance. They are not intended for benchmarking purposes but more so to accurately track your testing and to compare with benchmark testing.

The rate statistics and cumulative statistics pages update at 60-second intervals and sampled rates become valid after two minutes. The real-time statistics page updates at 2-second intervals. Thus, real-time statistics give you an instant look at system performance, while rate and cumulative statistics average the performance numbers over a longer period.

Statistics might not be accurate across events that alter virtual disks and volumes, including additions, deletions, and component failures. After such events, or if you are monitoring performance or changing how you are using volumes, you should reset the statistics; see “Resetting Statistics” on page 176.

Rate Statistics for Virtual Disks

You can view the following I/O statistics for all virtual disks:

- The total IOPS and bandwidth for all virtual disks
- The IOPS and bandwidth for each virtual disk

To view overall rate statistics for virtual disks:

- Select Monitor > Statistics > Overall Rate Stats.

  Statistics shown are based on host-side activity in the interval since the page was last refreshed. The page automatically refreshes at a 60-second interval.

Cumulative Statistics for Virtual Disks

You can view the following I/O statistics for all virtual disks:

- The Statistics For All Virtual Disks panel shows the total number of host read and write operations, sectors read and written, and queue depth for each controller module's host ports.
- The Total For All Virtual Disks - Host Read I/O Size Histogram panel shows how many host read operations fell into a particular size range for all virtual disks. The I/O ranges are based on powers of two.
- The Total For All Virtual Disks - Host Write I/O Size Histogram panel shows many host write operations fell into a particular size range for all virtual disks. The I/O ranges are based on powers of two.
- The Host Port Queue Depth & I/O Details panel shows the activity for each port connected to a host. Queue depth is the number of host-originated commands currently queued for the port. Last I/O size is the size of the last host access in sectors. An operation that is not a read or write sets this value to zero.

To view cumulative statistics for virtual disks:

- Select Monitor > Statistics > Cumulative Stats.

Statistics shown are based on host-side activity in the interval since the page was last refreshed. The page automatically refreshes at a 60-second interval.

Rate Statistics for Volumes

You can view the following I/O statistics for a selected virtual disk:

- The total IOPS and bandwidth for all volumes in the virtual disk
- The IOPS and bandwidth for each volume in the virtual disk

To view volume rate statistics.

1. Select Monitor > Statistics > Volume Rate Stats.
2. Select the virtual disk whose statistics you want to view.

   Statistics shown are based on host-side activity in the interval since the page was last refreshed. The page automatically refreshes at a 60-second interval.

Cumulative Statistics for Volumes

You can view the following I/O statistics for volumes of a selected virtual disk:

- The Select A Virtual Disk and Volume Menu panels show all virtual disks in the system and the selected virtual disk's volume name, size, and LUN.
- The Details For Virtual Disk Volume panel shows the total number of host read and write operations, sectors read and written, and queue depth for related host ports for the selected volume.
- The Host Read I/O Size Histogram shows how many host read operations fell into a particular size range for the selected volume. The I/O ranges are based on powers of two.
- The Host Write I/O Size Histogram shows how many host write operation fell into a particular size range for the volume. The I/O ranges are based on powers of two.
- The Host Port Details For Volume panel shows the activity for the volume of each port connected to a host. Queue depth is the number of host-originated commands currently queued for the port. The last I/O size is the size of the last read or write host access in sectors. An operation that is not a read or write sets this value to zero.

To view cumulative statistics for volumes:

1. Select Monitor > Statistics > Cumulative Volume Stats.
2. Select a virtual disk
3. Select the volume whose statistics you want to view.

   Statistics shown are based on host-side activity in the interval since the page was last refreshed. The page automatically refreshes at a 60-second interval.

Real-Time Statistics for Volumes

As an Advanced user, you can view the overall performance of volumes and related ports. This information is updated at a two-second interval.

- The Select A Virtual Disk and Volume Menu panels show all virtual disks in the system and the selected virtual disk's volume name, size, and LUN.
- The Statistics For Volume panel shows the IOPS bandwidth in Mbyte per second, the number of read and write operations, and the number of sectors (512-byte blocks) read and written. All statistics are based on host-side activity.
- The Port Statistics For Selected Volume panel shows the activity for the volume for each port connected to a host. Queue depth is the number of host-originated commands currently queued for the port. Last I/O size is the size of the last host access in sectors. An operation that is not a read or write sets this value to zero.

To view real-time statistics for volumes:

1. Select Monitor > Statistics > Real-Time Volume Stats.
2. Select a virtual disk.
3. Select the volume whose statistics you want to view.

Disk Drive Error Statistics

As an Advanced user, you can view the following disk drive error statistics, which are maintained by the controller for each drive. You can clear all error statistics except for Bad Block List Size.

Field	Description
SMART Event Count	The number of SMART (Self-Monitoring, Analysis, and Reporting Technology) events that the drive recorded. These events are often used by the vendor to determine the root cause of a drive failure. Some SMART events may indicate imminent electromechanical failure.
I/O Timeout Count	The number of times the drive accepted an I/O request but did not complete it in the required amount of time. Excessive timeouts can indicate potential device failure (media retries or soft, recoverable errors)
No Response Count	The number of times the drive failed to respond to an I/O request. A high value can indicate that the drive is too busy to respond to further requests.
Spin-up Retries	The number of times the drive failed to start on power-up or on a software request. Excessive spin-up retries can indicate that a drive is close to failing.
Media Errors	The number of times the drive had to retry an I/O operation because the media did not successfully record/retrieve the data correctly.
Non Media Errors	The number of soft, recoverable errors that are not associated with drive media.
Bad Block Reassignments	The number of block reassignments that have taken place since the drive was shipped from the vendor. A large number of reallocations in a short period of time could indicate a serious condition.
Bad Block List Size	The number of blocks that have been deemed defective either from the vendor or over time due to reallocation.

To view disk drive information and error statistics:

1. Select Monitor > Statistics > Disk Error Stats.
2. Select the disk drive whose error statistics you want to view.
3. Click Show Disk Drive Error Statistics.

To clear the disk drive error data for the selected drive:

- Click Clear Selected Disk Drive Error Statistics.

To clear the disk drive error data for all drives:

- Click Clear All Disk Drive Error Statistics.

Disk Space Usage Statistics

As an Advanced user, you can view information about overall disk space usage for all disk drives in the storage system.

The following information is displayed about virtual disk space, excluding spares.

Field	Description
Volume Space	Space for user data storage.
Free Space	Space allocated for a virtual disk but not used by volumes. Free space in a virtual disk can be used to add an additional volume to that virtual disk or to expand a volume in that virtual disk. There are three ways free space can be created: • When a virtual disk is created, space that is reserved for volume expansion is free space. • If a virtual disk is expanded by adding a disk then the new capacity becomes free space. • If a volume is deleted then the freed capacity will revert to free space.

Field	Description
RAID Protection Space	Space used for mirroring or Error Correction Code (ECC) data. This is the data that enables the virtual disk to continue to function even if a disk is lost.
Backoff Space	Space reserved to compensate for minor capacity differences between disk drives so that they can be used interchangeably. The backoff value can be changed by service technicians for testing but cannot be changed through the end-user interfaces.
Not Usable Because of Different Drive Sizes	This category is displayed if a virtual disk contains different size drives. Because the usable amount of space on any disk drive in a RAID virtual disk is equal to the size of the smallest disk drive, space on larger disk drives is unusable. For example, if a virtual disk contains 500-Gbyte disks and a 250-Gbyte disk, half of the space on each larger disk is unusable.

The following information is displayed about spares and unused space.

Field	Description
Virtual Disk Spare Space	Space on spare disk drives that are designated for use by a specific virtual disk.
Global Spare Space	Space on spare disk drives that are designated for use by any virtual disk.
Available Drive Space	Space on unassigned disk drives that are available for creating new virtual disks, for expanding virtual disks, or for use as spares.
Other space	This category is displayed if a disk drive was previously a member of a virtual disk and contains old metadata, making it a Leftover. When the metadata is cleared, the drive becomes Available. See “Clearing Metadata From a Disk Drive” on page 130.

The bottom of the page shows the total disk space and a color-coded bar representing the relative sizes of each space category.

To view disk usage:

- Select Monitor > Statistics > Disk Usage.

Resetting Statistics

You should reset statistics when you are monitoring performance, when you change how you are using volumes, or when an event occurs that alters virtual disks and volumes, including additions, deletions, and component failures.

As an Advanced user, you can reset to zero either or both of the following:

- All virtual disk and volume statistics
- All controller disk-drive error statistics, which are maintained by the controller for each disk drive

To reset statistics for specific drives only, see "Disk Drive Error Statistics" on page 173.

Note – You cannot reset port queue depth and last I/O size, which always show the current values.

To reset statistics:

1. Select Monitor > Statistics > Reset All Statistics.
2. Click the button for the statistics you want to reset.

   A message is displayed indicating whether the reset succeeded.

Displaying Notification Events

The Show Notification Events panel shows events that have occurred that were selected for Visual Notification. This panel specifies how many notification events are pending and shows up to a configured maximum number of events. To change the maximum number of events to show, see “Configuring Visual Alerts” on page 49.

To show visual events:

1. Select Monitor > Status > Show Notification.

   If events have occurred, the following information is displayed:

   - Severity Level – Critical, Warning, or Info.
   - Date/Time – Year, month, day, and time when the event occurred.
   - Event Code – Event code that assists service personnel when diagnosing problems.
   - Event Serial Number – An identifier for the event. The prefix (A or B) indicates which controller logged the event.
   - Message – Information about the event.
   - Alert Method – Icons representing the notification methods configured for this event type:
     - – Visual alert
     - – Email alert
     - – SNMP trap

2. Acknowledge the events by clicking one of the following buttons:
   - Acknowledge Above Events – Acknowledges the events on the page and if there are more notification events, these subsequently are displayed.
   - Acknowledge All Events – Acknowledges all of the events without necessarily showing them on the page.

Additional Status Information

The following additional status information will help you monitor the system:

- Using the debug log as explained in the *Troubleshooting Guide*.
- LED status descriptions in the *Troubleshooting Guide*.

CHAPTER 6

Additional Utilities and Configuration Functions

This chapter describes how to use RAIDar to run system utilities and perform advanced configuration tasks. Topics covered in this chapter are:

Updating Software

You can update controller software by loading a software package file. A software package file contains the following software components:

- Storage Controller and its loader
- Memory controller FPGA
- Management Controller and its loader
- Expander Controller
- PSU
- CPLD

RAIDar automatically updates only those types of software that require updating.

Note – By default the storage system's Partner Firmware Upgrade option is enabled, so when you upgrade one controller the system automatically upgrades the partner controller. If Partner Firmware Upgrade is disabled or if the Independent Cache Performance Mode option is enabled, after updating software on one controller you must manually upgrade the partner controller.

Caution – Do not turn off or restart the system during this process. If the code load is interrupted or there is a power failure, the unit might not be operational. If this occurs, contact technical support to attempt a serial code load recovery. In some cases the unit might need to be returned to the factory for reprogramming.

To update controller software:

1. Ensure that the software package file is saved to a location on your network that the system can access.
2. Select Manage > Update Software > Controller Software.

   The Load Software panel is displayed, which describes the update process and lists your current software versions.
3. Click Browse and select the software package file.
4. Click Load Software Package File.

   If the system finds a problem with the file, it shows a message at the top of the page. To resolve the problem, try the following:

   - Be sure to select the software package file that you just downloaded.
   - Download the file again, in case it got corrupted. Do not attempt to edit the file.

   After about 30 seconds, the Load Software To Controller Module panel is displayed. This page lets you know whether the file was validated and what software components are in the file. The system only updates the software that has changes.
5. Click Proceed With Code Update.

   A Code Load Progress window is displayed to show the progress of the update, which can take several minutes to complete. Do not power off the system during the code load process. When the firmware upload is complete, the controller resets after which the opposite controller automatically repeats the process to load the new firmware. When the update completes on the connected controller, you are logged out.
6. Wait one minute for the controller to start and then click Log In to reconnect to RAIDar.

Disabling Partner Firmware Upgrade

If a service technician tells you to disable partner firmware upgrade:

1. Select Manage > General Config > System Configuration.
2. Set Partner Firmware Upgrade to Disabled.

Changing the 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.

Priority Value	Description
High	Use when your highest priority is to get the system back to a fully fault-tolerant state. This causes heavy I/O with the host to be slower than normal. This value is the default.
Medium	Use when you want to balance data streaming with data redundancy.
Low	Use when streaming data without interruption, such as for a web server, is more important than data redundancy. This enables the Reconstruct or other utility to run at a slower rate with minimal effects on host I/O.

To change utility priority:

1. Select Manage > General Config > System Configuration.
2. For Utility Priority, select a priority.
3. Click Change System Configuration.

