Sun StorEdge™ RAID Manager 6.22 User's Guide

THE NETWORK IS THE COMPUTER™

Sun Microsystems, Inc.
901 San Antonio Road
Palo Alto, CA 94303-4900 USA
650 960-1300 Fax 650 969-9131

Part No. 806-0478-10
September 1999, Revision A

Send comments about this document to: docfeedback@sun.com

Contents

Figures

Tables

CHAPTER 1

Getting Started

This chapter describes the documents you received with your RAID Modules and how to use them.

About This User's Guide

This manual provides the information necessary to set up, monitor, tune, and repair your RAID Modules using the Sun StorEdge™ RAID Manager software. The manual is divided into seven chapters and an appendix.

- Chapter 1 provides an overview of the chapters contained in this book as well as descriptions of companion documentation.
- Chapter 2 describes the RAID Manager software, defines key terms used throughout this book, describes the different hardware configurations supported by the software, describes the types of component failures that may occur, and lists the routine tasks you need to perform to manage the RAID Modules.
- Chapter 3 describes the software features common to all applications.
- Chapter 4 describes how to configure the drives in the RAID Modules into logical units (LUNs) and hot spares.
- Chapter 5 describes how to check the status and performance of the RAID Modules.
- Chapter 6 describes how to recover after device failures. These failures may involve drives, controllers, fans, batteries, or power supplies. In some cases, you also need the information found in the hardware manuals accompanying the RAID Modules for specifics on how to do the physical tasks required for recovery.
- Chapter 7 describes the routine maintenance and tuning that you need to perform to manage the RAID Modules.
- Appendix A describes the command line interface to the storage management software.

Figures in the User's Guide

Many of the procedures in this manual show figures and diagrams to help explain how to perform the procedure. The exact appearance of the screen depicted in the figure may vary according to your operating system. This is especially true of buttons you need to click to select an option. For example, see FIGURE 1-1.

Note the difference in buttons between the new screens. How the screen appears on your system depends on your operating system. Functionally, however, the displays are the same. Regardless of the display, you use the mouse to click a button to select an option.

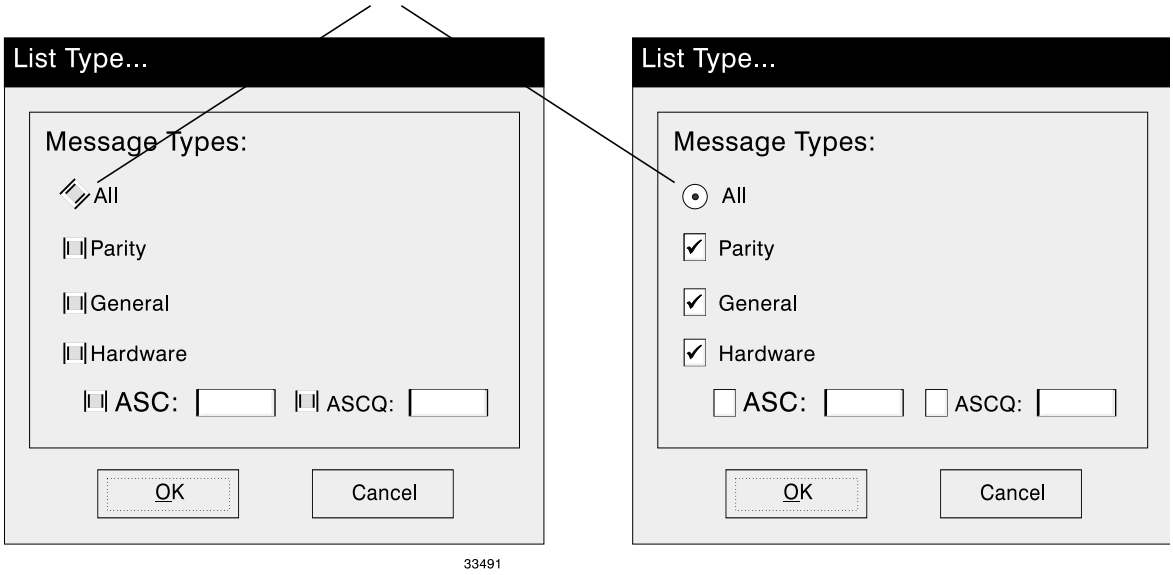

FIGURE 1-1 Screen Differences Between Operating Systems

Although the figures may not show the exact screen that appears on your system, the functionality is still the same as that described in the figure or step text.

About the Installation and Support Guide

In addition to this book, you also received the *Sun StorEdge RAID Manager Installation and Support Guide* and the *Sun StorEdge RAID Manager Release Notes* for the particular operating system you are using to access the RAID Modules.

The *Sun StorEdge RAID Manager Installation and Support Guide* contains:

- The installation procedure for installing the RAID Manager software on that operating system.
- Special restrictions on the use of the RAID Manager software on that operating system.
- Troubleshooting information for the RAID Manager software.
- Important information on controller firmware upgrading.
- Other notes on how the RAID Manager software runs in that specific operating system.

Refer to the *Sun StorEdge RAID Manager Installation and Support Guide* for installation instructions before performing any operations on the RAID Modules as described in this manual.

The *Sun StorEdge RAID Manager Release Notes* includes important information about the installation and operation of the Sun StorEdge RAID Manager software that was not available at the time this manual was published.

About the Hardware Guides

Along with your RAID Modules, you received a set of hardware guides describing how to install, operate, and service the RAID Modules and replace hardware components. These guides also contain information on troubleshooting the modules and determining which components have failed.

You may need these hardware guides to supplement the recovery information given here. They cover the *physical* replacement of the hardware components. The software procedures in this *User's Guide* work hand in hand with the hardware procedures in the hardware guides.

Always follow the software recovery procedures described in this *User's Guide*, if any, when replacing a hardware component. Use the hardware guides, if necessary, at the appropriate time to perform the physical replacement.

About Software Installation

The software installation procedure is described in the *Sun StorEdge RAID Manager Installation and Support Guide* that you received along with this manual. The *Sun StorEdge RAID Manager Installation and Support Guide* is specific to the operating system you are using. Refer to this manual for complete information on installation, and return to this *User's Guide* for information on how to run the installed software.

Note – Before you attempt to install the RAID Manager software, be sure to read the *Sun StorEdge RAID Manager Release Notes*.

What to Do Next

To set up and run the RAID Modules and the storage management software, you need to perform the following steps:

1. **Install the RAID Module hardware, using the information given in the hardware guides.**
2. **Install the RAID Manager software, using the information in the *Sun StorEdge RAID Manager Installation and Support Guide* and in the *Sun StorEdge RAID Manager Release Notes*.**
3. **If necessary, read Chapter 2 to become familiar with the RAID Manager software. Be sure to read the task summary list for a description of the routine tasks that you need to perform to manage the RAID Modules.**
4. **If necessary, read Chapter 3 to familiarize yourself with the common program features.**
5. **Go to Chapter 4 to configure the RAID Modules into logical units and hot spares, according to your system needs.**

CHAPTER 2

Overview

This chapter defines the terms used by the storage management software, describes the hardware configurations supported, describes component failures, and lists the tasks you must perform to maintain the RAID Modules.

Common Definitions

This section describes the concepts used by the storage management software. You need to understand these concepts to use your RAID Modules effectively.

Cache Memory

Cache memory is memory on the controller used for intermediate storage of read and write data. By using cache, you can increase system performance because the data for a read from the host may already be in the cache from a previous operation (thus eliminating the need to access the drive itself), and a write operation is completed once it is written to the cache, rather than to the drives.

When you create a logical unit (LUN), you can specify various caching parameters for it. If you need to change any caching parameters after LUN creation, use the Maintenance and Tuning application (see “Changing Cache Parameters” on page 167).

The following cache parameters are available through the storage management software. You can enable these parameters when creating LUNs using the Configuration application, or later using the Maintenance and Tuning application.

- Write caching – Allows write operations from the host to be stored in the controller’s cache memory. The use of write caching increases overall performance because a write operation from the host is completed when data is put in the cache, instead of when the data is actually written to the drive.
- Write cache mirroring – Allows cached data to be mirrored across two redundant controllers with the same cache size. The data written to the cache memory of one controller is also written to the cache memory of the other controller. Therefore, if one controller fails, the other can complete all outstanding write operations.
- Cache without batteries – Allows write caching to continue even if the batteries are discharged, not fully charged, or if no batteries are present. If you select this parameter without a UPS for protection, you could lose data if power fails.

You can change other cache parameters using the `rdacutil` utility. For more information about the `rdacutil` utility, see Appendix A.

Device Name

The storage management software uses device names to identify the controllers and logical units in the RAID Module. You use these names in the storage management software applications to access the logical units and controllers. In some operating systems, you can assign a name to these items.

Device names are determined by the location of the subsystem hardware, and vary according to the operating system you are using. For example, most UNIX operating systems use a *cXtXdXsX* naming scheme, and the Windows NT operating system uses a *driveX* naming scheme.