Rescanning for Drive Changes

Normally, you don't need to manually scan the system for drive module changes. The RAID controllers automatically detect that drives have been added or removed. When drives are inserted they are detected after a short delay, which allows the drives to spin up.

If you perform a manual rescan, it temporarily pauses all I/O processes, then resumes normal operation.

To rescan disk ports, as an Advanced Manage user:

1. Select Manage > Utilities > Disk Drive Utilities > Rescan.
2. In the Rescan For Devices panel, click Rescan.

   The Rescan Summary panel shows the number of drives that were present before the rescan; are detected by the rescan; are newly discovered; are not detected and presumed removed; and are present with a changed address.

Resetting Host Channels

For a Fibre Channel system using FC-AL (loop) topology, you might need to reset a host link to fix a host connection or configuration problem. As an Advanced Manage user, you can use this command to remotely issue a loop initialization primitive (LIP) on specified controller channels.

To reset host channels:

1. Select Manage > Utilities > Host Utilities > Reset Host Channel.
2. Set the channel and controller options.
3. Click Reset Host Channel.

Clearing Unwritable Cache Data

Unwritable cache data is data in the controller cache that cannot be written out to a virtual disk because that virtual disk is no longer accessible. The virtual disk may be offline or missing. Unwritable cache data can exist if I/O to the virtual disk does not complete because drives or enclosures fail or are removed before the data can be written. Recovery is possible if the missing devices can be restored so that the cached data can be written to the virtual disk.

Unwritable cache data might affect performance because it ties up the cache space and prevents that space from being used by other virtual disks that might be performing I/O. The percentage of the cache filled with unwritable cache data appears. This data can be from one or multiple virtual disks. If the data is from only one virtual disk, then the serial number for this virtual disk appears. If the data is from multiple virtual disks, then only the serial number of one virtual disk appears. If the unwritable cache data for this virtual disk is cleared, then the serial number of the next virtual disk appears.

Caution – Make sure that the data is no longer needed before clearing it. Once writable cache is cleared, it cannot be recovered.

To remove data from the cache:

1. Select Manage > Utilities > Recovery Utilities > Cache Data Status.
2. Click Clear Unwritable Cache Data.

   If there is unwritable cache data, this page specifies the percentage of both controllers' cache that this data occupies. Otherwise, the page shows that there is no unwritable cache data.

Note – Event log entries specify the percentage of the owning controller's cache that unwritable data occupies. Therefore in a dual-controller system, the percentage that the Cache Data Status page shows is half the percentage that the event log shows.

Restoring a Saved Configuration File

As an Advanced Manage user, if you have created a backup configuration file as explained in “Saving the Configuration to a File” on page 54, you can load (restore) the configuration data to either:

- The same system to revert its current configuration to the saved configuration
- A second system to “clone” the first system's configuration

---

Note – The file does not include configuration data for virtual disks and volumes. This data is saved as metadata in the first sectors of associated disk drives.

To restore a configuration file:

1. Select Manage > Utilities > Configuration Utilities > Restore Config File.
2. Select the IP address option you want to use for the restore:

Option	Description
Use IP addresses and network information as currently found on RAID controller A and RAID controller B	Restores the configuration file to the system that RAIDar is currently connected to and retains the currently assigned IP addresses. Use this to restore a configuration file to the current system without changing IP addresses of the system. This option ignores the IP addresses in the configuration file.
Use new IP addresses and network information	Restores the configuration file to the system that RAIDar is currently connected to and changes the system’s IP addresses to what you enter on the next page. Use this option to clone the system and change the IP addresses to what you enter. This option ignores the IP addresses in the configuration file. Enter the IP address, IP subnet mask, and gateway IP address values for each controller in the same system. After the file is restored, you must reconnect to the system using one of the new IP addresses.
Use IP address and network information as found within configuration file	Restores the configuration file to the system that RAIDar is currently connected to and changes the IP addresses to those contained in the configuration file. Use this option to restore a configuration file to the current system when the IP addresses in the configuration file are the IP addresses you want assigned to the system. After the file is restored, you might need to reconnect to the system using one of the IP addresses from the file.

3. Click Continue.

   A new page is displayed whose content depends on the IP address option you selected.

4. If you selected the second option in Step 2:

   a. Enter network information in the fields.

   b. Click Continue Restore Process.

5. Click Browse to navigate to a previously saved configuration file.

6. Click Restore Configuration File.

Viewing and Restoring Default Settings

You can view current and default settings as well as restore the system's default settings.

Viewing Changed Settings

To view the storage system parameter settings that have been changed from the default configuration, and their default settings:

- Select Manage > Utilities > Configuration Utilities > Show Changed Settings.

Restoring All Defaults

As an Advanced Manage user, if the system is not working properly and you cannot determine why, you can restore its default configuration settings. This restores all defaults except the following:

- Settings related to virtual disks and volumes
- IP settings (address, subnet mask, and gateway)
- System time and date

You then can change the settings that are critical to your configuration.

Caution – Restoring default settings replaces your current configuration changes with the original manufacturer configuration settings. Some of these settings take effect immediately while others take effect after you restart the RAID controllers. Restoring default settings cannot be undone.

To restore all defaults:

1. Select Manage > General Config > Restore Defaults.
2. (Optional) To see a list of the current settings and default settings, click See Restore Defaults Changes. When done, click Return to Restore Defaults Page.
3. In the Restore Defaults panel, click Restore Defaults.

   Changes take effect immediately, except for Requested Loop ID for host ports (one per controller), which requires a controller restart. Select Manage > Restart System > Shut Down/Restart**.**

Enabling and Disabling Background Scrub for Disks

You can enable or disable whether the system analyzes disk drives associated with specified virtual disks to detect, report, and store information about disk drive defects. At the vdisk level, hard errors, medium errors, and bad block replacements (BBRs) are reported. At the drive level, metadata read errors, SMART events during scrub, bad blocks during scrub, and new drive defects during scrub are reported. Any errors found are reported as events. Background scrub always runs at background utility priority. The default is Disabled.

To enable or disable background scrub:

1. Select Manage > General Config > System Configuration.
2. Set Background Scrub to Enabled or Disabled.
3. Click Change System Configuration.

Controlling Host Access to the System's Write-Back Cache Setting

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

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

To enable or disable host control of write-back cache:

1. Select Manage > General Config > System Configuration.
2. Click Advanced Options.
3. In the SCSI Configuration Options panel, set Host Control Of Write-Back Cache to Enabled or Disabled.
4. Click Change SCSI Configuration Options.

Changing the Sync Cache Mode Option

Sync Cache Mode controls how the SCSI `SYNCHRONIZE CACHE` command is handled. Typically, you do no need to change this option. However, if the system has performance problems or problems writing to databases or other applications, contact technical support to determine if you should change this option.

To change the cache synchronization mode:

1. Select Manage > General Config > System Configuration.
2. Click Advanced Options.
3. In the SCSI Configuration Options panel, set Sync Cache Mode to one of the following options:
   - Immediate – Good status is returned immediately and cache content is unchanged. This option is the default.
   - Flush To Disk – Good status is returned only after all write-back data for the specified volume is flushed to disk.
4. Click Change SCSI Configuration Options.

   Changes take effect immediately.

Changing the Missing LUN Response Option

Some operating systems do not look beyond LUN 0 if they do not find a LUN 0 or cannot handle noncontiguous LUNs. Missing LUN Response 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 virtual disks). Use the default value unless a service technician asks you to change it to work around a host driver problem.

To change the missing LUN response:

1. Select Manage > General Config > System Configuration.
2. Click Advanced Options.
3. In the SCSI Configuration Options panel, set Missing LUN Response to one of the following options:
   - 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 sensekey = 2, code = 4, qualifier = 3. This option is the default.
   - Illegal Request – Sends a reply that there is a LUN but that the request is "illegal." Sense data returned is sensekey = 5, code = 25h, qualifier = 0.
4. Click Change SCSI Configuration Options.

APPENDIX A

Configuring SNMP

This appendix describes the Simple Network Management Protocol (SNMP) capabilities that Phoenix storage systems support. This includes standard MIB-II, the Fibre Alliance SNMP Management Information Base (MIB) version 2.2 objects, and enterprise traps. This appendix also describes how to use RAIDar to configure SNMP and how to integrate a Phoenix storage system into a network management system.

Topics covered in this appendix are:

Introduction

Phoenix storage systems can report their status through SNMP. SNMP provides basic discovery using MIB-II, more detailed status with the FA MIB 2.2, and asynchronous notification using enterprise traps.

SNMP is a widely used network monitoring and control protocol. It is an application layer protocol that facilitates the exchange of management information between network devices. It is part of the Transmission Control Protocol/Internet Protocol (TCP/IP) protocol suite.

SNMP enables network administrators to manage network performance, find and solve network problems, and plan for network growth. Data is passed from SNMP agents reporting activity on each network device to the workstation console used to oversee the network. The agents return information contained in a Management Information Base (MIB), which is a data structure that defines what is obtainable from the device and what can be controlled (turned on and off, etc.).

An SNMP object identifier (OID) is a number assigned to devices in a network for identification purposes. OID numbering is hierarchical. Using the IETF notation of digits and dots resembling very long IP addresses, various registries such as ANSI assign high-level numbers to vendors and organizations. They, in turn, append digits to the number to identify individual devices or software processes.

Phoenix systems use SNMPv2c, which improves on SNMPv1 features and uses its community-based security scheme.

Standard MIB-II Behavior

MIB-II is implemented to support basic discovery and status.

In the system group, all objects can be read. The contact, name, and location objects can be set.

The system object identifier (`sysObjectID`) is based on the vendor name followed by ".2." and the identifier for the particular product model. For example, the object identifier for a Phoenix International System Inc. model 69501 storage system is 1.3.6.1.4.1.347.2.2730. System uptime is an offset from the first time this object is read.

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. 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 following example shows the elements of a trap for an FC link down event:

The text of the trap MIB is included at the end of this appendix.

FA MIB 2.2 SNMP Behavior

The FA2.2 MIB objects are in compliance with the Fibre Alliance MIB v2.2 Specification (FA MIB2.2 Spec). For a full description of this MIB, go to: http://www.fibrealliance.org/fb/mib/mib2_2.htm

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

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

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

The traps section is not supported. It has been replaced by an ability to configure trap destinations using the CLI or RAIDar. The statistics section is not implemented.

The following table lists the MIB objects, their descriptions and the value set in a Phoenix storage system. Unless specified otherwise, objects are *not* settable.

Table A-1 FA MIB 2.2 Objects, Descriptions, and Values

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

Table A-1 FA MIB 2.2 Objects, Descriptions, and Values *(Continued)*

Object	Description	Value
connUnitTable	Includes the following objects as specified by the FA MIB2.2 Spec	
connUnitId	Unique identification for this connectivity unit	Total of 16 bytes comprised of 8 bytes of the node WWN or similar serial number-based identifier (for example, 1000005013b05211) with the trailing 8 bytes equal to zero
connUnitGlobalId	Same as connUnitId	Same as connUnitId
connUnitType	Type of connectivity unit	storage-subsystem[11]
connUnitNumports	Number of host ports in the connectivity unit	Number of host ports
connUnitState	Overall state of the connectivity unit	online[2] or unknown[1], as appropriate
connUnitStatus	Overall status of the connectivity unit	ok [3], warning[4], failed[5], or unknown[1], as appropriate
connUnitProduct	Connectivity unit vendor's product model name	Model string
connUnitSn	Serial number for this connectivity unit	Serial number string
connUnitUpTime	Number of centiseconds since the last unit initialization	0 at startup
connUnitUrl	Same as systemURL	Same as systemURL
connUnitDomainId	Not used; set to all 1s as specified by the FA MIB2.2 Spec	0xFFFF
connUnitProxyMaster	Stand-alone unit returns yes for this object	yes [3] since this is a stand-alone unit
connUnitPrincipal	Whether this connectivity unit is the principal unit within the group of fabric elements. If this value is not applicable, returns unknown.	unknown[1]

Table A-1 FA MIB 2.2 Objects, Descriptions, and Values *(Continued)*

Object	Description	Value
connUnitNumSensors	Number of sensors in the connUnitSensorTable	33
connUnitStatusChangeTime	Same as statusChangeTime	Same as statusChangeTime
connUnitConfiguration ChangeTime	Same as configurationChangeTime	Same as configurationChange Time
connUnitNumRevs	Number of revisions in the connUnitRevsTable	16
connUnitNumZones	Not supported	0
connUnitModuleId	Not supported	16 bytes of 0s
connUnitName	Settable: Display string containing a name for this connectivity unit	Default: Uninitialized Name
connUnitInfo	Settable: Display string containing information about this connectivity unit	Default: Uninitialized Info
connUnitControl	Not supported	invalid[2] for an SNMP GET operation and not settable through an SNMP SET operation.
connUnitContact	Settable: Contact information for this connectivity unit	Default: Uninitialized Contact
connUnitLocation	Settable: Location information for this connectivity unit	Default: Uninitialized Location
connUnitEventFilter	Defines the event severity that will be logged by this connectivity unit. Settable only through 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

Table A-1 FA MIB 2.2 Objects, Descriptions, and Values *(Continued)*

Object	Description	Value
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 201
connUnitRevsRevId	Vendor-specific string identifying a revision of a component of the connUnit	String specifying the code version. Reports "Not Installed or Offline" if module information is not available.
connUnitRevsDescription	Description of a component to which the revision corresponds	See "External Details for connUnitRevsTable" on page 201
connUnitSensorTable	Includes the following objects as specified by the FA MIB2.2 Spec	
connUnitSensorUnitId	connUnitId of the connectivity unit that contains this sensor table	Same as connUnitId
connUnitSensorIndex	Unique value for each connUnitSensorEntry between 1 and connUnitNumSensors	See "External Details for connUnitSensorTable" on page 202
connUnitSensorName	Textual identification of the sensor intended primarily for operator use	See "External Details for connUnitSensorTable" on page 202
connUnitSensorStatus	Status indicated by the sensor	ok[3], warning[4], or failed[5] as appropriate for FRUs that are present, or other[2] if FRU is not present.
connUnitSensorInfo	Not supported	Empty string

Table A-1 FA MIB 2.2 Objects, Descriptions, and Values *(Continued)*

Object	Description	Value
connUnitSensorMessage	Description the sensor status as a message	connUnitSensorName followed by the appropriate sensor reading. Temperatures display in both Celsius and Fahrenheit; for example, CPU Temperature (Controller Module A): 48C 118F). Reports "Not installed" or "Offline" if data is not available.
connUnitSensorType	Type of component being monitored by this sensor	See "External Details for connUnitSensorTable" on page 202
connUnitSensor Characteristic	Characteristics being monitored by this sensor	See "External Details for connUnitSensorTable" on page 202
connUnitPortTable	Includes the following objects as specified by the FA MIB2.2 Spec	
connUnitPortUnitId	connUnitId of the connectivity unit that contains this port	Same as connUnitId
connUnitPortIndex	Unique value for each connUnitPortEntry between 1 and connUnitNumPorts	Unique value for each port, between 1 and the number of ports
connUnitPortType	Port type	not-present[3], or n-port[5] for point-to-point topology, or l-port[6]
connUnitPortFCClassCap	Bit mask that specifies the classes of service capability of this port. If this is not applicable, returns all bits set to zero.	Fibre Channel ports return 8 for class-three
connUnitPortFCClassOp	Bit mask that specifies the classes of service that are currently operational. If this is not applicable, returns all bits set to zero.	Fibre Channel ports return 8 for class-three

Table A-1 FA MIB 2.2 Objects, Descriptions, and Values *(Continued)*

Object	Description	Value
connUnitPortState	State of the port hardware	unknown[1], online[2], offline[3], bypassed[4]
connUnitPortStatus	Overall protocol status for the port	unknown[1], unused[2], ok[3], warning[4], failure[5], notparticipating[6], initializing[7], bypass[8]
connUnitPortTransmitterType	Technology of the port transceiver	unknown[1] for Fibre Channel ports
connUnitPortModuleType	Module type of the port connector	unknown[1]
connUnitPortWwn	Fibre Channel World Wide Name (WWN) of the port if applicable	WWN octet for the port, or empty string if the port is not present
connUnitPortFCId	Assigned Fibre Channel ID of this port	• Fibre Channel ID of the port • All bits set to 1 if the Fibre Channel ID is not assigned or if the port is not present
connUnitPortSn	Serial number of the unit (for example, for a GBIC). If this is not applicable, returns an empty string.	Empty string
connUnitPortRevision	Port revision (for example, for a GBIC)	Empty string
connUnitPortVendor	Port vendor (for example, for a GBIC)	Empty string
connUnitPortSpeed	Speed of the port in KByte per second (1 KByte = 1000 Byte)	Port speed in KByte per second, or 0 if the port is not present
connUnitPortControl	Not supported	invalid[2] for an SNMP GET operation and not settable through an SNMP SET operation

Table A-1 FA MIB 2.2 Objects, Descriptions, and Values *(Continued)*

Object	Description	Value
connUnitPortName	String describing the addressed port	See "External Details for connUnitPortTable" on page 204
connUnitPortPhysical Number	Port number represented on the hardware	Port number represented on the hardware
connUnitPortStatObject	Not supported	0 (No statistics available)
connUnitEventTable	Includes the following objects as specified by the FA MIB2.2 Spec	
connUnitEventUnitId	connUnitId of the connectivity unit that contains this port	Same as connUnitId
cconnUnitEventIndex	Index into the connectivity unit's event buffer, incremented for each event	Starts at 1 every time there is a table reset or the unit's event table reaches its maximum index value
connUnitEventId	Internal event ID, incremented for each event, ranging between 0 and connUnitMaxEvents	Starts at 0 every time there is a table reset or connUnitMaxEvents is reached
connUnitREventTime	Real time when the event occurred, in the following format: DDMMYYYY HHMMSS	0 for logged events that occurred prior to or at startup
connUnitSEventTime	sysuptime timestamp when the event occurred	0 at startup
connUnitEventSeverity	Event severity level	error[5], warning[6] or info[8]
connUnitEventType	Type of this event	As defined in CAPI
connUnitEventObject	Not used	0
connUnitEventDescr	Text description of this event	Formatted event, including relevant parameters or values
connUnitLinkTable	Not supported	N/A
connUnitPortStatFabric Table	Not supported	N/A
connUnitPortStatSCSITable	Not supported	N/A

Table A-1 FA MIB 2.2 Objects, Descriptions, and Values *(Continued)*

Object	Description	Value
connUnitPortStatLANTable	Not supported	N/A
SNMP TRAPS	The following SNMP traps are supported	
trapMaxClients	Maximum number of trap clients	1
trapClientCount	Number of trap clients currently enabled	1 if traps enabled; 0 if traps not enabled
connUnitEventTrap	This trap is generated each time an event occurs that passes the connUnitEventFilter and the trapRegFilter	N/A
trapRegTable	Includes the following objects per the FA MIB2.2 Spec	
trapRegIpAddress	IP address of a client registered for traps	IP address set through Telnet
trapRegPort	User Datagram Protocol (UDP) port to send traps to for this host	162
trapRegFilter	Settable: Defines the trap severity filter for this trap host. The connUnit will send traps to this host that have a severity level less than or equal to this value.	Default: warning[6]
trapRegRowState	Specifies the state of the row	• READ: rowActive[3] if traps are enabled through Telnet; otherwise rowInactive[2] • WRITE: Not supported

External Details for Certain FA MIB 2.2 Objects

Tables in this section specify values for certain objects described in Table A-1.

External Details for connUnitRevsTable

The following table provides external details for the `connUnitRevsTable` objects `connUnitRevsIndex` and `connUnitRevsDescription`.

Table A-2 connUnitRevsTable Index and Description Values

Revs Index	Revs Description
1	Firmware revision for Storage Controller (Controller Module A)
2	Firmware revision for Storage Controller (Controller Module B)
3	Firmware revision for Memory Controller (Controller Module A)
4	Firmware revision for Memory Controller (Controller Module B)
5	Firmware revision for Storage Controller loader (Controller Module A)
6	Firmware revision for Storage Controller loader (Controller Module B)
7	Firmware revision for Management Controller (Controller Module A)
8	Firmware revision for Management Controller (Controller Module B)
9	Firmware revision for MC loader (Controller Module A)
10	Firmware revision for MC loader (Controller Module B)
11	Firmware Revision for Unified CPLD (Controller Module A)
12	Firmware Revision for Unified CPLD (Controller Module B)
13	Firmware Revision for Expander (Controller Module A)
14	Firmware Revision for Expander (Controller Module B)
15	Hardware Revision for Controller Module A
16	Hardware Revision for Controller Module B

External Details for connUnitSensorTable

The following table provides external details for the `connUnitSensorTable` objects `connUnitSensorIndex`, `connUnitSensorName`, `connUnitSensorType`, and `connUnitSensorCharacteristic`.

Table A-3 connUnitSensorTable Index, Name, Type, and Characteristic Values

Sensor Index	Sensor Name	Sensor Type	Sensor Characteristic
1	CPU Temperature (Controller Module A)	board [8]	temperature[3]
2	CPU Temperature (Controller Module B)	board [8]	temperature[3]
3	FPGA Temperature (Controller Module A)	board [8]	temperature[3]
4	FPGA Temperature (Controller Module B)	board [8]	temperature[3]
5	Onboard Temperature 1 (Controller Module A)	board [8]	temperature[3]
6	Onboard Temperature 1 (Controller Module B)	board [8]	temperature[3]
7	Onboard Temperature 2 (Controller Module 1)	board [8]	temperature[3]
8	Onboard Temperature 2 (Controller Module 2)	board [8]	temperature[3]
9	Capacitor Temperature (Controller Module 3)	board [8]	temperature[3]
10	Capacitor Temperature (Controller Module 4)	board [8]	temperature[3]
11	CM Temperature (Controller Module A)	enclosure[7]	temperature[3]
12	CM Temperature (Controller Module A)	enclosure[7]	temperature[3]
13	Power Supply 1 Temperature	enclosure[7]	temperature[3]
14	Power Supply 2 Temperature	enclosure[7]	temperature[3]
15	Capacitor Pack Voltage (Controller Module A)	board [8]	power[9]
16	Capacitor Pack Voltage (Controller Module B)	board [8]	power[9]
17	Capacitor Cell 1 Voltage (Controller Module A)	board [8]	power[9]
18	Capacitor Cell 1 Voltage (Controller Module B)	board [8]	power[9]
19	Capacitor Cell 2 Voltage (Controller Module A)	board [8]	power[9]
20	Capacitor Cell 2 Voltage (Controller Module B)	board [8]	power[9]
21	Capacitor Cell 3 Voltage (Controller Module A)	board [8]	power[9]
22	Capacitor Cell 3 Voltage (Controller Module B)	board [8]	power[9]
23	Capacitor Cell 4 Voltage (Controller Module A)	board [8]	power[9]

Table A-3 connUnitSensorTable Index, Name, Type, and Characteristic Values *(Continued)*

Sensor Index	Sensor Name	Sensor Type	Sensor Characteristic
24	Capacitor Cell 4 Voltage (Controller Module B)	board [8]	power[9]
25	Capacitor Charge Current (Controller Module A)	board [8]	currentValue[6]
26	Capacitor Charge Current (Controller Module B)	board [8]	currentValue[6]
27	Power Supply 1 Voltage, 12V	power-supply[5]	power[9]
28	Power Supply 1 Voltage, 5V	power-supply[5]	power[9]
29	Power Supply 1 Voltage, 3.3V	power-supply[5]	power[9]
30	Power Supply 2 Voltage, 12V	power-supply[5]	power[9]
31	Power Supply 2 Voltage, 5V	power-supply[5]	power[9]
32	Power Supply 2 Voltage, 3.3V	power-supply[5]	power[9]
33	Overall Status	enclosure[7]	other[2]

External Details for connUnitPortTable

The following table provides external details for the `connUnitPortTable` objects `connUnitPortIndex` and `connUnitPortName`.

Table A-4 connUnitPortTable Index and Name Values

Port Index	Port Name
1	FC Host Port 1 (Controller Module A)
2	FC Host Port 2 (Controller Module B)
3	FC Host Port 1 (Controller Module A)
4	FC Host Port 2 (Controller Module B)

Configuring SNMP Event Notification in RAIDar

As a Manage user you can configure and enable SNMP event notification. To do so:

1. Select the level of events to include in the FA2.2 event table; see "Setting the SNMP Event Table Filter" on page 44.
2. Verify that the storage system's SNMP service is enabled; see "Configuring Network Management Services" on page 46.
3. Select event levels for notification; see "Selecting Event Categories to Monitor" on page 48.
4. Configure and enable SNMP traps; see "Configuring SNMP Traps" on page 52.

SNMP Management Using HP OpenView

There are numerous network management systems that can manage storage devices using SNMP. A commonly used system is HP OpenView. As a general example of integrating a Phoenix storage system into a network management system, some aspects of configuring HP OpenView are described here. For more information about using HP OpenView, see its documentation.

Note – A Phoenix storage system has an internal network interface with an IP address of 192.168.0.1 for controller A and 192.168.0.2 for controller B. These addresses can appear on the HP OpenView network map as unreachable nodes. This is normal.

Loading MIBs

By itself, OpenView can listen for and dispatch SNMP traps. However, MIBs are supplied to make the best use of the management feature.

It is assumed that HP OpenView has discovered the Phoenix storage system. Refer to your HP OpenView documentation for details on node discovery.