See the *Sun StorEdge RAID Manager Installation and Support Guide* for more information on the device names used by your operating system.

Drive Group

A drive group is a physical set of drives in the RAID Module. There are three types of drive groups:

- Unassigned drive group – drives in the RAID Module that have not been configured into logical units or hot spares. This drive group is only displayed in the Configuration application.
- Hot spare drive group – drives that have been assigned as hot spares. This drive group is only displayed in the Configuration application.
- Configured drive group – drives that have been configured into one or more logical units (LUNs) with the same RAID Level. A configured drive group is created the first time you create a logical unit on the selected drives. Each configured drive group is designated by a number. The drive groups are displayed by number in all applications. See the next section for a discussion of how drive groups are numbered.

Figure Figure 2-1 shows eight-drive trays organized into drive groups and LUNs.

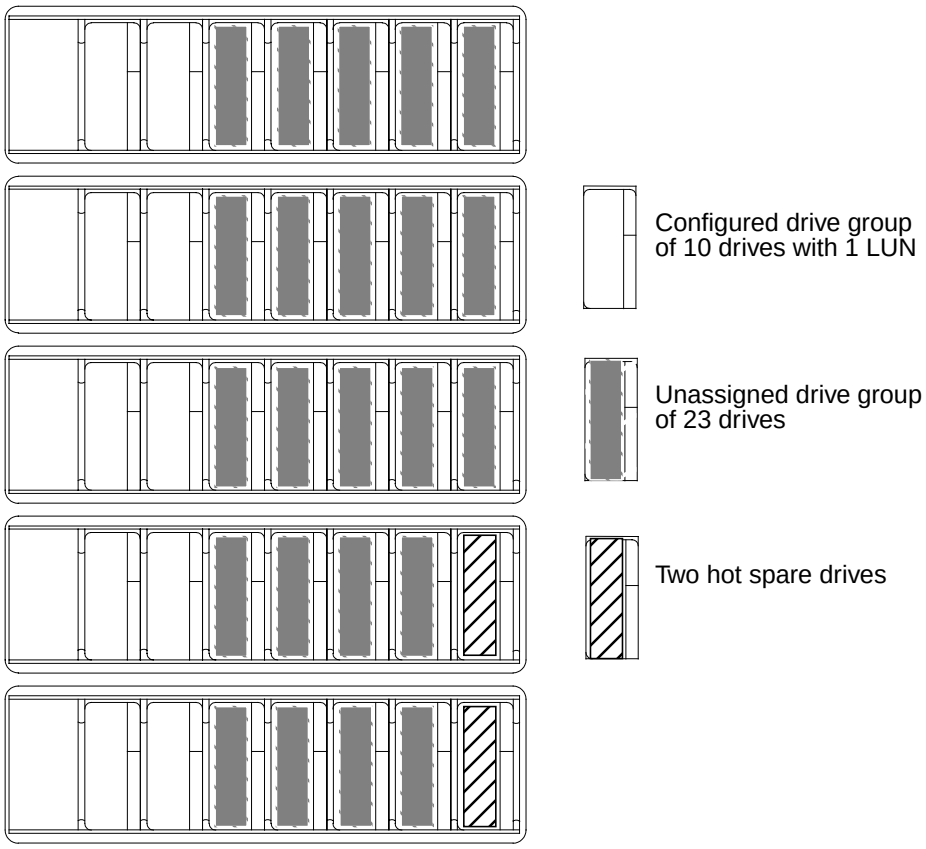

FIGURE 2-1 Drive Groups and Logical Units (LUNs)

Drive Group Renumbering

Each drive group in a RAID Module has a number from 1 to 32, depending on how many logical units (LUNs) are on your system and how you have configured them. (See the *Sun StorEdge RAID Manager Installation and Support Guide* for information on LUN limits.)

Drive groups are numbered according to the lowest numbered LUN in the group. For example, the drive group containing LUN 0 is always drive group 1.

Tip – Drive group numbering can change when you create and delete logical units.

For example, suppose you had the following drive groups:

Set*	Drive Group Number	LUN Number
A	1	0, 1
B	2	2
C	3	3, 4, 5

*"Set" is used here for purposes of this example; it is not a storage management software concept.

Now, you delete LUN 1. Because LUN 0 is the lowest numbered LUN in set A, renumbering would not occur. The drive groups would be as follows:

Set	Drive Group Number	LUN Number
A	1	0
B	2	2
C	3	3, 4, 5

Next, you create a new LUN in a new drive group (Set "D"). The new LUN uses the first available LUN number, LUN 1. Because 1 is a lower number than the lowest numbered LUNs in sets B and C (2 and 3), the drive groups would be renumbered as follows:

Set	Drive Group Number	LUN Number
A	1	0
D	2	1
B	3	2
C	4	3, 4, 5

As you can see, LUN 1 (Set D) becomes drive group 1. Set B is changed from drive group 2 to drive group 3, and Set C is changed from drive group 3 to drive group 4.

If you create another LUN on Set D, it will be LUN 6. This will *not* cause drive group renumbering, because the lowest LUN in the drive group is still 1.

Drive Names

In a RAID Module, drives are identified by a channel number and a SCSI ID. For example, a drive may be listed in a storage management software display as [1,3], indicating that it is channel 1, SCSI ID 3. The SCSI ID is independent of the Host Side ID of the controller.

See your hardware documentation for a description of the drive numbering used on your RAID Modules.

Hot Spares

A hot spare is a drive containing no data that acts as a standby in case a drive fails in a RAID 1, 3, or 5 logical unit. The hot spare drive adds another level of redundancy to your RAID Module. If a drive fails, the hot spare takes over for the failed drive until you replace it. When you replace the failed drive, data from the hot spare is copied back to the replaced drive, and the hot spare returns to a Standby status.

A hot spare drive is not dedicated to a specific drive group/logical unit, but can instead be used for any failed drive in the RAID Module with the same or smaller capacity. This is true even if the RAID Modules have independent controllers.

The hot spare drive group is displayed only in the Configuration application.

See “Creating Hot Spares” on page 93 for information on how to create a hot spare. See “Protecting Yourself From Component Failures” on page 25 for more information on the advantages of hot spares.

Logical Unit (LUN)

A logical unit (called a “LUN” for “logical unit number”) is the basic structure you create on the RAID Module to retrieve and store your data. A LUN is a set of physical drives that have been configured into either RAID level 0, 1, 3, or 5. Multiple LUNs may occupy the same set of drives (making up a drive group). See FIGURE 2-1 on 11 for an example of the relationship between drive groups and LUNs.

Each logical unit (not the drive group) is seen by the operating system as *one* drive. The operating system accesses each logical unit in the same way it accesses a standard disk drive, and you must follow the same procedures on the LUN (creating partitions, file systems, volumes, and so on) that you would use to access a standard disk drive.

Each logical unit has several properties:

- A logical unit number – This is a number (from 0 to 32, depending on your system) assigned to the logical unit by the controller when you create it.
- A RAID level – This determines how data is written to the drives and determines whether or not the logical unit has data redundancy (RAID level 1, 3, and 5 have data redundancy; RAID 0 does not.)
- A set number of drives – Each logical unit is made up of a set number of drives. The number of drives is determined when you create the logical unit and is dependent on the RAID level selection. For example, a RAID 1 LUN must have an even number of drives. You can add drives to an existing drive group/LUN, but you can’t decrease the number of drives in a drive group/LUN without deleting and recreating the drive group.
- A capacity – Each logical unit has a fixed amount of space. The capacity is set when you create the logical unit. The capacity depends both on the size of the drives and the selected RAID level. For example, a six-drive RAID 0 LUN has twice the capacity of a six-drive RAID 1 LUN, even if both LUNs use the same size drive (but RAID 1 has redundancy and RAID 0 does not).
- A set of parameters – These parameters include segment size, cache parameters, reconstruction rate, and so on. Some of these parameters can be changed for individual LUNs, while others apply to the whole drive group (one or more LUNs occupying the same drives).

Parity

Parity is additional information stored along with the data that enables the controller to rebuild (“reconstruct”) data after a drive failure. Only RAID 3 and 5 have parity. RAID 1 uses mirroring, not parity, which allows the controller to reconstruct data after a drive failure.

Parity Check / Repair

Parity check scans the segments in a logical unit and checks the parity for each segment. RAID 1 does not use parity, but parity check can still be performed on a RAID 1 unit. In this case, a parity check compares the data on the mirrored drives.

Parity repair corrects any parity inconsistencies found during parity check. However, parity repair corrects only parity errors. If the errors were caused by corrupted data, the data is still corrupted.

You can set a time to run an automatic parity check / repair by using Options → Auto Parity Settings in the Maintenance / Tuning application. You can also initiate a manual check / repair using the Recovery application. See“Running Parity Check” on page 153 for more information.