To launch HP OpenView on Microsoft Windows, perform the following steps:

1. From the Windows Start menu, select HP OpenView.
2. Select Network Node Manager.
3. From the Root dialog, navigate to the network segment on which the Phoenix nodes reside.

   To load MIBs, perform the following steps:

4. Choose Options > Load/Unload MIBs: SNMP.

Note – The FibreAlliance MIB 2.2 available from the FibreAlliance web site contains an error and might not load correctly in HP OpenView. Before loading this MIB, remove the extra hyphen character on line 14.

---

5. Click Load.

6. Select the MIBs to be loaded from the dialog, and click OK.

   Because the MIB contains TRAP/NOTIFICATION information, OpenView detects this. A dialog is displayed requesting confirmation to load the definitions into the OpenView event system.

7. Click OK to continue.

   If the definitions load successfully, a dialog is displayed confirming that the macro definitions have successfully loaded into `trapd.conf`.

8. Click OK.

9. Repeat the preceding steps for both the trap MIB and the FibreAlliance MIB 2.2.

Note – You can also unload MIBs by selecting them from the list and clicking Unload.

---

10. After both MIBs have been loaded, the Load/Unload dialog shows the MIBs.

11. Click Close to exit the Load/Unload dialog.

Configuring Events

Because the MIBs contain information on traps, it is possible to configure these events. Events can be displayed in the Alarm browser, in a pop-up window, forwarded to other hosts, and logged to files. Refer to your HP OpenView documentation for details.

If the MIBs have been loaded successfully, the Enterprise Name and Enterprise ID will be displayed in the Event Configuration dialog.

To view and modify events, perform the following steps:

1. Select Options > Event Configuration.

   The Event Configuration dialog is displayed.

2. To modify an event, select the Enterprise Name in the upper pane, and double-click on the Event Name in the lower pane.

   The Modify Events tabs are displayed.

3. Select the Event Message tab.

   The Event Message dialog is displayed.

4. To display the event in the Alarm browser, select an appropriate category.

   Traps can be categorized by event type. In the following example, `dhEventCriticalTrap`, a category of “Error Alarms” is used. A category of “Status Alarms” could be used for `dhEventInfoTrap`.

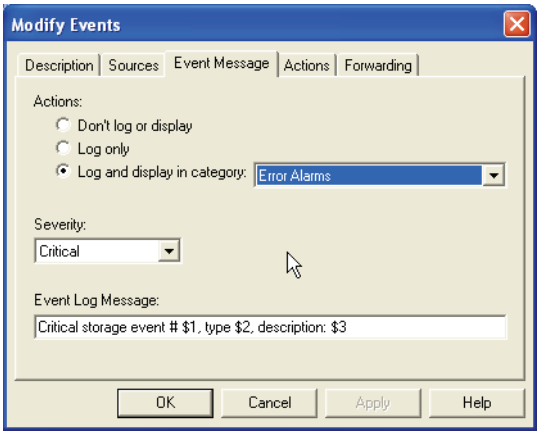

5. Enter a message for the event in the Event Log Message text box.

   The MIBs should provide HP OpenView with a default message string. In this example, $1, $2, and $3 represent the values of the variables sent as part of the trap PDU. Refer to your HP OpenView documentation for details on other “$” variables available as part of the event message.

6. Click OK to save changes and exit the dialog.

7. To view alarms from the main menu bar, select Fault > Alarms.

   The All Alarms Browser dialog is displayed and lists any current alarm messages.

8. Repeat for all events shown in the lower pane.

Viewing and Setting System Group Objects

SNMP must be enabled on the system to view and set system group objects.

In RAIDar, perform the following steps:

1. Select Manage > General Config > Services Security.

2. Set Simple Network Management Protocol (SNMP) to Enabled.

   This setting is Enabled by default.

3. Click Update Network Management Services.

In HP OpenView, perform the following steps:

1. Browse the system group objects for a node by selecting the node on the segment map.
2. Select Tools > SNMP MIB Browser.

   Confirm that the correct Name or IP Address is displayed.
3. Navigate to the following MIB Object ID.

   `iso.org.dod.internet.mgmt.mib-2`
4. Select `system` from the list, and click Start Query.

   Read/Write values can be set from this dialog.

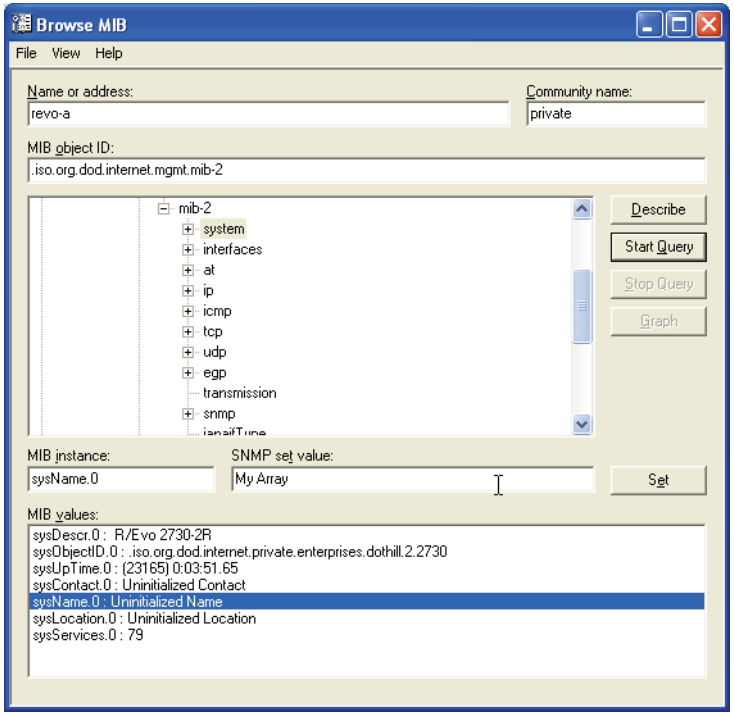

5. To set the system name, enter the community name **`private`**, select `sysName.0` from the list, and enter a new value in the SNMP Set Value field.

6. Click Set.

   You can use the community name `public` for queries as shown here. If the community name is blank, `public` is used by default.

   A new query on the system group shows the new value.

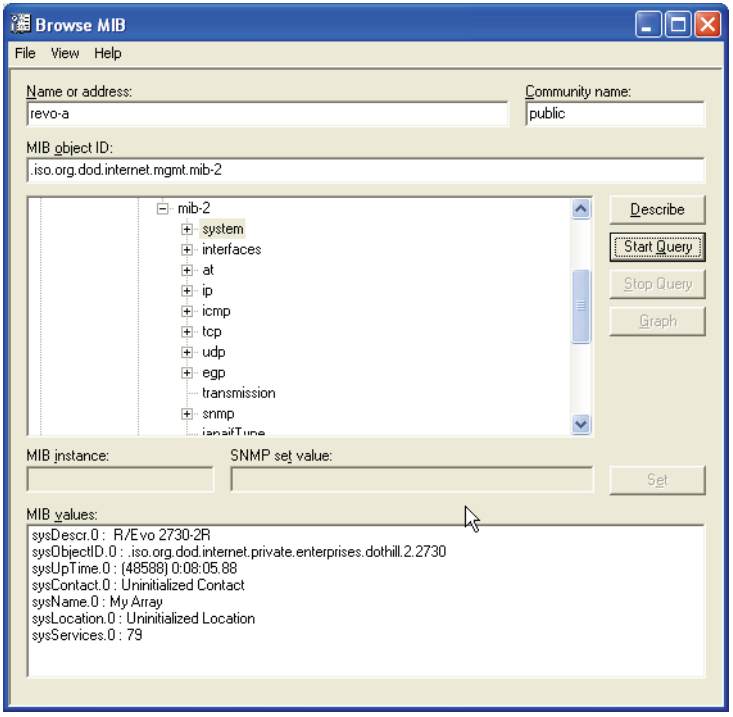

Enterprise Trap MIB

The following pages show the source for the Phoenix International System Inc. traps MIB. This MIB defines the content of the SNMP traps that a Phoenix storage system generates.

```
-- ------------------------------------------------------
-- Dot Hill Low Cost Array MIB for SNMP Traps
--
-- $Revision: 1.1 $
--
-- 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}

-- 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
```

APPENDIX B

RAID Levels

This appendix describes the different RAID levels that virtual disks in your system can use.

Topics covered in this appendix are:

Introduction

The RAID controllers enable you to set up and manage virtual disks, whose storage may be spread across multiple disk drives. This is accomplished through software resident in the RAID controller. RAID (Redundant Array of Independent Disks) refers to virtual disks in which part of the storage capacity may be used to store redundant information. The redundant information enables the system to reconstruct data if a drive in the virtual disk fails.

Hosts see each partition of a virtual disk, known as a volume, as a single disk drive. A volume is actually a portion of the storage space on disk drives behind a RAID controller. The RAID controller software makes each volume appear as a single, very large disk drive. Depending on the RAID level used for a virtual disk, the disk drive presented to hosts has advantages in fault-tolerance, cost, performance, or a combination of these. This section explains the different RAID levels and the disk requirements for each level.

Note – Choosing the right RAID level for your needs improves performance. The following table includes examples of storage needs and appropriate RAID levels.

Table B-1 Example Applications and RAID Levels

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

RAID Level Descriptions

RAID levels are numbered from 0 through 6; a higher RAID level does not necessarily indicate a higher level of performance or fault tolerance. The RAID controllers support RAID levels that have proven to be the most useful for RAID applications: RAID 0, 1, 10, 3, 5, 50, and 6. You can use Non-RAID for a virtual disk that will have a single drive and not need the data redundancy or performance benefits of RAID.

RAID 0

In a RAID 0 virtual disk, data is distributed, or *striped*, across the disk drives in the virtual disk. The virtual disk appears to the host as one large disk with a capacity approximately equal to the combined capacity of the disk drives. Because multiple reads and writes can be handled in parallel, the I/O performance of the virtual disk is much better than that of a single disk drive.

RAID 0 virtual disks do not store redundant data, so they are not true RAID applications. If one disk drive fails, the entire virtual disk fails and all virtual disk data is lost. The fault tolerance of a RAID 0 virtual disk, therefore, is less than that of any single disk drive in the virtual disk. The term RAID 0 is widely used for these virtual disks, however, because they are conceptually similar to true RAID applications.

RAID 1, RAID 10

In RAID 1 and RAID 10 virtual disks (commonly called *mirrored* virtual disks), disk drives are paired, with both disk drives in a pair containing the same data. When data is written to a mirrored virtual disk, it is written twice—once to each disk drive in the pair. A RAID 1 virtual disk has only one set of paired disk drives. A RAID 10 virtual disk has multiple pairs, across which data is striped.

The read performance of RAID 1 virtual disks can be much better than that of a single disk drive, while the write performance is slightly lower. In RAID 10 virtual disks, both read performance and write performance are better than those of a single disk drive.

A mirrored virtual disk is also highly reliable, because both disk drives in a pair must fail for the virtual disk to fail. In an virtual disk with five pairs of mirrored disk drives, for example, the virtual disk can maintain its integrity even if five disks fail—as long as each pair is left with one good disk. The main disadvantage of a mirrored virtual disk is its cost. Because all disk drives must have a twin, you must use twice the number of disk drives that actually contribute to the virtual disk capacity. In an eight-disk virtual disk, for example, you have only four disks of usable capacity.

RAID 3

RAID 3 virtual disks contain redundant information in the form of parity data, which is calculated block-by-block for all user data. The user data is distributed across all but one of the disk drives in the virtual disk. The parity data is written exclusively to the parity disk (also known as the check disk). In the event of a disk drive failure, the data can be reconstructed from corresponding data stripes on the remaining disk drives in the virtual disk.

RAID 3 provides excellent I/O performance for applications that require high data transfer rates such as image processing, video processing, scientific data collection, batch data processing, or sequential reads and writes.

RAID 3 is not well suited for transaction processing or other applications that require simultaneous reads and writes.

RAID 5

RAID 5 virtual disks contain redundant information in the form of parity data, which is calculated block-by-block for all user data. The parity information is distributed across the disk drives in the virtual disk and occupies the equivalent capacity of approximately one disk drive. Data is interspersed with the parity information. If one disk drive in the virtual disk fails, the data on the failed disk drive can be reconstructed from the parity data and user data on the remaining disk drives. Two disk drives must fail before the entire virtual disk fails.

The read performance of a RAID 5 virtual disk is excellent—comparable to that of a RAID 0 virtual disk. Write performance is lower than that of a RAID 0 virtual disk, because write operations involve calculating and writing new parity data as well as writing the new user data.