RAID Level

The RAID level determines how the controller reads and writes data and parity on the drives. The storage management software supports 4 RAID levels:

- RAID 0 – In RAID 0, data is striped across the drives in segments. There is no parity data, so RAID 0 uses the full capacity of the drives. However, there is no redundancy; therefore, if a single drive fails, all data on the logical unit is lost.
- RAID 1 – In RAID 1, data is mirrored. Each data drive has a corresponding mirrored drive with identical data. If one drive in the mirrored pair fails, the data from the other drive is used. In fact, a RAID 1 logical unit can suffer multiple drive failures without losing data, as long as none of the failed drives are mirrored pairs. Because RAID 1 mirrors data, a RAID 1 logical unit has only half the capacity of the assigned drives (for example, if you create a four-drive RAID 1 logical unit with 18 MB drives, the resulting data capacity is 2 x 18 = 36 MB).
- RAID 3 – In this version of the storage management software, RAID 3 is identical to RAID 5.

- RAID 5 – In RAID 5, data and parity are striped across the drives in the logical unit in segments. Because of this parity, if a single drive fails, data can be recovered from the remaining good drives. Two drive failures cause all data to be lost. A RAID 5 logical unit has the data capacity of all the drives in the logical unit less one (for example, a five-drive RAID 5 logical unit with 18 MB drives has a data capacity of 4 x 18 = 72 MB).

RAID Module

A RAID Module is a set of drives, one or two controllers (single active, active/passive, or active/active), and applicable power supplies, fans, batteries, and so on, that are physically connected together. You select a RAID Module to perform the various RAID tasks (such as configuring, obtaining status, recovering, and so on).

A RAID Module may be made up of one or more physical hardware components. For example, several drive trays connected to a controller tray make up one RAID Module. A unit that contains drives and a controller in its enclosure is also considered one RAID Module. A rackmount subsystem may contain more than one RAID Module.

RDAC Driver

The Redundant Disk Array Controller (RDAC) driver is part of the storage management software package and manages the I/O data connection for RAID Modules with redundant controllers. If a component (cable, controller, host adapter, and so on) fails along the connection, causing the host to lose communication with a controller, the RDAC driver automatically reroutes all I/O operations to the other controller.

Consult the *Sun StorEdge RAID Manager Installation and Support Guide* for your operating system for specific information.

Reconstruction

Reconstruction is the process of using data and parity on all operational drives within a drive group to rebuild the data on a replaced drive or a hot spare drive. Only data on a RAID 1, 3, or 5 logical unit can be reconstructed.

Reconstruction occurs when you replace a failed drive in a degraded RAID 1, 3, or 5 logical unit (a degraded logical unit has suffered a drive failure but is still operational).

The rate of reconstruction is determined by the Reconstruction Rate settings. The settings define how much processing time is allocated for reconstruction and for system performance. The faster the rate, the faster reconstruction occurs, but the slower system I/O is serviced.

Redundant Controllers

Most RAID Modules contain two controllers set up as redundant controllers. That is, if one controller fails, the other controller in the pair takes over the failed controller's functions, and the RAID Module continues to process data. You can then replace the failed controller, often without shutting down the RAID Module, to resume normal operation.

This feature involves more than just the controller. It concerns the entire data connection, the route data takes from the host system to the RAID Module controller. If any part of the connection fails (for example, if the cable connecting the two units fails), the controller redundancy feature reroutes I/O to the remaining good connection.

The redundant controller feature is managed by the RDAC software, which controls data flow to the controller pairs independent of the operating system. This software keeps track of the current status of the connections and can perform the switch-over without any changes in the operating system.

Whether your RAID Modules have the redundant controller feature depends on a number of things:

- Whether the hardware supports it – See the hardware documentation for your RAID Modules to determine whether the hardware supports the redundant controllers.
- Whether your operating system supports RDAC – See the *Installation and Support Guide* for your operating system to determine if it supports redundant controllers.
- How the RAID Modules are connected – See "RAID Module Configurations" on page 19 for information on the possible configurations and their effect on redundant controller support.

Redundant controllers can be configured in two ways:

- Active/active – In this configuration, both controllers handle I/O from the host. Each controller is assigned specific drive groups/LUNs to service. If one controller fails, the second controller takes over the failed controller's assigned drive groups/LUNs. This is the recommended configuration, because using two controllers for I/O is usually faster than using a single controller.
- Active/passive – In this configuration, one controller (the active controller) handles all the I/O from the host, and the other controller (the passive controller) is held as a spare, taking over if the active controller fails.

Segment Size

A segment is the amount of data the controller writes on a single drive in a logical unit before writing data on the next drive. Segment size is given in blocks of 512 bytes. For example, if the segment size is 128 blocks, the controller writes 128 blocks of data on drive 1, then 128 blocks of data on drive 2, and so on.

When you create a logical unit, the default segment size is the optimal size for that RAID level. You can change this default segment size when you create a logical unit using the Options selection. You can also change it later using the Modify LUN option in the Configuration Application.

SNMP

The Simple Network Management Protocol (SNMP) notification is an option that you may enable when installing this software. It directs the software to send remote notification of RAID events to a designated network management station (NMS) using SNMP traps.

See the *Sun StorEdge RAID Manager Installation and Support Guide* for your operating system for details on enabling this notification option.

RAID Module Configurations

The storage management software supports three main configurations from the host systems to the RAID Modules. These configurations may be over a SCSI or over a Fibre Channel connection. Be sure to consult the *Sun StorEdge RAID Manager Installation and Support Guide* specific to your operating system for details on hardware requirements required to use this software with that operating system.

Caution – No configurations or combinations are supported beyond those described in this section. Furthermore, the software's operation cannot be guaranteed to work as intended and described in the operating-system-specific *Sun StorEdge RAID Manager Installation and Support Guide*, this *User's Guide*, or the onlline help if other configurations are used.

Most RAID Modules supported by this software use active/active redundant controllers. Each drive group/LUN is owned by *only* one of the active controllers in a RAID Module. Furthermore, the combined total of LUNs configured for both controllers cannot exceed the maximum number of LUNs that the module can handle (that is, 8, 16, or 32) regardless of which configuration is used. For information on LUN limits per module, see your operating-system-specific *Installation and Support Guide*.

The following sections describe the three configurations supported:

Single-Host Configuration

In a single-host configuration, the host machine contains two host adapters, with each host adapter connected to one of the controllers in the RAID Module. The storage management software is installed on the host. The two connections are required for maximum RDAC failover support for redundant controllers (see FIGURE 2-2).

Note – This is the recommended configuration. It provides the fullest functionality and complete RDAC failover support with dual controllers.

Tip – Although this configuration also supports RAID Modules that have a single-controller or dual controllers on the same connection, you do *not* have complete RDAC data path protection with either of these configurations. The host adapter and cable become a single point of failure, and any data path failure could result unpredictable effects on the host system. For the greatest level of I/O protection, provide each controller in a RAID Module with its own connection to a separate host adapter in the host system.

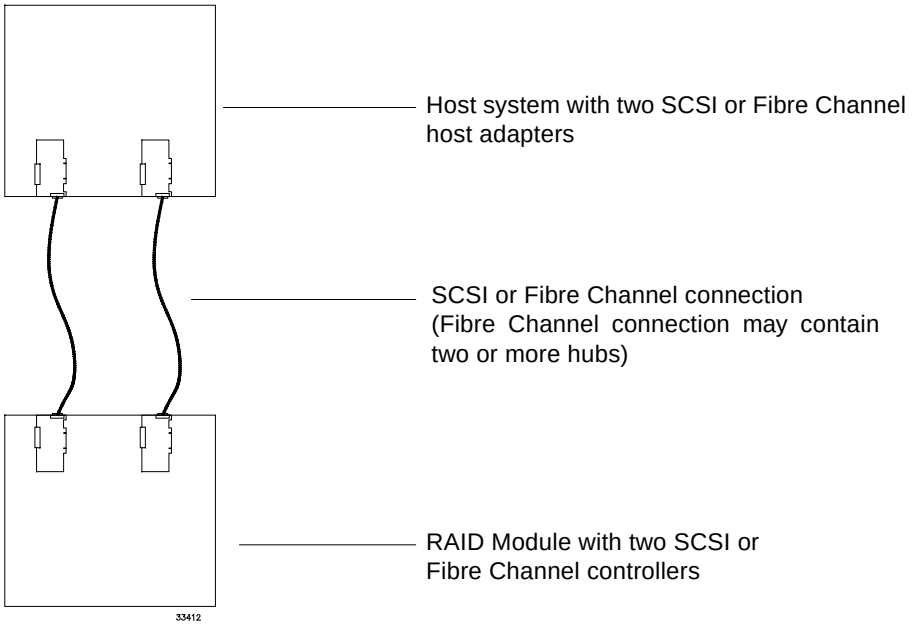

FIGURE 2-2 Single-Host to RAID Module Configuration

Multi-Host Configuration

In a multi-host configuration, two host machines are each connected by two connections to both of the controllers in a RAID Module (see FIGURE 2-3). The storage management software is installed on each host.

Caution – Not every operating system supports this configuration. Consult the restrictions in the *Sun StorEdge RAID Manager Installation and Support Guide* specific to your operating system for more information. Also, the host machines must be able to handle the multi-host configuration; refer to your hardware guides.

On some systems, the storage management software can only support multi-initiator configurations with Sun™ Cluster software. Refer to the *Sun Cluster Software Planning and Installation Guide*, or check with your local solution center, for multi-initiator issues.

Both hosts have complete visibility of both controllers, all data connections, and all configured drive groups/LUNs in a RAID Module, plus RDAC failover support for the redundant controllers. However, in this configuration, you must use caution when performing storage management tasks (especially creation and deletion of LUNs) to ensure that the two hosts do not send conflicting commands to the controllers in the RAID Modules.

The following items are unique to this configuration:

- Both hosts *must* have the same operating system and storage management software versions installed.
- Both host machines should have the same LUNs-per-host adapter capacity (that is, either both are limited to eight LUNs or both can have 16 to 32 LUNs). This is important for RDAC failover situations so that each controller can take over for the other and display all configured drive groups/LUNs.
- If the operating system on the host machine is capable of creating reservations, the storage management software will honor them. This means that each host could have reservations to specified drive groups/LUNs and *only* that host's software can perform operations on the reserved drive group/LUN. Without reservations, the software on either host machine is able to begin any operation. Therefore, you must use caution when performing certain tasks that need exclusive access. Especially when creating and deleting logical units, you should be sure to have only one configuration session open at a time (from only one host) or the operations could fail.
- This software does not provide failover protection at the host level. Failover protection requires third-party software.

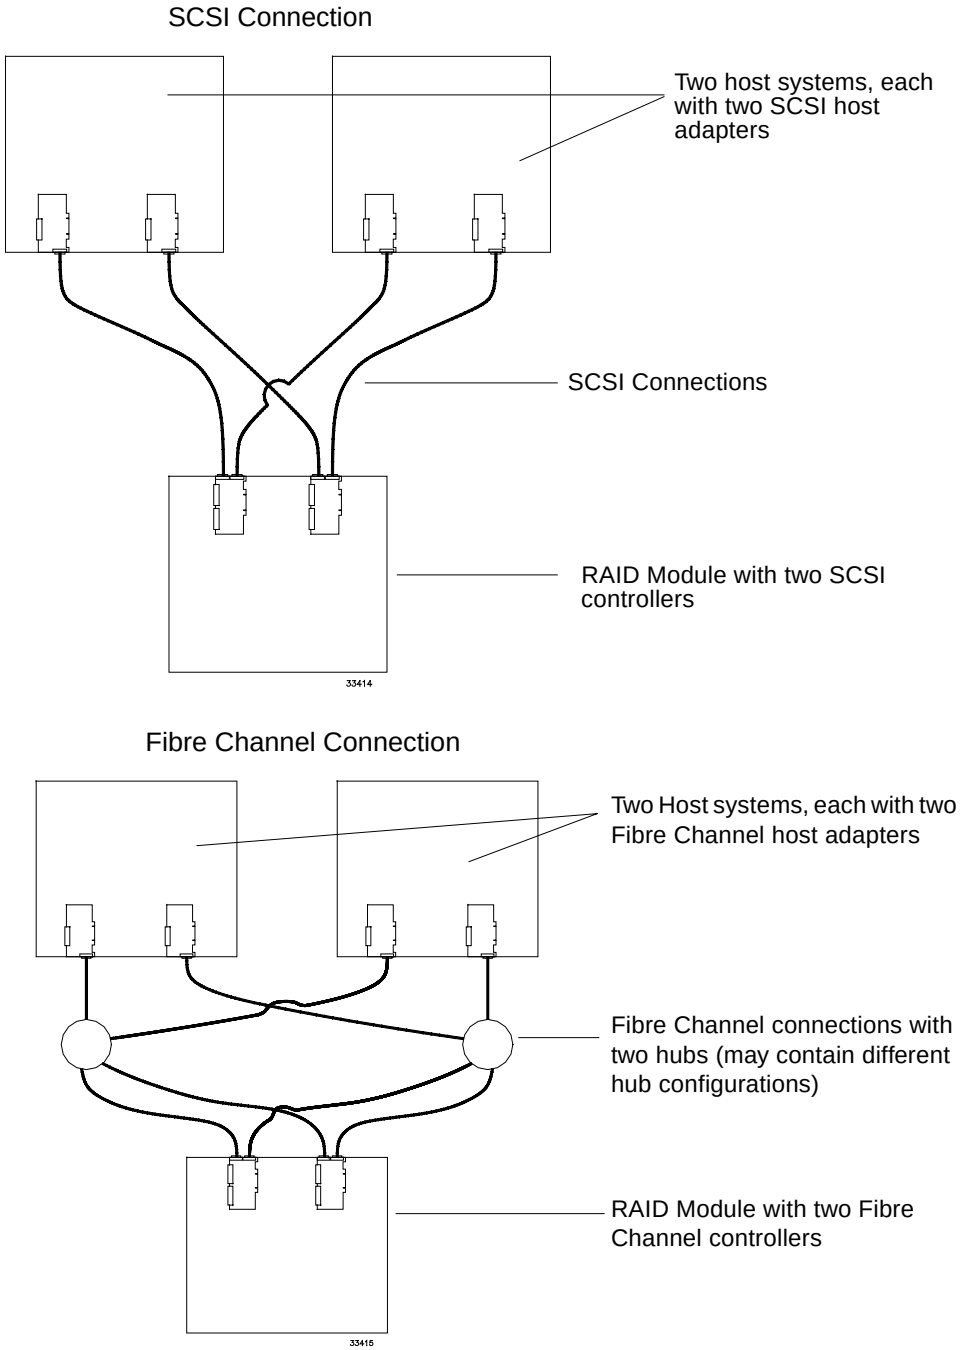

FIGURE 2-3 Multi-host to RAID Module Configuration

Independent Controller Configuration

In the independent controller configuration, two host machines are connected to a dual-controller RAID Module. One host machine is connected to one RAID Module controller, and the second host machine is connected to the other controller (see FIGURE 2-4.) The storage management software is installed on each host. Each host machine sees the controller and the drive groups/LUNs that it owns as independent of the other controller. Therefore, each host machine acts as if it were connected to a single-controller RAID Module.

Also, when an independent controller is selected with Select Module, the storage management software has knowledge of the alternate controller and displays all configured drive groups/LUNs. However, it only reports real-time statuses for the host-controller data connection on which it is installed and reports an "Inaccessible" status for drive groups/LUNs owned by the alternate controller.

Note – It is possible to open a Configuration application from both hosts. Therefore, when creating and deleting LUNs be sure that you have only one configuration session open at a time (from only one host) or the operation will fail for one of the hosts.

The following items are unique to this configuration:

- Both hosts *must* have the same operating system and the same storage management software versions installed.
- Both host machines should have the same LUNs-per-host adapter capacity (that is, either both are limited to eight LUNs or both can have 16 to 32 LUNs).
- This configuration is indicated by a special setting in Select Module (the "Indep. Cntrls?" column says Yes).
- The controllers in the RAID Module do not have RDAC failover protection.
- The storage management software reports the alternate controller and its drive group/LUNs as "Inaccessible."
- Recovery Guru/Health Check (Recovery application) detects data connection-related failures *only* for the controller that is connected to that host machine running the storage management software. For example, if host 1 has a controller failure, host 1 reports the failure, but host 2 does *not* report a controller failure when you run its Health Check. Also, these applications detect drive-related failures *only* for configured drive groups/LUNs that are owned by the controller (connected to that host machine) or for any unassigned or hot spare drive.
- Hot spares and unassigned drives are always available to LUNs owned by either controller.

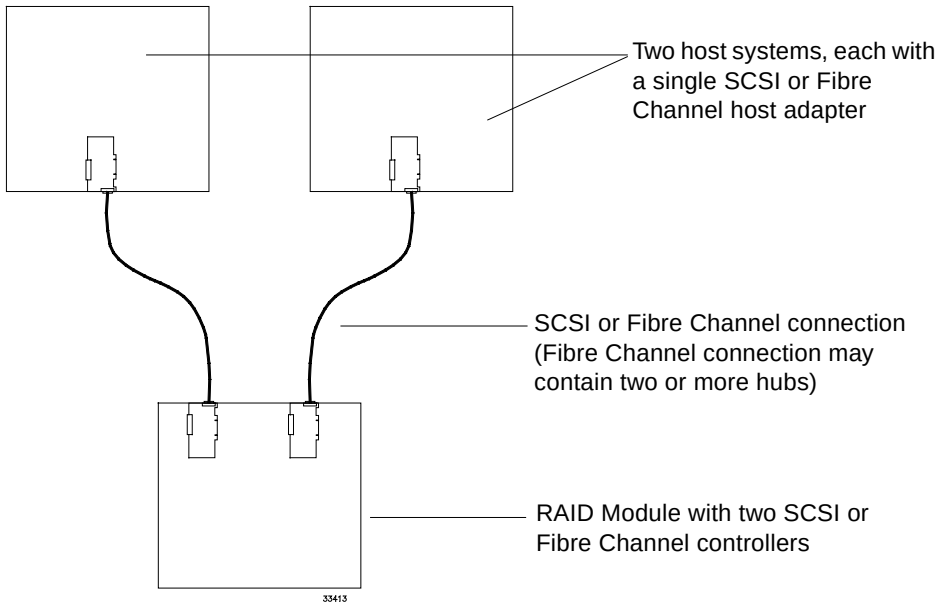

FIGURE 2-4 Independent Controller Configuration

Component Failures

The RAID Modules are designed to be fault tolerant. All modules using 1, 3, or 5 LUNs provide drive redundancy, so that if a drive fails, the data contained on the logical unit is still available. In most cases, the modules also feature redundant controllers, fans, and power supplies as well, so that the RAID Modules can keep functioning even if one of these components fails. The controllers in some RAID Modules also feature caching, to both speed up I/O operations and to preserve data in the event of a controller failure and battery backup, to prevent cache data loss in the event of a power failure. See the hardware guides accompanying your RAID Modules for complete information on the redundant features of your hardware.