RAID 50

RAID 50 virtual disks are made up of two or more RAID 5 virtual disks, across which data is striped. RAID 50 virtual disks contain redundant information in the form of parity data, which is calculated block-by-block for all user data. As in a RAID 5 virtual disk, the parity information is distributed across the disk drives in the virtual disk and occupies the equivalent capacity of one disk drive per RAID 5. Data is interspersed with the parity information. If one disk drive in the virtual disk fails, the data on the failed disk drive can be reconstructed from the parity data and user data on the remaining disk drives. Two disk drives in one RAID 5 subset must fail before the entire virtual disk fails.

The read performance of a RAID 50 virtual disk is excellent—better than a RAID 5 virtual disk—along with better data protection. Write performance is lower than that of a RAID 0 virtual disk, because write operations involve calculating and writing new parity data as well as writing the new user data.

RAID 6

RAID 6 virtual disks contain redundant information in the form of parity data, which is calculated block-by-block for all user data. The parity information is distributed across the disk drives in the virtual disk and occupies the equivalent capacity of approximately two disk drives. Data is interspersed with the parity information. If one or two disk drives in the virtual disk fail, the data on the failed disk drives can be reconstructed from the parity data and user data on the remaining disk drives. Three disk drives must fail before the entire virtual disk fails.

Non-sequential read and sequential read/write performance is comparable to RAID 5, however non-sequential write performance is less than RAID 5.

Non-RAID

Non-RAID virtual disks provide the ability to create a host-accessible volume consisting of a single disk drive in the system. A Non-RAID virtual disk is nonredundant and its capacity equals the disk drive capacity. Non-RAID virtual disks are useful if you have a single disk drive available and you do not want to use it as a spare.

Comparing RAID Levels

Table A-2 illustrates the differences between the different RAID levels.

Table B-2 RAID Level Comparison

RAID Level	Min. Number of Drives	Description	Strengths	Weaknesses
0	2	Data striping without redundancy	Highest performance	No data protection: if one drive fails all data is lost
1	2	Disk mirroring	Very high performance and data protection; minimal penalty on write performance	High redundancy cost overhead: because all data is duplicated, twice the storage capacity is required
10	4	Combination of RAID 0 (data striping) and RAID 1 (mirroring)	Highest performance and data protection (can tolerate multiple drive failures)	High redundancy cost overhead: because all data is duplicated, twice the storage capacity is required; requires minimum of four drives
3	3	Block-level data striping with dedicated parity drive	Excellent performance for large, sequential data requests (fast read)	Not well-suited for transaction-oriented network applications: single parity drive does not support multiple, concurrent write requests
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	Write performance is slower than RAID 0 or RAID 1

Table B-2 RAID Level Comparison *(Continued)*

RAID Level	Min. Number of Drives	Description	Strengths	Weaknesses
50	6	Combination of RAID 0 (data striping) and RAID 5 with distributed parity	Better random read and write performance and data protection than RAID 5; supports more drives than RAID 5	Lower storage capacity than RAID 5
6	4	Block-level data striping with distributed parity	Best suited for large sequential workloads; non-sequential read and sequential read/write performance is comparable to RAID 5	Higher redundancy cost than RAID 5 because the parity overhead is twice that of RAID 5; not well-suited for transaction-oriented network applications; non-sequential write performance is slower than RAID 5
Non-RAID	1	Non-RAID, nonstriped mapping to a single disk drive	Ability to use a single disk drive to store additional data	Not protected, lower performance (not striped)

Mixing Disk Drive Models

A virtual disk can contain different models of disk drives, even disk drives with different capacities. For example, a virtual disk can include a 250-Gbyte disk drive and a 500-Gbyte disk drive. If you mix disk drives with different capacities, the smallest disk drive determines the logical capacity of all other disk drives in the virtual disk, regardless of RAID level. For example, if a RAID 0 virtual disk contains one 250-Gbyte disk drive and four 500-Gbyte disk drives, the capacity of the virtual disk is equivalent to approximately five 250-Gbyte disk drives. To maximize capacity, use disk drives of similar size.

For greatest reliability, use disk drives of the same size and rotational speed.

APPENDIX C

Host Access to Storage

This appendix describes how the controllers present volumes to data hosts in direct attach and switch attach configurations, during normal operation and after failover. Failover information only applies to a controller enclosure with two controller modules installed.

Topics covered in this appendix are:

Data Presented for Mapped Volumes

A volume in a virtual disk can be mapped through all controller host ports to all data hosts, or through specific controller host ports to specific data hosts. Each mapping between a volume and a data host includes a logical unit number (LUN) that identifies the mapping.

Each controller has a unique, permanent node World Wide Name (WWN) assigned to it. Each FC controller host port has a unique port WWN (or WWPN). The WWN format is:

2*<port-ID>*7*<A/B><multiID><OUI><midplane-SN>*

- *port-ID* – 0 for a node WWN; 0 or 1 for a port WWN.
- *A/B* – 0 for controller A; 8 for controller B.
- *multiID* – 0 for a node WWN; for a port WWN, 0 for the first ID on each port per controller.
- *OUI* – The manufacturer's unique identifier, composed of six ASCII hex digits.
- *midplane-SN* – A serial number derived from the last six ASCII hex digits of the midplane serial number.

Each iSCSI controller host port has a port IP address assigned to it.

Table C-1 shows example node WWNs and FC port WWNs used in this appendix. Notice that the node WWNs differ in the fourth digit and, for a given controller, the port WWNs differ in the second digit.

Table C-1 Example Node WWNs and FC Port WWNs

Controller	Node WWN	Port WWN
A	207000c0ff0a408a	0: 207000C0FF0A408A 1: 217000C0FF0A408A 2: 227000C0FF0A408A 3: 237000C0FF0A408A
B	207800c0ff0a408a	0: 207800C0FF0A408A 1: 217800C0FF0A408A 2: 227800C0FF0A408A 3: 237800C0FF0A408A

Table C-2 shows example node WWNs and iSCSI port IP addresses used in this appendix.

Table C-2 Example Node WWNs and iSCSI Port IP Addresses

Controller	Node WWN	Port IP Address
A	207000c0ff0a408a	0: 10.11.10.4 1: 10.10.10.5
B	207800c0ff0a408a	0: 10.11.10.2 1: 10.10.10.3

69501 Direct Attach Configuration

When a data host is directly connected to controller host ports, loop topology must be used. The host should have one HBA port connected to each controller.

When availability is more important than performance, enable the host port interconnects to connect the host ports in controller A to those in controller B. When the interconnects are enabled, the host has access to both controllers' mapped volumes. This dual access makes it possible to create a redundant configuration without using an external switch.

If one controller fails in this configuration, the interconnects remain active so hosts can continue to access all mapped volumes without the intervention of host-based failover software. The controllers accomplish this by means of FC target multi-ID: while a controller is failed over, each surviving controller host port presents its own port WWN and the port WWN of the interconnected, failed controller host port that was originally connected to the loop. The mapped volumes owned by the failed controller remain accessible until it is removed from the enclosure.

Figure C-1 identifies the host paths in this configuration. For each path, Table C-3 specifies how mapped volumes and port WWNs are presented when both controllers are active and when either controller has failed.

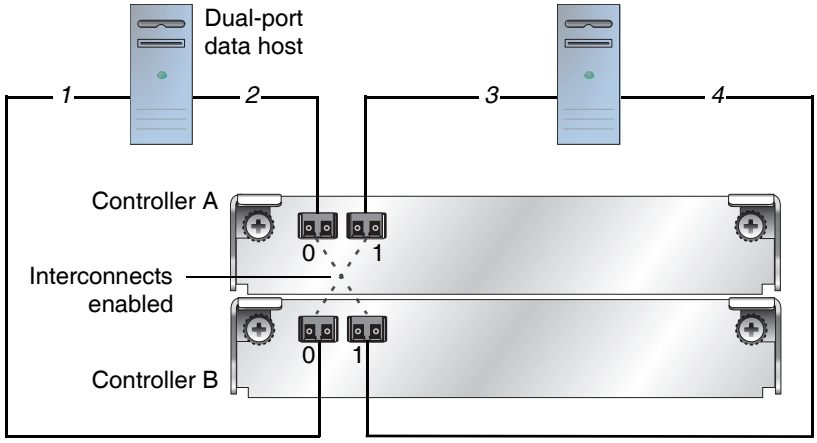

Figure C-1 Direct Attach Configuration with Interconnects Enabled

Table C-3 Storage Presentation in a Direct Attach Configuration with Interconnects Enabled

Path	Both Controllers Active	Controller A Failed	Controller B Failed
1	B volumes, 207800C0FF0A408A A volumes, 217000C0FF0A408A	No change	Offline
2	A volumes, 207000C0FF0A408A B volumes, 217800C0FF0A408A	Offline	No change
3	A volumes, 207000C0FF0A408A B volumes, 217800C0FF0A408A	Offline	No change
4	B volumes, 217800C0FF0A408A A volumes, 207000C0FF0A408A	No change	Offline

69501 Switch Attach Configuration

When a data host is connected through one or more switches to controller host ports, either loop topology or point-to-point topology can be used. All controller module host ports must be set to use the same topology.

Whichever topology is used, each data host has dual-ported access to volumes through both controllers; the topology only affects how mapped volumes and port WWNs are presented if one controller fails.

- **Failover in a switch attach, loop configuration.** If one controller fails in a switch attach configuration using loop topology, the host ports on the surviving controller present the port WWNs for both controllers. Each controller's mapped volumes remain accessible. For example, if controller B fails, data access is essentially unchanged because controller A already had access to all mapped volumes before the failure.
- **Failover in a switch attach, point-to-point configuration.** If one controller fails in a switch attach configuration using point-to-point topology, the surviving controller presents its mapped volumes on its primary host port (FC0) and the mapped volumes owned by the failed controller on the secondary port (FC1).

  Unlike loop topology where both controllers' mapped volumes are presented on all host ports of the surviving controller, point-to-point topology requires host-based failover software to redirect access to the failed controller's mapped volumes through the surviving controller's secondary port. Further redundancy can be added by linking the fabric switches together, however this approach requires detailed path management configuration through the host adapter software.

For high availability, two data hosts can be connected through two switches to a dual-controller system. The host port interconnects must be disabled and host-based multipathing software is not required. Figure C-2 identifies the host paths in this configuration. For each path, Table C-4 specifies how mapped volumes and port WWNs are presented when both controllers are active and when either controller has failed.

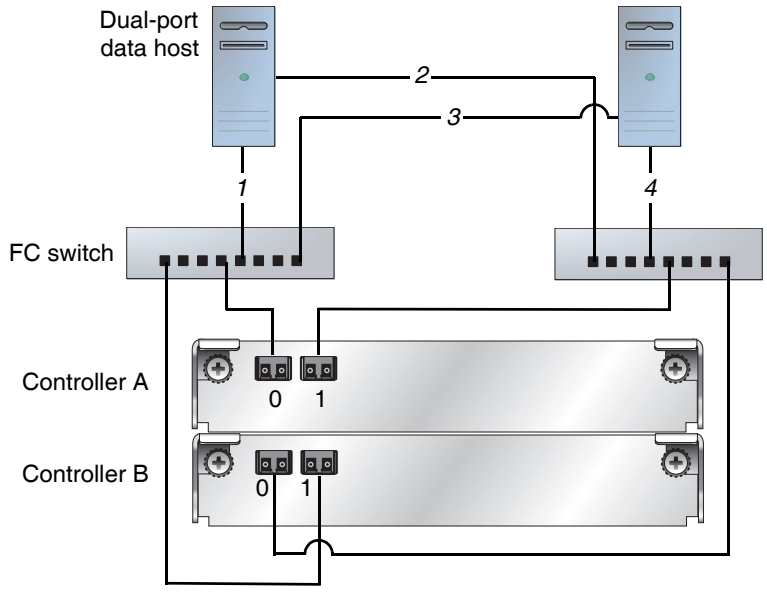

Figure C-2 Switch Attach Configuration with Two Switches and Two Hosts

Table C-4 Storage Presentation in a Switch Attach Configuration with Two Switches and Two Hosts

Path	Both Controllers Active	Controller A Failed	Controller B Failed
1	A volumes, 207000C0FF0A408A B volumes, 217800C0FF0A408A	B volumes, 217800C0FF0A408A	A volumes, 207000C0FF0A408A
2	A volumes, 217000C0FF0A408A B volumes, 207800C0FF0A408A	B volumes, 207800C0FF0A408A	A volumes, 217000C0FF0A408A
3	A volumes, 207000C0FF0A408A B volumes, 217800C0FF0A408A	B volumes, 217800C0FF0A408A	A volumes, 207000C0FF0A408A
4	A volumes, 217000C0FF0A408A B volumes, 207800C0FF0A408A	B volumes, 207800C0FF0A408A	A volumes, 217000C0FF0A408A

69503 Switch Attach Configuration

The high-availability configuration requires two gigabit Ethernet (GbE) switches, as shown in the following figure. During active-active operation, both controllers' mapped volumes are visible to both data hosts.