The storage management software provides notification of most component failures. When you replace the failed component, follow the instructions given in Recovery Guru. For more information on component failures, see Chapter 6, "Recovery".

In some cases, you may need to consult the hardware guides for information on how to identify and replace the failed component. This usually does not involve the software. For example, there is no procedure for replacing a failed power supply given in this *User's Guide*.

Protecting Yourself From Component Failures

Using the features of the storage management software, you can take steps to further increase your protection against component failures.

- Use RAID levels 1, 3, or 5. These RAID levels offer data redundancy and can survive a single drive failure without losing data. RAID 0 does not have redundancy. A RAID 0 logical unit loses all data if a single drive fails.
- Use sufficient hot spares. A hot spare automatically replaces a failed drive in a RAID 1, 3, or 5 logical unit (assuming sufficient time between failures). Therefore, a hot spare protects your data from *two* drive failures instead of one. Two hot spares protect your data from three drive failures, and so on.
- Use write cache mirroring, if available on the RAID Module. This feature is available only with active/active controller pairs. Cache mirroring stores cache information for a controller on both controllers in the pair, so if a single controller in the pair fails, the cache information is still present in the other controller, and will be written to the logical units after controller failover occurs.

- When creating logical units, try to assign drives on different drive channels, if possible. This prevents a channel failure from preventing access to multiple drives. This may require you to use the Options selection when creating logical units. Note that this may not be possible if you create logical units with more drives than there are drive channels.

Application Summary Charts

The storage management software consists of four separate applications:

- Configuration
- Recovery
- Status
- Maintenance and Tuning

These applications can be run independently of each other or at the same time.

The following pages contain charts showing the functions in each application. Use the charts as a quick reference to review each application.

For a list of tasks, see “Task List” on page 31.