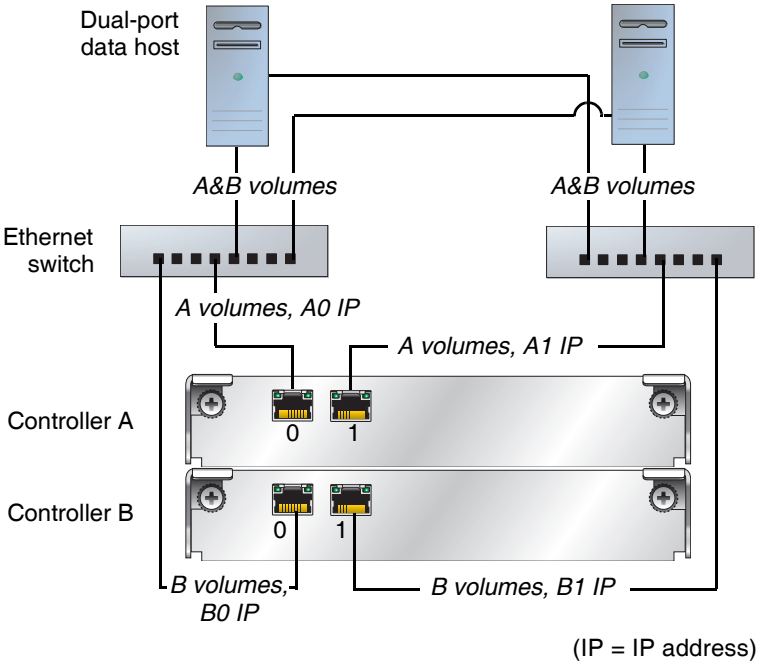

Figure C-3 iSCSI Storage Presentation During Normal, Active-Active Operation

A dual-controller 69503 storage system uses port 0 of each controller as one failover pair and port 1 of each controller as a second failover pair. If one controller fails, all mapped volumes remain visible to all hosts. Dual IP-address technology is used in the failed over state, and is largely transparent to the host system. However, for complete fault tolerance, host-based path failover software is recommended.

The following figure shows an example of storage presentation during failover of controller B.

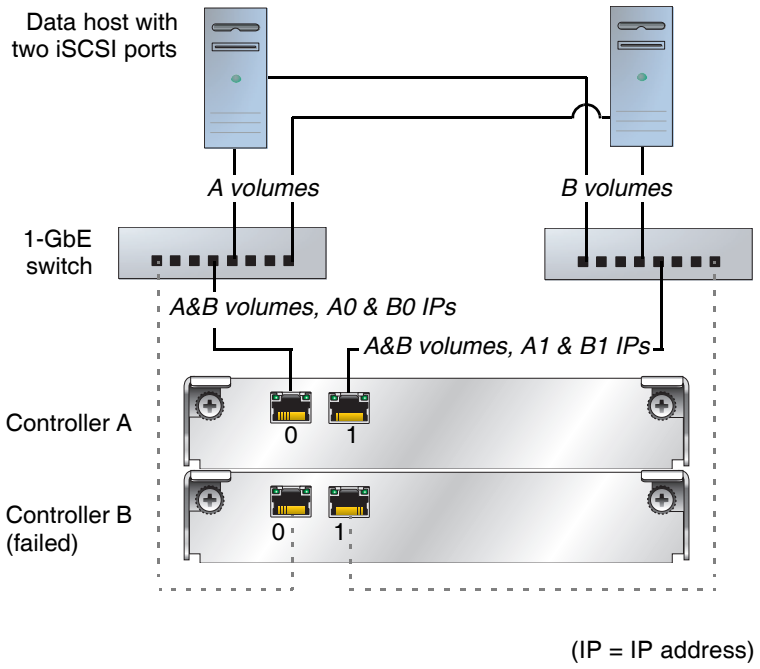

Figure C-4 iSCSI Storage Presentation During Failover

APPENDIX D

RAIDar Menu Reference

This appendix shows the RAIDar menu hierarchy. As described in "Introducing and Using RAIDar" on page 13, user configuration affects the RAIDar menu.

- "Standard and Advanced User Functions" on page 231 lists the RAIDar functions available to Standard and Advanced users.
- "Diagnostic User Functions" on page 238 lists the RAIDar functions available to Diagnostic users only.

If users do not have access to a function, the specified user type might be preventing access. You can increase access privileges as described in "Modifying Users" on page 27.

Standard and Advanced User Functions

RAIDar menu pages that Standard and Advanced users can access are listed in the following two tables. Pages in *italics* can be accessed by Advanced users only.

Table D-1 Monitor Menu – Standard and Advanced User Functions

Submenu	Page	See
Status	Status Summary	"Status Summary" on page 143
	Vdisk Status	"Virtual Disk Status" on page 144
	Module Status	"Module Status" on page 153
	Enclosure View	"Disk Drives by Enclosure" on page 150
	Enclosure Status	"Enclosure Status" on page 156
	Show Notification	"Displaying Notification Events" on page 177
	View Event Log	"Displaying the Event Log" on page 165

Table D-1 Monitor Menu – Standard and Advanced User Functions *(Continued)*

Submenu	Page	See
	Advanced Settings	
	• Controller Versions	"Controller Versions" on page 154
	• FRU Information	"FRU Information" on page 155
	• Disk Drive List	"Disk Drive List" on page 149
	• Host Port Status	"Host Port Status" on page 146
	• LUN Information	"LUN Information" on page 158
	• *Misc Configuration*	"Misc Configuration" on page 159
	• Expander Status	"Expander Status" on page 161
	• LAN Information	"LAN Information" on page 152
	• *Temperature Status*	"Temperature Status" on page 157
	• *Power Status*	"Power Status" on page 157
Statistics	Overall Rate Stats	"Rate Statistics for Virtual Disks" on page 169
	Cumulative Stats	"Cumulative Statistics for Virtual Disks" on page 170
	Volume Rate Stats	"Rate Statistics for Volumes" on page 170
	Cumulative Volume Stats	"Cumulative Statistics for Volumes" on page 171
	Real-Time Volume Stats	"Real-Time Statistics for Volumes" on page 172
	Disk Error Stats	"Disk Drive Error Statistics" on page 173
	Disk Usage	"Disk Space Usage Statistics" on page 174
	Reset All Statistics	"Resetting Statistics" on page 176
Help	Getting Started	"Help Menu" on page 23
	Subject Index	"Help Menu" on page 23
	Support Information	"Help Menu" on page 23

Table D-2 Manage Menu – Standard and Advanced User Functions

Submenu	Page	See
Virtual Disk Config	Vdisk Configuration	
	• Vdisk Status	"Virtual Disk Status" on page 67
	• Disk Drive Status	"Viewing Virtual Disk and Disk Drive Status Information" on page 67
	• Verify Virtual Disk	"Starting Virtual Disk Verification" on page 72
	• Expand Virtual Disk	"Expanding Virtual Disk Capacity" on page 69
	• Add Vdisk Spares	"Adding Vdisk Spares" on page 78
	• Delete Vdisk Spares	"Deleting Vdisk Spares" on page 79
	• Change Vdisk Name	"Changing a Virtual Disk Name" on page 75
	• Change Vdisk Owner	"Changing Virtual Disk Ownership" on page 74
	Create A Vdisk	"Creating a Virtual Disk Automatically" on page 61 and "Creating a Virtual Disk Manually" on page 63
	Delete A Vdisk	"Deleting a Virtual Disk" on page 75
	Abort A Vdisk Utility	"Stopping Virtual Disk Verification" on page 73
	Vdisk Utility Progress	"Checking the Progress of a Utility" on page 70
	Global Spare Menu	
	• Show Global Spares	"Displaying Global Spares" on page 80
	• Add Global Spares	"Adding Global Spares" on page 79
	• Delete Global Spares	"Deleting Global Spares" on page 80

Table D-2 Manage Menu – Standard and Advanced User Functions *(Continued)*

Submenu	Page	See
Volume Management	Volume Menu	
	• Volume Status	"Viewing Volume Status Information" on page 84
	• Add Volume	"Adding a Volume" on page 82
	• Delete Volume	"Deleting a Volume" on page 97
	• Expand Volume	"Expanding a Volume" on page 83
	• Change Volume Name	"Changing a Volume Name" on page 85
	• *Read Ahead Cache*	"Changing a Volume's Read-Ahead Cache Settings" on page 92
	• *Write Back Cache*	"Changing a Volume's Write-Back Cache Setting" on page 94
	Snapshot Services	"Using Snapshot Services" on page 98
	• Snapshot Overview	"Viewing Information About All Snap Pools, Master Volumes, and Snapshots" on page 114
	• Create Snap-Pool	"Creating a Snap Pool" on page 103
	• Create Master Volume	"Creating a New Volume as a Master Volume" on page 107
	• Set Snap-Pool Policy	"Setting Snap Pool Policies and Thresholds" on page 104
	• Snapshot-Enable Volume	"Converting a Standard Volume to a Master Volume" on page 108
	• Take Snapshot	"Taking a Snapshot" on page 109
	• Reset Snapshot	"Updating a Snapshot by Resetting" on page 110
	• Delete Snapshot	"Deleting a Snapshot" on page 113
	• Delete Modified Data	"Deleting Modified Data" on page 111
	• Rollback Volume	"Rolling Back a Master Volume" on page 112
	Volume-Copy Services	"Using Volume Copy Services" on page 117
	• Volume-Copy	"Copying a Volume" on page 117
	• Abort Volume-Copy	"Canceling a Volume Copy" on page 119
	• Volume-Copy Status	"Viewing the Status of a Volume Copy" on page 118

Table D-2 Manage Menu – Standard and Advanced User Functions *(Continued)*

Submenu	Page	See
	Volume Mapping	
	• Map Hosts To Volume	"Managing Volume Mappings" on page 90
	• Manage Host List	"Managing the Global Host Port List" on page 87
Scheduler	Manage Scheduler	"Using the Scheduler" on page 120
General Config	LAN Configuration	"Configuring Ethernet Management Ports" on page 42
	Host Port Configuration	"Configuring Host Ports" on page 35
	Enclosure Management	"Using the Enclosure Management Page" on page 136
	License Management	
	• Installed Licenses	"Viewing Installed Licenses" on page 31
	• Install A License	"Installing a License" on page 32
	Disk Configuration	"Enabling or Disabling SMART Changes" on page 131
	Services Security	"Configuring Network Management Services" on page 46; "Configuring In-band Management Services" on page 127
	User Configuration	
	• Modify Users	"Modifying Users" on page 27
	• Add Users	"Adding Users" on page 28
	• Delete Users	"Deleting Users" on page 30
	System Preferences	"Configuring Preferences" on page 24
	System Information	"Setting System Information" on page 34
	Set Date/Time	"Setting Date and Time" on page 34

Table D-2 Manage Menu – Standard and Advanced User Functions *(Continued)*

Submenu	Page	See
	System Configuration	• "Changing the Cache Mirroring Mode" on page 53 • "Managing Dynamic Spares" on page 77 • "Changing Auto-Write-Through Triggers and Behaviors" on page 96 • "Changing the Utility Priority" on page 181 • "Enabling and Disabling Background Scrub for Disks" on page 186 • "Disabling Partner Firmware Upgrade" on page 181 • "Controlling Host Access to the System's Write-Back Cache Setting" on page 187 • "Changing the Sync Cache Mode Option" on page 187 • "Changing the Missing LUN Response Option" on page 188
	Restore Defaults	"Restoring All Defaults" on page 185
Event Notification	Notification Summary	"Configuring Event Notification" on page 47
	Visual Configuration	"Configuring Visual Alerts" on page 49
	Email Configuration	"Configuring Email Alerts" on page 51
	SNMP Configuration	"Configuring SNMP Traps" on page 52
Utilities	Recovery Utilities	
	• Cache Data Status	"Clearing Unwritable Cache Data" on page 183
	• Vdisk Quarantine	"Dequarantining a Virtual Disk" on page 71
	Host Utilities	
	• *Reset Host Channel*	"Resetting Host Channels" on page 182
	Disk Drive Utilities	
	• *Rescan*	"Rescanning for Drive Changes" on page 182
	• Locate Disk Drive	"Illuminating a Drive Module LED" on page 132
	• Clear Metadata	"Clearing Metadata From a Disk Drive" on page 130
	• Display Disk Cache	"Viewing Disk Drive Read-Cache Status" on page 132

Table D-2 Manage Menu – Standard and Advanced User Functions *(Continued)*

Submenu	Page	See
	Configuration Utilities	
	• Show Changed Settings	"Viewing Changed Settings" on page 185
	• *Save Config File*	"Saving the Configuration to a File" on page 54
	• *Restore Config File*	"Restoring a Saved Configuration File" on page 184
	• *Custom Configuration*	*Phoenix Storage System Interface Customization Guide*
	Debug Utilities	
	• Save Logs to File	"Saving Log Information to a File" on page 166
	• *Debug Log Setup*	"Setting Up the Debug Log" on page 168
Restart System	Shut Down/Restart	"Restarting and Shutting Down a Controller" on page 55
Update Software	Controller Software	"Updating Software" on page 179
	Disk Drive Firmware	
	• Show Disk Drives	"Viewing Disk Drive Types and Firmware Versions" on page 133
	• Show Disk Drive Types	"Viewing Disk Drive Types and Firmware Versions" on page 133
	• Update Firmware	"Updating Disk Drive Firmware" on page 134
	Enclosure Firmware	
	• Show Enclosures	"Updating Expansion Enclosure Firmware" on page 140
	• Show Enclosure Types	"Updating Expansion Enclosure Firmware" on page 140
	• Update Firmware	"Updating Expansion Enclosure Firmware" on page 140

Diagnostic User Functions

The RAIDar menu options listed in the following table are available to Diagnostic Manage users for troubleshooting purposes. For information about using these functions, refer to the *Troubleshooting Guide*.