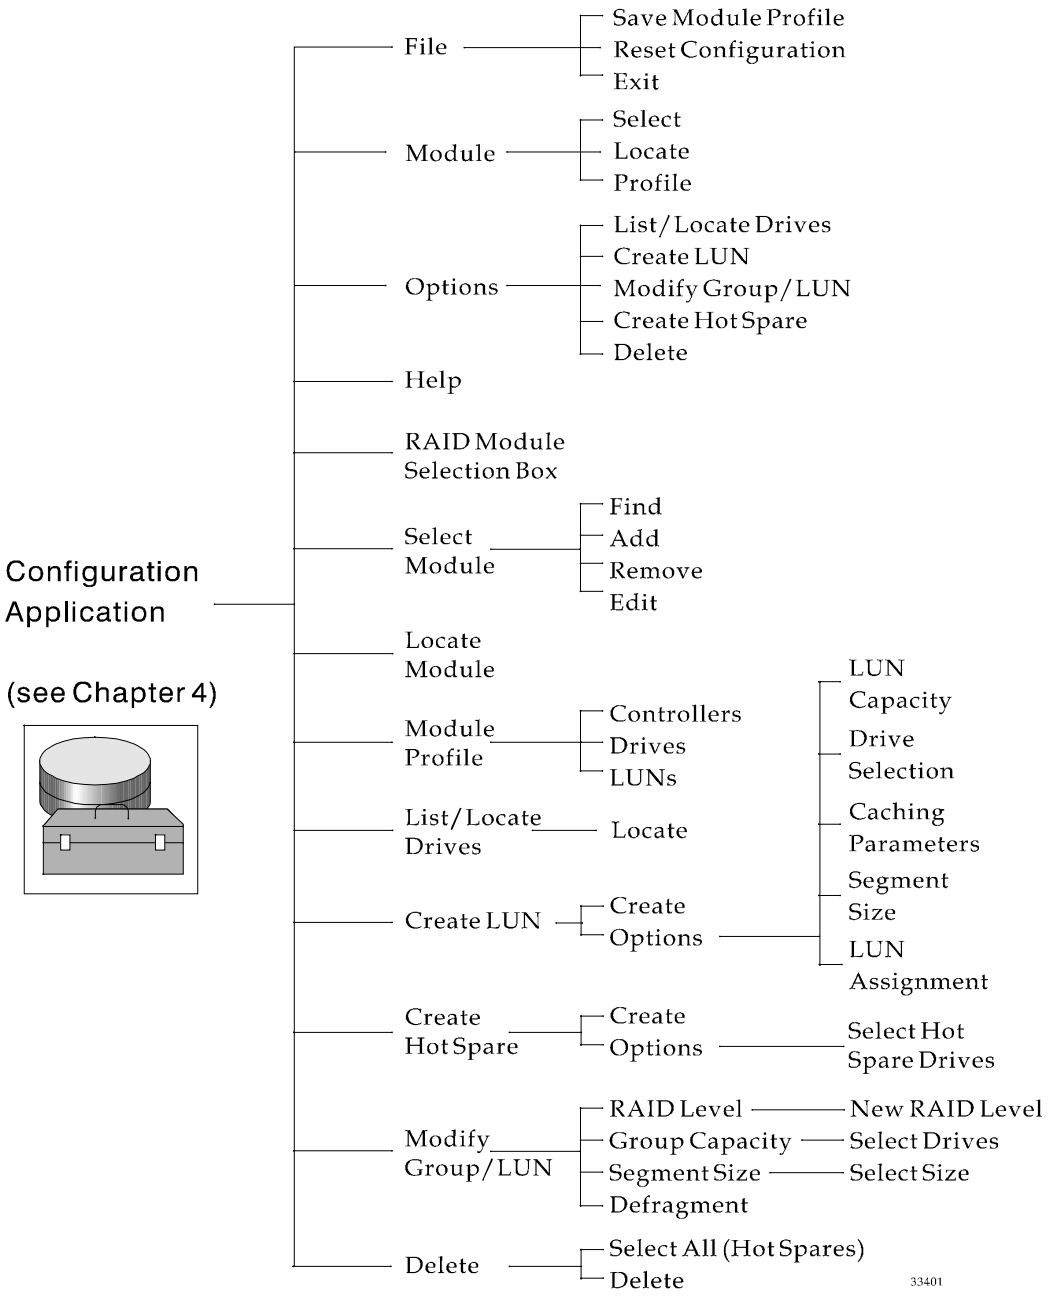

FIGURE 2-5 Configuration Application Summary

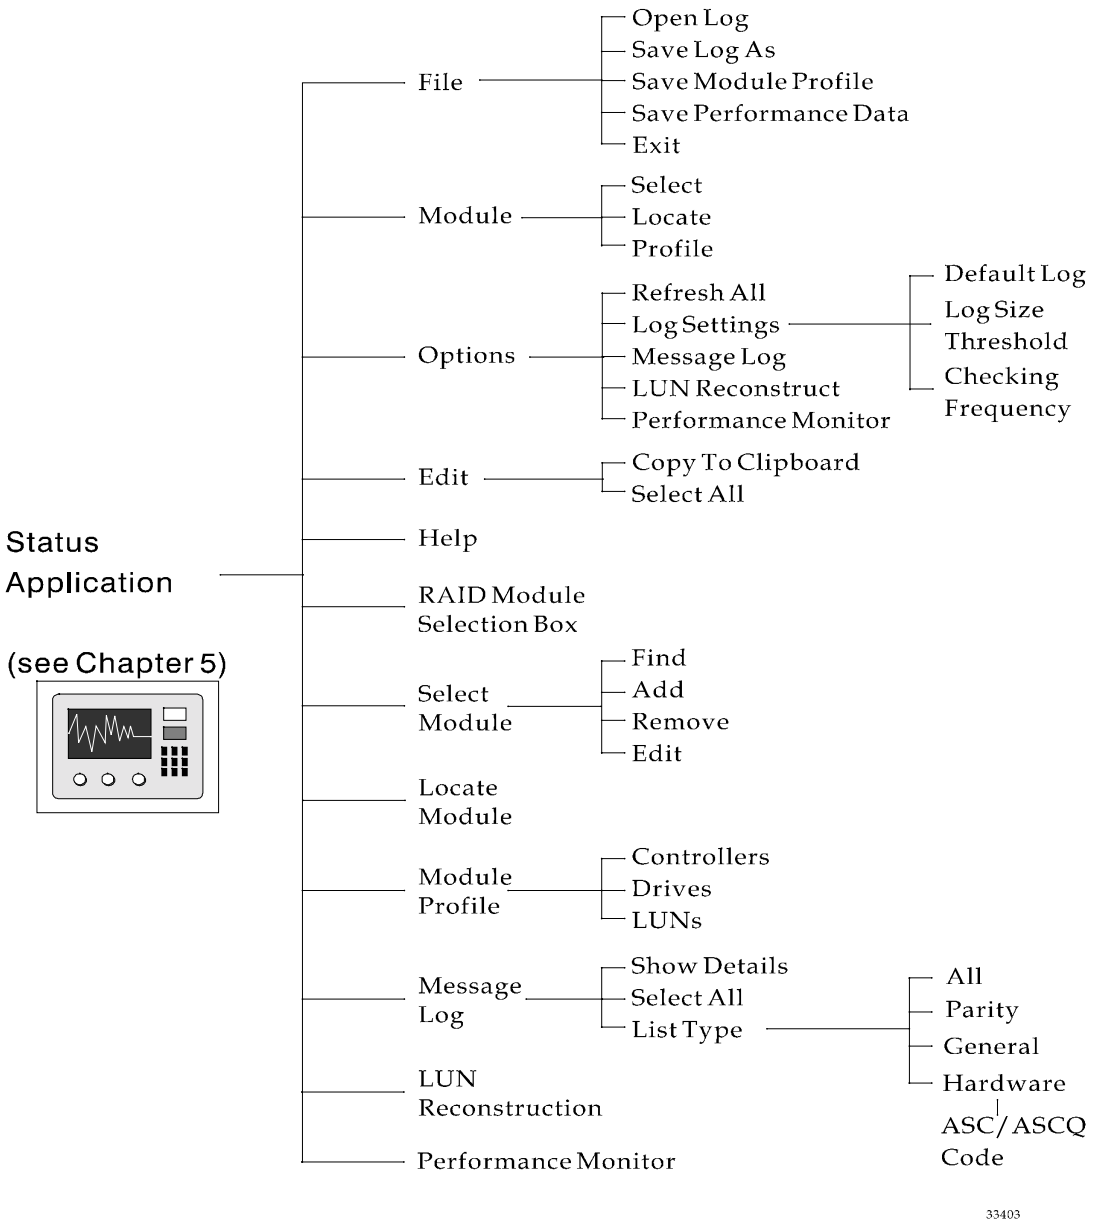

FIGURE 2-6 Status Application Summary

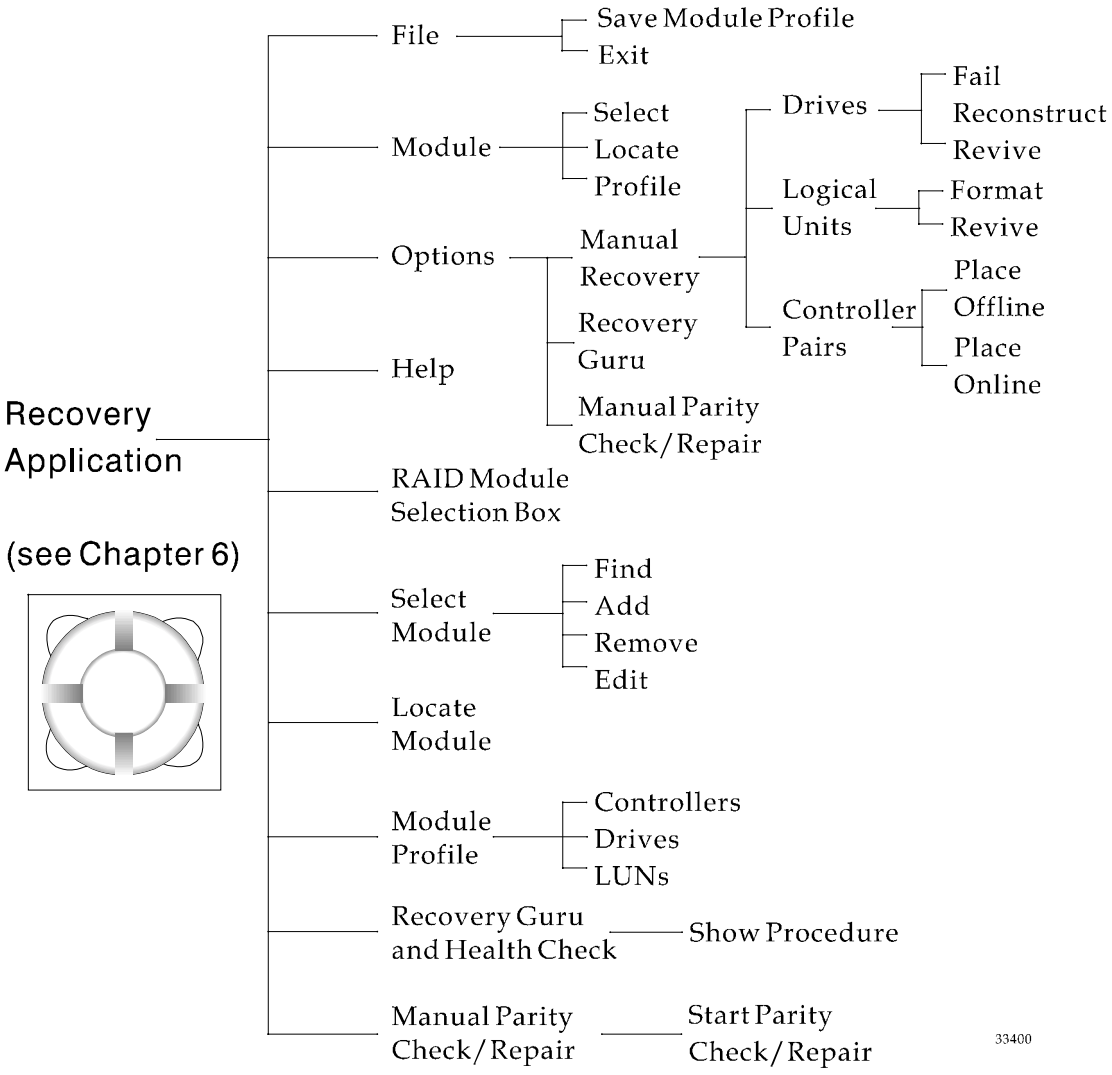

FIGURE 2-7 Recovery Application Summary

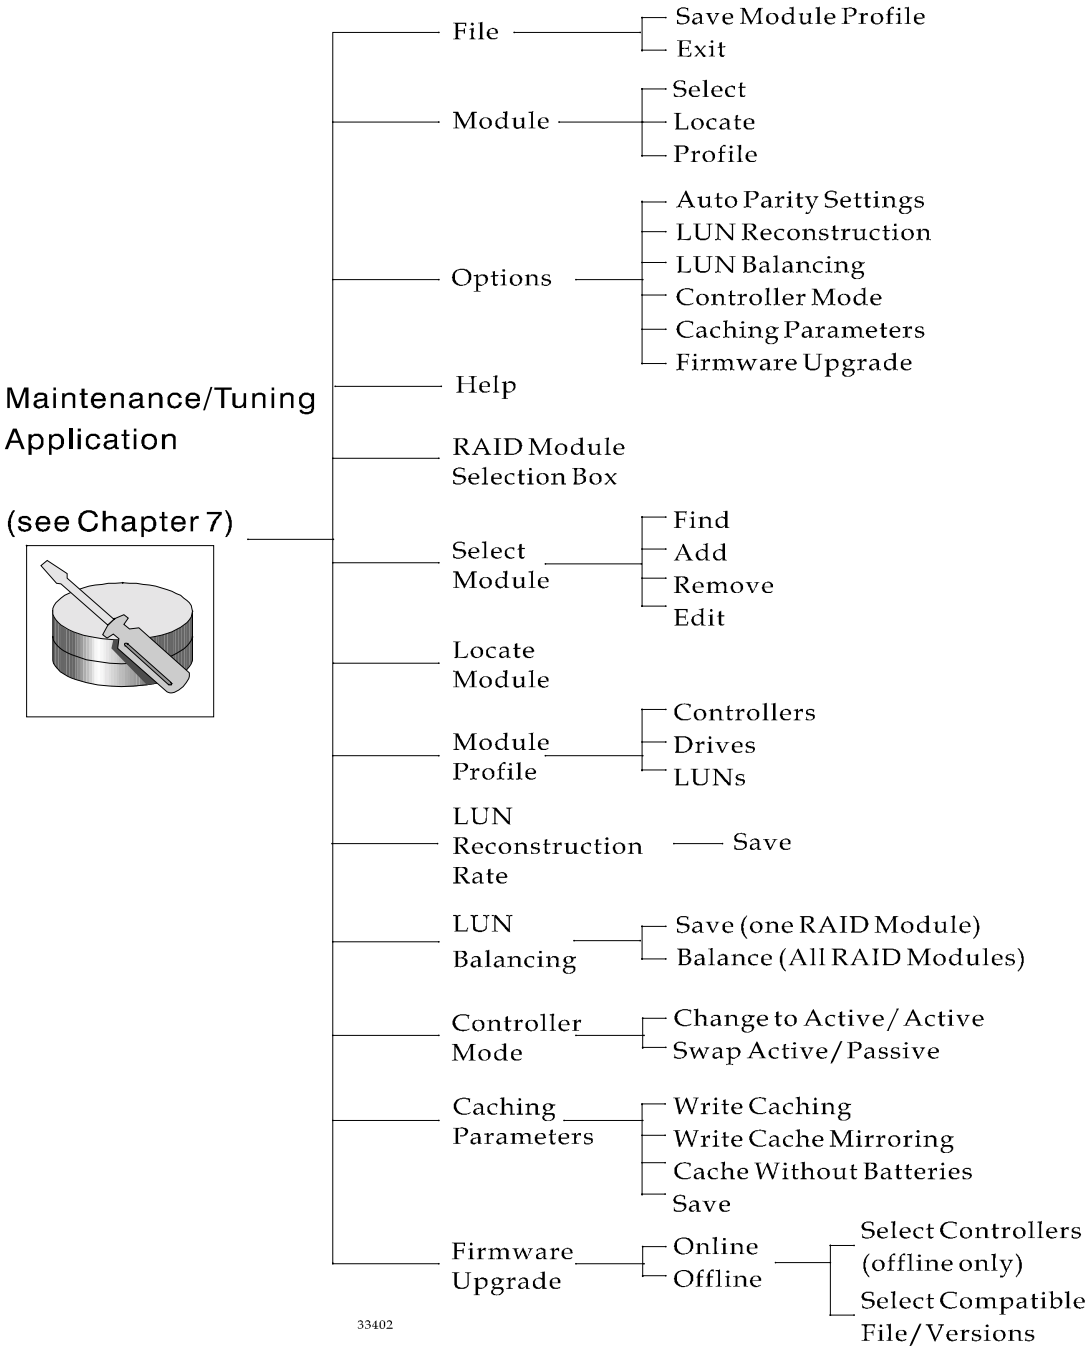

FIGURE 2-8 Maintenance and Tuning Application Summary

Task List

TABLE 2-1 lists the tasks you must perform and when to perform those tasks in order to manage and support the RAID Modules.

TABLE 2-1 Storage Management Task List

Task	When to Perform	Reference
Informational Tasks (Locating Raid Modules, Checking Configurations)		
Locate a RAID Module	When you want to physically locate a RAID Module. Perform this task when you first install the module and mark the module for future reference.	"Locating a Module" on page 44
Locate a logical unit or specific drives	When you want to physically locate a drive or a logical unit.	"Locating Drive Groups and Drives" on page 45
Assign names to RAID modules and controllers	To make the RAID Module components easily identifiable. Only the networked version allows you to name controllers.	"Selecting a Module" on page 41
View RAID module configuration information	When you need information on your RAID Modules, such as current firmware version or drive information.	"Viewing a Module Profile" on page 47
Check the Status Log	When you want to see if any errors have occurred on the RAID Module.	"Viewing the Status Log" on page 104
Run the Performance Monitor	To gather information on the performance of the RAID Module, with the goal of discovering any changes that could be made to speed performance.	"Message Log Save Message" on page 111
Configuration Tasks (Creating and Modifying Logical Units)		
Create logical units (LUNs)	After you install a RAID Module.	"Creating Logical Units" on page 69
Create a hot spare	When you configure a RAID Module. You should always have at least one hot spare on your system to protect against failures.	"Creating Hot Spares" on page 93
Add drives to an existing drive group	When you want to add capacity in a drive group.	"Modifying Drive Groups and Logical Units" on page 80

TABLE 2-1 Storage Management Task List *(Continued)*

Task	When to Perform	Reference
Change the RAID level of an existing drive group	If you want a different configuration in your RAID Module after you have created a logical unit.	"Modifying Drive Groups and Logical Units" on page 80
Change the segment size of existing logical units	To change the striping on all the logical units in a drive group. This may improve performance.	"Modifying Drive Groups and Logical Units" on page 80
Delete a logical unit	If you want to make some drives available to reconfigure your system. Note that this deletes all data from the logical unit.	"Deleting Logical Units" on page 91
Delete a hot spare	If you need the hot spare drive for a new logical unit.	"Deleting a Hot Spare" on page 96
Maintenance Tasks (Loading new firmware, changing parameters)		
Download controller firmware	• When you receive new firmware from your local solution center. • When adding RAID Modules with new firmware into an existing site (you may need to upgrade the existing controllers).	"Upgrading Controller Firmware" on page 143
Download NVSRAM	• When you receive new firmware from your local solution center. • When adding RAID Modules with new firmware into an existing site (you need to upgrade the existing controllers).	"Upgrading Controller Firmware" on page 143
Reset the configuration	If all attempts to use the RAID Module have failed. This is a last resort option that will delete all data and configuration information from your RAID Module and return it to a default configuration.	"Resetting the Configuration" on page 98
Change controller modes	If the controllers in a RAID Module are currently active/passive and you want to change them to active/active (recommended).	"Changing Controller Mode" on page 165
Balance drive groups/logical units across controllers	• If you have changed controller mode from active/passive to active/active and did not balance the logical units at that time. • If you want to reassign logical units to controllers.	"Balancing Drive Groups/LUNs" on page 160