Table D-3 Manage Menu – Diagnostic User Functions

Submenu	Page	Option
General Config	Services Security	Service Interface
		Service Debug
	Restore Defaults	Restore Management Controller Defaults
Event Notification	Select Individual Events	
	• Critical Events	
	• Warning Events	
	• Info Vdisk Events	
	• Info Drive Events	
	• Info Health Status	
	• Info Status Events	
	• Info Config Events	
	• Info Misc Events	
	• Set/Clear All Events	
Utilities	Recovery Utilities	
	• Enable Trust Vdisk	
	• Trust Vdisk	
	Debug Utilities	
	• View Debug Log	
	• View Error Buffers	
	• View CAPI Trace	
	• View Mgmt Trace	
	Diagnostic Tools	
	• Expander Isolation	

APPENDIX E

RPD Drive Canister Installation Procedure

REQUIRED TOOLS: #1 PHILLIPS HEAD SCREWDRIVER

1. Orient the Drive Canister so the handle is facing you and offset to the right (The Phoenix logo is in the lower right corner).

2. Holding the canister by the handle (**handle oriented to the right**) place it into

the desired bay of the chassis enclosure.

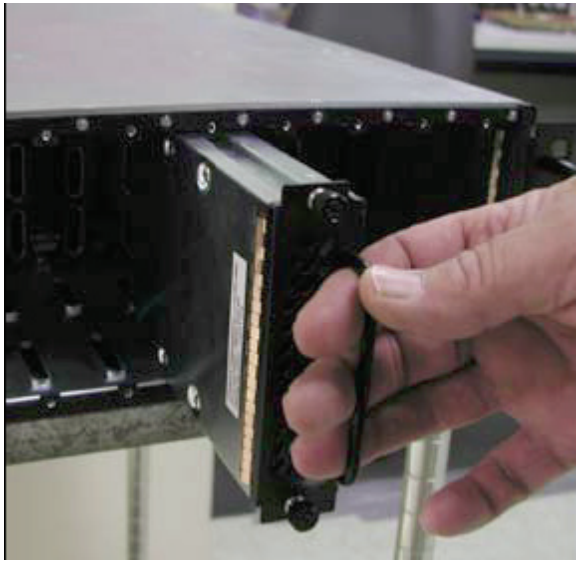

3. Slowly slide the canister all the way into the chassis until it bottoms out. A slight push on the handle or front of canister may be needed to fully mate the connectors.

4. Thread the top canister thumbscrew into the chassis 3 to 4 turns by hand.

5. Thread the bottom canister

thumbscrew into the chassis 3 to 4 turns by hand.

6. Go back to the top canister thumbscrew and thread it by hand until seated.

7. Thread the bottom canister thumbscrew by hand until seated.

8. Using a #1 phillips screwdriver, finish threading the top and bottom thumbscrews an additional quarter to half turn until both thumbscrews are fully seated. Do not exceed 6 inch/lbs.

Operation complete

Glossary

The glossary defines terms and acronyms used in Phoenix storage system documentation. Definitions obtained from the Storage Networking Industry Association (SNIA) Dictionary are indicated with "(SNIA)" at the end. For the complete SNIA Dictionary, go to www.snia.org/education/dictionary.

active-active Synonym for *dual active* components or controllers. A pair of components, such as the controllers in a failure tolerant storage subsystem that share a task or class of tasks when both are functioning normally. When one of the components fails, the other takes on the entire task. Dual active controllers are connected to the same set of storage devices, improving both I/O performance and failure tolerance compared to a single controller. (SNIA)

address A data structure or logical convention used to identify a unique entity, such as a particular process or network device.

AL_PA See *arbitrated loop physical address (AL_PA)*.

ANSI American National Standards Institute.

arbitrated loop physical address (AL_PA) An 8-bit value used to identify a participating device in an Arbitrated Loop. (SNIA)

API Application programming interface.

ARP Address Resolution Protocol.

array See *storage system*.

block The unit in which data is stored to or retrieved from a disk. For R/Evolution storage systems a block is 512 bytes, equivalent to the size of a disk sector.

broadcast write Technology that provides simultaneous caching of write data to both RAID controllers' cache memory with positive direct memory access acknowledgement (certified DMA).

cache A high speed memory or storage device used to reduce the effective time required to read data from or write data to a lower speed memory or device. Read cache holds data in anticipation that it will be requested by a client. Write cache holds data written by a client until it can be safely stored on more permanent storage media such as disk or tape. (SNIA)

See also *write-back cache*, *write-through cache*.

capacitor pack The controller module component that provides backup power to transfer unwritten data from cache to Compact Flash memory in the event of a power failure. Storing the data in Compact Flash provides unlimited backup time. The unwritten data can be committed to the disk drives when power is restored.

CAPI Dot Hill Configuration API.

channel A physical path used for the transfer of data and control information between storage devices and a RAID controller or a host; or, a SCSI bus in a controller module.

chassis An enclosure's metal housing.

chunk size The amount of contiguous data that is written to a virtual disk member before moving to the next member of the virtual disk. The default chunk size is 64 Kbyte. The number can be adjusted to improve performance. Generally, larger chunks are more effective for sequential reads.

CLI The command-line interface that system administrators can use to configure, monitor, and manage Phoenix storage systems. The CLI is accessible from any management host that can access a controller module through an out-of-band Ethernet or RS-232 connection.

controller The control logic in a storage subsystem that performs command transformation and routing, aggregation (RAID, mirroring, striping, or other), high-level error recovery, and performance optimization for multiple storage devices. (SNIA)

A controller is also referred to as a RAID controller.

controller enclosure An enclosure that contains disk drives and one or two controller modules. See *controller module*.

controller module A FRU that contains: a Storage Controller processor; a Management Controller processor; a SAS expander and Expander Controller processor; management interfaces; a LAN subsystem; cache protected by a capacitor pack and Compact Flash memory; host, expansion, management, and service ports; and midplane connectivity. If a controller enclosure contains redundant controller modules, the upper one is designated *A* and the lower one is designated *B*.

copy-on-write (COW) A technique for maintaining a point in time copy of a collection of data by copying only data that is modified after the instant of replicate initiation. The original source data is used to satisfy read requests for both the source data itself and for the unmodified portion of the point in time copy. (SNIA)

See also *snap pool*.

CPLD Complex programmable logic device. A generic term for an integrated circuit that can be programmed in a laboratory to perform complex functions.

CPU Central processing unit. The CPU is where most calculations take place, and the type of CPU in a controller module affects its performance capability. In Phoenix storage systems CPU is also referred to as the Storage Controller processor or the RAID controller processor.

DAS See *direct attach storage (DAS)*.

data host A host that reads/writes data to the storage system. A 69501 can be directly connected to multiple data hosts for direct attach storage (DAS). A 69501 or 69503 can be connected to multiple data hosts through switches for a storage area network (SAN).

data mirroring Data written to one disk drive is simultaneously written to another disk drive. If one disk fails, the other disk can be used to run the virtual disk and reconstruct the failed disk. The primary advantage of disk mirroring is 100 percent data redundancy: since the disk is mirrored, it does not matter if one of the disks fails; both disks contain the same data at all times and either can act as the operational disk. The disadvantage of disk mirroring is that it is expensive because each disk in the virtual disk is duplicated. RAID 1 and 10 use mirroring.

data striping The storing of sequential blocks of incoming data on all the different disk drives in a virtual disk. This method of writing data increases virtual disk throughput because multiple disks are working simultaneously, retrieving and storing. RAID 0, 3, 5, 6, 10, and 50 use striping.

DHCP Dynamic Host Configuration Protocol.

direct attach storage (DAS) A dedicated storage device that connects directly to one or more servers. (SNIA)

Supported for the 69501.

disk mirroring See *data mirroring*.

DMA Direct memory access.

drive module A FRU consisting of a disk drive and drive sled.

dynamic spare An available disk drive that is used to replace a failed drive in a virtual disk, if the Dynamic Spares feature is enabled and no vdisk spares or global spares are designated.

EC See *Expander Controller (EC)*.

ECC Error correcting code. (SNIA)

EIA Electronic Industries Alliance.

EMP See *enclosure management processor (EMP)*.

enclosure A physical storage device that contains disk drives. If the enclosure contains integrated RAID controllers it is known as a controller enclosure; otherwise it is an expansion enclosure.

enclosure management processor (EMP) An Expander Controller subsystem that provides data about an enclosure's environmental conditions such as temperature, power supply and fan status, and the presence or absence of disk drives.

Ethernet adapter An adapter that connects an intelligent device to an Ethernet network. Usually called an Ethernet network interface card, or Ethernet NIC. (SNIA)

Expander Controller (EC) The processor (located in the SAS expander in each controller module and expansion module) that is primarily responsible for enclosure management and SES.

expansion enclosure An enclosure that contains disk drives and one or two expansion modules. See *expansion module*.

expansion module A FRU that contains: a SAS expander and Expander Controller processor; host, expansion, and service ports; and midplane connectivity. If a system contains redundant expansion modules, the upper one is designated A and the lower one is designated B.

fabric A Fibre Channel switch or two or more Fibre Channel switches interconnected in such a way that data can be physically transmitted between any two N_Ports on any of the switches. (SNIA)

fabric port (F_Port) An F_Port that can support an attached arbitrated loop. An FL_Port on a loop has the AL_PA hex '00' and is the gateway to the fabric for NL_Ports on a loop.

fabric switch A Fabric switch functions as a routing engine that actively directs data transfer from source to destination and arbitrates every connection. Bandwidth per node via a Fabric switch remains constant when more nodes are added, and a node on a switch port uses a data path of up to 100 Mbyte/sec to send or receive data.

fabric-loop port (FL_Port) An F_Port can support an attached arbitrated loop. An FL_Port on a loop has the AL_PA hex'00', giving the fabric highest priority access to the loop. An FL_Port is the gateway to the fabric for NL_Ports on a loop.

failback See *recovery*.

failover In an active-active configuration, failover is the act of temporarily transferring ownership of controller resources from a failed controller to a surviving controller. The resources include virtual disks, cache data, host ID information, and LUNs and WWNs. See also *recovery*.

fault tolerance The capacity to cope with internal hardware problems without interrupting the system's data availability, often by using backup systems brought online when a failure is detected. Many systems provide fault tolerance by using RAID architecture to give protection against loss of data when a single disk drive fails. Using RAID 1, 3, 5, 6, 10, or 50 techniques, the RAID controller can reconstruct data from a failed disk drive and write it to a spare or replacement disk drive.

fault-tolerant virtual disk A virtual disk that provides protection of data in the event of a single disk drive failure by employing RAID 1, 10, 3, 5, or 50.

FC See *Fibre Channel (FC)*.

FC-AL See *Fibre Channel-Arbitrated Loop (FC-AL)*.

Fibre Channel (FC) A set of standards for a serial I/O bus capable of transferring data between two ports at up to 100 Mbyte/sec, with standards proposals to go to higher speeds. Fibre Channel supports point-to-point, arbitrated loop, and switched topologies. Fibre Channel was completely developed through industry cooperation, unlike SCSI, which was developed by a vendor and submitted for standardization after the fact. (SNIA)

Fibre Channel-Arbitrated Loop (FC-AL) A form of Fibre Channel network in which up to 126 nodes are connected in a loop topology, with each node's L_Port transmitter connecting to the L_Port receiver of the node to its logical right. Nodes connected to a Fibre Channel Arbitrated Loop arbitrate for the single transmission that can occur on the loop at any instant using a Fibre Channel Arbitrated Loop protocol that is different from Fibre Channel switched and point-to-point protocols. An arbitrated loop may be private (no fabric connection) or public (attached to a fabric by an FL_Port). (SNIA)

field-replaceable unit (FRU) An assembly component that is designed to be replaced on site, without the system having to be returned to the manufacturer for repair.