Change cache parameters	To turn cache on and off after you have created a logical unit.	"Changing Cache Parameters" on page 167

TABLE 2-1 Storage Management Task List *(Continued)*

Task	When to Perform	Reference
Change the reconstruction rate	To change the speed at which data is reconstructed after a drive replacement and the speed used to format logical units.	"Changing the Reconstruction Rate" on page 158
Set the time for Automatic Parity Check/Repair	If you want to change the default time for automatic parity check. You might want to do this if the current time is interfering with system operations.	"Setting Automatic Parity Time" on page 154
Determine whether Automatic Parity Check will repair any errors found.	You may want to turn off the repair option to speed up parity check, and to give yourself a chance to determine which files are affected by the parity error.	"Setting Automatic Parity Time" on page 154
View the results of Automatic Parity Check/Repair	If you want to see if automatic parity check/repair detected and repaired any parity errors.	"Viewing the Status Log" on page 104
Recovery and Repair Tasks (recovering from failures, checking data)		
Perform a Health Check	If you have indications that a component failure or other error has occurred. Health Check locates the problem and gives information on how to correct it.	"Running Recovery Guru/Health Check" on page 128
Replace a failed component	If Health Check has detected a failed component. Do not replace a failed component unless the Recovery Guru recommends it. If necessary, see the manuals accompanying your hardware for information on how to replace a failed component.	Chapter 6, "Recovery"
Run Manual Parity Check	If Recovery Guru indicates that you should run parity check after a component failure, or you have indications that I/O errors have occurred.	"Running Parity Check" on page 153

What's New?

If you are upgrading from an earlier version of Sun StorEdge RAID Manager, the following new features are available:

- **Completely different user interface** – The user interface is considerably changed from previous versions of the storage management/RAID Manager software.
- **Performance Monitor** – A software option is provided to track logical unit and controller usage.
- **Dynamic RAID changes** – You can now change the capacity of a Drive Group and the RAID level or segment size of an existing logical unit without deleting the logical unit (with controller firmware 3.0 or higher).
- **Fibre Channel support** – RAID Manager 6.2x supports selected Fibre Channel host adapters.
- **Changed recovery functions** – The Recovery Guru has changed. The function is now tied with the Health Check function and provides recovery information without providing the step-by-step procedures found in earlier versions of the software. In addition, Recovery Guru does not send any commands to the controller or initiate any action automatically. All Recovery operations are performed by the user.
- **Faster LUN creation** – Logical units are created faster. The controller completes formatting in background, while you are using the LUN.
- **Autosynch of controller firmwar**e – If you replace a controller with a firmware level of 3.0 or greater, the firmware on the new controller automatically synchronizes with the firmware on the remaining controller.

Note – The controller autosynch feature requires the setting of two NVSRAM bits for it to work. At the time this manual was published, this feature was not supported by Sun.

- **SMART (Self-Monitoring Analysis and Reporting Technology) drive support** – the software now reports PFA (Predictive Failure Analysis) flags on drives.
- **Parity Check/Repair enhancements** – You can now choose not to repair errors during automatic parity check/repair.

For other new features, see the *Sun StorEdge RAID Manager Installation and Support Guide*.

CHAPTER 3

Common Application Features

This chapter describes the software features common to all applications, including how to display and save RAID Module configuration information.

Navigation

This software requires that you use a mouse (or similar pointing device) for fullest functionality; however, you can also use your keyboard to access the task options. TABLE 3-1 describes how to use the mouse and keyboard.