FRU See *field-replaceable unit (FRU)*.

Gbyte (GB) Gigabyte. Equivalent to 1000 Kbyte for data storage and statistics, or 1024 Kbyte for memory.

global spare A spare disk drive that is available to all virtual disks in a system.

HBA See *host bus adapter (HBA)*.

HIM Host interface module.

host bus adapter (HBA) An I/O adapter that connects a host I/O bus to a computer's memory system (SNIA).

host port A host-interface port on a controller module or an expansion module.

host port interconnect A dual-controller Fibre Channel enclosure includes host port interconnect circuitry which can be used to connect the host ports on the upper controller module to those on the lower controller module. When enabled, the port interconnect gives each host access to all the volumes assigned to both controllers and makes it possible to create a redundant configuration without using an external FC switch. The port interconnect should only be enabled when the system is used in direct attach configurations. When using a switch attached configuration, the port interconnect must be disabled.

hot swap The ability to remove and replace a FRU while the system is powered on and operational.

in-band management Transmission of a protocol other than the primary data protocol over the same medium as the primary data protocol. Management protocols are a common example of in-band transmission. (SNIA)

This type of access is available through use of the Dot Hill Configuration API (CAPI) to develop a programmed interface.

independent cache performance mode (ICPM) An operating mode in which a pair of controllers can process host I/Os and share disk channels but cannot fail over and assume responsibilities of a failed controller, because no mirroring of write-back cache occurs.

initialization The process of writing a specific pattern to all data blocks on all disk drives in a virtual disk. This process overwrites and destroys existing data on the disk drives and the virtual disk. Initialization is required to make the entire virtual disk consistent at the onset. Initialization ensures that virtual-disk verifications performed in the future are executed correctly.

I/O Input/output.

IP Internet Protocol.

iSCSI Internet Small Computer System Interface.

JBOD Just a Bunch of Disks. An expansion enclosure that is directly attached to a host.

Kbyte (KB) Kilobyte. Equivalent to 1000 bytes for data storage and statistics, or 1024 bytes for memory.

LAN See *local area network (LAN)*.

leftover drive A disk drive that contains metadata but is no longer part of a virtual disk.

local area network (LAN) Local Area Network. A communications infrastructure designed to use dedicated wiring over a limited distance (typically a diameter of less than five kilometers) to connect to a large number of intercommunicating nodes. (SNIA)

logical unit number (LUN) The SCSI identifier of a logical unit with a target. (SNIA)

For example, a LUN identifies the mapping between a storage system volume and a port on a switch or FC HBA.

loop address Indicates the unique ID of a node in FC loop topology. A loop address is sometimes referred to as a Loop ID.

loop port (L_Port) A "Loop" port is capable of performing arbitrated loop functions and protocols. NL_Ports and FL_Ports are examples of loop-capable ports. (SNIA)

loop topology See *Fibre Channel-Arbitrated Loop (FC-AL)*.

LUN See *logical unit number (LUN)*.

Management Controller (MC) The processor (located in a controller module) that is primarily responsible for human-computer interface and computer-computer interface functions, and interacts with the Storage Controller.

management host A workstation with direct or network connections to a storage system's management ports and that is used to manage the system.

management information base (MIB) A database of managed objects accessed by network management protocols. An SNMP MIB is a set of parameters that an SNMP management station can query or set in the SNMP agent of a network device (for example, a router).

master volume A volume that is enabled for snapshots. A master volume must be owned by the same controller as the associated snap pool.

Mbyte Megabyte (MB).

MC See *Management Controller (MC)*.

metadata Data in the first sectors of a disk drive that the system uses to identify virtual disk members.

MIB See *management information base (MIB)*.

network interface card (NIC) See *Ethernet adapter*.

NIC See *network interface card (NIC)*.

node port (N_Port) A port on a computer, disk drive, or other device through which the device does its FC communication.

node-loop port (NL_Port) An N_Port that can operate on FC-AL topology.

node WWN See *world wide node name (WWNN)*.

Non-RAID The RAID level option that can be used for a virtual disk having a single disk drive and that does not need the data redundancy or performance benefits of RAID. The capacity of a non-RAID virtual disk equals the capacity of its disk drive. For fault tolerance, use RAID 0 or above.

out-of-band management Method of accessing and managing a system using the RS-232 or Ethernet connection.

ownership In an active-active configuration, one controller has ownership of the following resources: virtual disks and vdisk spares. When a controller fails, the other controller assumes temporary ownership of its resources.

PHY Hardware component that converts between digital and analog in the signal path between the Storage Controller, Expander Controller, disk drives, and SAS ports.

PID Primary controller identifier number.

point-to-point Point-to-point is an alternative to FC-AL topology and is required in some fabric switch configurations. The controller enclosure supports point-to-point connections only to fabric ports (F_Ports). Loop topology is appropriate for most fabric switches, as it provides more flexibility when considering fault-tolerant designs.

port WWN See *world wide port name (WWPN)*.

power-and-cooling module A FRU that includes an AC power supply and two cooling fans. An enclosure has two power-and-cooling modules for failure tolerance and can operate with only one module.

priority Priority enables controllers to serve other I/O requests while running jobs (utilities) such as rebuilding virtual disks. Priority ranges from low, which uses the controller's minimum resources, to high, which uses the controller's maximum resources.

RAID Redundant Array of Independent Disks, a family of techniques for managing multiple disks to deliver desirable cost, data availability, and performance characteristics to host environments. (SNIA)

RAIDar The web browser interface that system administrators can use to configure, monitor, and manage Phoenix storage systems. RAIDar is accessible from any management host that can access a system through an out-of-band Ethernet connection.

RAID controller See *controller*.

RAS Reliability, availability, and serviceability. These headings refer to a variety of features and initiatives all designed to maximize equipment uptime and mean time between failures, minimize downtime and the length of time necessary to repair failures, and eliminate or decrease single points of failure in favor of redundancy.

rebuild The regeneration and writing onto one or more replacement disks of all of the user data and check data from a failed disk in a virtual disk with RAID level 1, 10, 3, 5, 6, and 50. A rebuild can occur while applications are accessing data on the system's virtual disks.

recovery In an active-active configuration, recovery (also known as failback) is the act of returning ownership of controller resources from a surviving controller to a previously failed (but now active) controller. The resources include virtual disks, cache data, host ID information, and LUNs and WWNs.

remote scripting CLI client A command-line interface (CLI) that enables you to manage the system from a remote management host. The client communicates with the management software through a secure out-of-band interface, HTTPS, and provides the same control and monitoring capability as the browser interface. The client must be installed on a host that has network access to the system.

rollback The process of resetting a volume's data to become identical to a snapshot taken of that volume.

SAN See *Storage Area Network (SAN)*.

SAS Serial Attached SCSI.

SATA Serial Advanced Technology Attachment.

SC See *Storage Controller (SC)*.

SCSI Small Computer System Interface. A collection of ANSI standards and proposed standards which define I/O buses primarily intended for connecting storage subsystems or devices to hosts through host bus adapters. (SNIA)

SCSI Enclosure Services (SES) An ANSI X3T10 standard for management of environmental factors such as temperature, power, voltage, etc. (SNIA)

In Phoenix storage systems, SES data is managed by the Expander Controller and EMP.

SFP Small form-factor pluggable connector, used in FC controller module host ports. An SFP is a FRU.

SID Secondary controller identifier number.

SMART Self-Monitoring Analysis and Reporting Technology. The industry-standard reliability prediction indicator for both the IDE/ATA and SCSI hard disk drives. Hard disk drives with SMART offer early warning of some hard disk failures so critical data can be protected.

SMI-S Storage Management Interface Specification.

SMTP Simple Mail Transfer Protocol. A protocol for sending email messages between servers and from mail clients to mail servers. The messages can then be retrieved with an email client using either POP or IMAP.

snap pool A volume that is configured to store snapshot data.

snapshot A fully usable copy of a defined collection of data that contains an image of the data as it appeared at the point in time at which the copy was initiated. (SNIA)

SNIA Storage Networking Industry Association.

SNMP Simple Network Management Protocol. An IETF protocol for monitoring and managing systems and devices in a network. The data being monitored and managed is defined by a MIB. The functions supported by the protocol are the request and retrieval of data, the setting or writing of data, and traps that signal the occurrence of events. (SNIA)

spare See *dynamic spare*, *global spare*, *vdisk spare*.

standard volume A volume that is not enabled for snapshots.

standby See *spare*.

state The current operational status of a disk drive, a virtual disk, or controller. A controller module stores the states of drives, virtual disks, and the controller in its nonvolatile memory. This information is retained across power interruptions.

Storage Area Network (SAN) A storage system consisting of storage elements, storage devices, computer systems, and/or appliances, plus all control software, communicating over a network. (SNIA)

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

storage system One or more enclosures, referred to in a logical (as opposed to physical) sense.

stripe size The number of data disks in a virtual disk multiplied by the chunk size.

sub-vdisk One of multiple RAID 1 virtual disks across which data is striped to form a RAID 10 virtual disk; or one of multiple RAID 5 virtual disks across which data is striped to form a RAID 50 virtual disk.

system See *storage system*.

Tbyte (TB) Terabyte. Equivalent to 1000 Gbyte for data storage and statistics, or 1024 Gbyte for memory.

TCP/IP Transmission Control Protocol/Internet Protocol.

topology The logical layout of the components of a computer system or network and their interconnections. Topology deals with questions of what components are directly connected to other components from the standpoint of being able to communicate. It does not deal with questions of physical location of components or interconnecting cables. (SNIA)

trap A type of SNMP message used to signal that an event has occurred. (SNIA)

UPS Uninterruptible Power Supply.

vdisk Abbreviation for virtual disk.

vdisk spare A disk drive that is marked as a spare to support automatic data rebuilding after a disk drive associated with a virtual disk fails. For a vdisk spare to take the place of another disk drive, it must be at least equal in size to the failed disk drive and all of the virtual disks dependent on the failed disk drive must be redundant—RAID 1, 10, 3, 5, 6, or 50.

VDS Virtual Disk Service. An API that enables virtual disks and volumes to be managed by third-party applications.

verify A process that checks the integrity of the redundant data on fault-tolerant virtual disks. For RAID 3, 5, 6, and 50, the verify process recalculates the parity of data stripes in each of the virtual disk's RAID stripe sets and compares it with the stored parity. If a discrepancy is found, an error is reported and the new correct parity is substituted for the stored parity. For RAID 1 and 10, the verify process checks for mirror mismatches. If an inconsistency is encountered, data is copied from the master disk drive to the slave disk drive. If a bad block is encountered when the parity is regenerated, the data is copied from the other disk drive, master or slave, to the reporting disk drive reallocating the bad block.

virtual disk For Phoenix storage systems, a set of disk drives that share a RAID level and drive type, and across which host data is spread for redundancy or performance.

volume A logical subdivision of a virtual disk. Multiple LUNs can be assigned to the same volume, one for each host port given access to the volume. See also *standard volume*.

volume mapping The process by which volume permissions (read only, read/write, or none) and LUNs are assigned to a host port.

VSS Volume Shadow Copy Service. An API that enables snapshots to be managed by third-party applications.

WBI See *RAIDar*.

web-browser interface (WBI) See *RAIDar*.

world wide name (WWN) A unique 64-bit number assigned by a recognized naming authority (often via block assignment to a manufacturer) that identifies a node process or node port. (SNIA)

Phoenix storage systems derive WWNs from the serial numbers of controller modules and expansion modules.

world wide node name (WWNN) A globally unique 64-bit identifier assigned to each Fibre Channel node process. (SNIA)

world wide port name (WWPN) A globally unique 64-bit identifier assigned to each Fibre Channel port. (SNIA)

write policy A cache-writing strategy used to control write operations. The write policy options are CIFS write-back and write-through cache.

write-back cache A caching technique in which the completion of a write request is signaled as soon as the data is in cache, and actual writing to non-volatile media occurs at a later time. Write-back cache includes an inherent risk that an application will take some action predicated on the write completion signal, and a system failure before the data is written to non-volatile media will cause media contents to be inconsistent with that subsequent action. For this reason, good write-back cache implementations include mechanisms to preserve cache contents across system failures (including power failures) and to flush the cache at system restart time. (SNIA)

This is how Phoenix storage systems operate. See also *write-through cache*.

write-through cache A caching technique in which the completion of a write request is not signaled until data is safely stored on non-volatile media. Write performance with a write-through cache is approximately that of a non-cached system, but if the data written is also held in cache, subsequent read performance may be dramatically improved. (SNIA)

Phoenix storage systems use write-through cache when write-back cache is disabled or when cache backup power is not working. See also *write-back cache*.

Index

D

E

F

G

H

I

L

M

T

U

V

W