TABLE 3-1 Using a Mouse and Keyboard

When using a mouse...	• To select an option, place the pointer over the desired option and single-click. • To receive information about a top menu option, click on the option and hold down the left-mouse button. • To receive information about a particular button option, move the pointer over the appropriate button and read the description at the bottom of the screen. • To highlight items, do one of the following: 1. Single-click to highlight a single item. 2. Hold down the <Shift> key and click to highlight a series of items. For example, single-click to highlight the top item in a list, then hold down the <Shift> key and click on the last item in the list to highlight all the items in that list. 3. Hold down the <Control> key and click to highlight items not in a series. For example, single-click to highlight one item in a list, then hold down the <Control> key and click on another item to highlight it as well. Do this for every item you want to highlight.
When using a keyboard...	• To select an option using the keyboard (such as Locate Module), press <Alt> and the key for the underlined letter that appears on the screen. • If you select a task button, the associated screen is launched. For example, <Alt> + <L> brings up the Locate Module screen. • If you select from the top-menu items, a drop-down menu displays the second-level menu options that are available. To select a second-level menu item, press the key for the underlined letter in that option. For example, to select Save Module Profile from the File menu, press <Alt> + <F>, then either press <S>, or use the arrow keys to highlight Save Module Profile and press <Enter>.

Starting and Exiting an Application

Starting an Application

Select the application (FIGURE 3-1) you want to start.

- The first screen you see prompts you to select the RAID Module you want to access (see “Selecting a Module” on page 41).
- If the message log file is currently over its maximum size limit, the storage management software prompts you to rewrite the file (see “Clearing the Message File” on page 110).

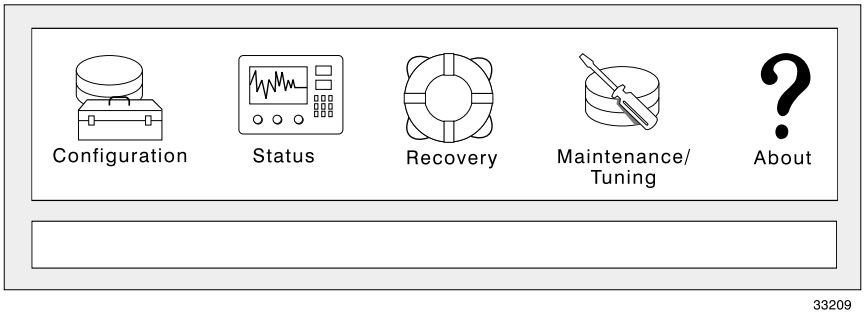

FIGURE 3-1 Application Icons

Exiting an Application

To exit any application, select File → Exit from the top menu. The application exits and you are returned to the Application icons.

Using Online Help

A powerful, hypertext online help system is available with this software. This help has information on features common to all the applications as well as topics that are specific to each application (Configuration, Status, Recovery, and Maintenance/Tuning).

You can access all of the help topics from any application. However, in situations where a new screen is overlaid on top of the main application screen, you cannot access help from within that specific application. However, it is still possible to obtain help by selecting Help from another application.

In many cases, the online help offers more specific information than is given in this *User's Guide*. If you have questions concerning a specific procedure, check the online help before coming back to this manual.

When you select Help from the top menu of any application, you see a screen similar to FIGURE 3-2. TABLE 3-2 details the features common to online help.

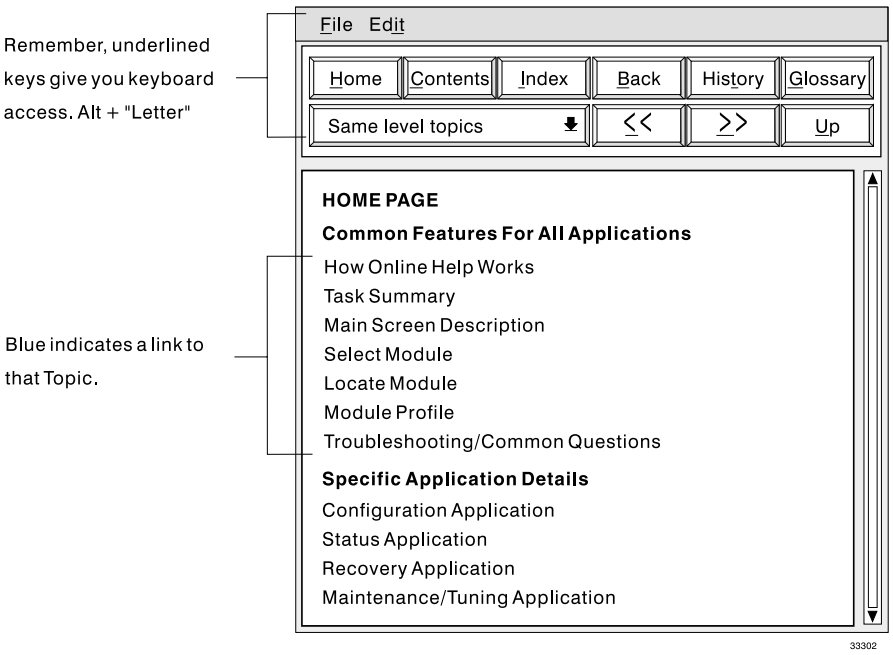

FIGURE 3-2 Main Online Help Screen

TABLE 3-2 Online Help Screen Elements

Selection	Description
File	Allows you to do the following: • Print the displayed topic to a file or printer. • Set up your printer (landscape/portrait, margins, and so on). • Exit Online Help.
Edit	Copies text to the clipboard. From the top menu, select Edit → Copy to Clipboard to copy the topic in the window you are viewing.
Home	Returns you to the Home Page. This screen appears whenever you select Help from the top menu in an application.
Contents	Displays all the help topics organized by hierarchy and appearance on the Home Page. Press a letter to quickly advance through the list of topics. For example, pressing <M> takes you to the first topic that begins with M. You can also use the <Home> and <End> keys on your keyboard to move through this list.
Index	Lists key words or phrases in alphabetical order in the top of the Index window. The bottom of the window displays the topics in which the highlighted index term appears. Press a letter to quickly advance through this alphabetical list. For example, pressing <M> takes you to the first word that begins with M. You can also use the <Home> and <End> keys on your keyboard to move through this list. To view one of these topics, you can either double-click the topic or simply highlight the topic and select <Go To>.
Back	Takes you back (one topic at a time) through the topics you have viewed since selecting Help.
History	Creates a list of all topics you view in the order you have selected them. This feature begins a new list each time you enter Help. To return to one of these topics, either double-click the topic or simply highlight the topic and select <Go To>. You can also use the <Home> and <End> keys on your keyboard to move through this list.

TABLE 3-2 Online Help Screen Elements *(Continued)*

Selection	Description
Glossary	Displays an alphabetical list of defined terms. Press a letter to quickly advance through this alphabetical list. For example, pressing <M> takes you to the first word that begins with M. To view a definition, click and hold the mouse button while pointing to the glossary term. **Note**: The ability to quickly advance through this list as described above is not available. To navigate the list, use the <Home>, <End>, <Page Up>, and <Page Down> keys.
Same level topics	Displays topics of the same level using the less than and greater than buttons to move forward or backward. You can also select Same Level Topics → All Topics from the drop-down menu to make the arrow buttons move you through every topic in help.
<< and >>	• If you are in Same Level Topics, these buttons move you to the previous / next topic within the level you are currently viewing. • If you are in All Topics, these buttons move you to the previous / next topic across all levels.
Up	Moves you to the next higher level of topics.

Selecting a Module

All storage management operations require you to first select the module you want to work on. Once you select the module, you can perform an operation on it, such as configuration or recovery.

Instead of selecting a single module, you can select All RAID Modules if you want to perform an operation on all of the modules. Note, however, that certain operations (such as configuration) cannot be performed on all of the RAID Modules simultaneously.

RAID Modules are identified by either a name assigned by the operating system or a user-specified name. You can name the module through the Module Selection screen.

There are several ways to select a RAID Module:

- If the correct parameter is turned on in your `rmparams` file (it is turned on by default), you are given the chance to select a RAID Module when you first start an application (see FIGURE 3-3). Either highlight a module and select OK or double-click on the module you want.
- Once in an application, you can:
  - Select the RAID Module you want from the RAID Module drop-down list box next to the tool bar.
  - Click on the Module Select button.
  - Select Module → Select from the drop-down menus.
  - Right click on any item in the Module Information area (in the Configuration application only).

In all cases except when selecting a module from the RAID Module drop-down list box, the RAID Module selection screen is displayed. See “RAID Module Selection Screen” on page 42 for details on that screen.

RAID Module Selection Screen

If you select a module through any procedure except the RAID Module drop-down list box, the Module Selection screen is displayed. FIGURE 3-3 gives an example of this screen.

This display shows all the RAID Modules connected to your system. In any application other than Configuration, you can also select All RAID Modules. You can double-click on an individual module to select that module, or you can highlight the module and click on OK.

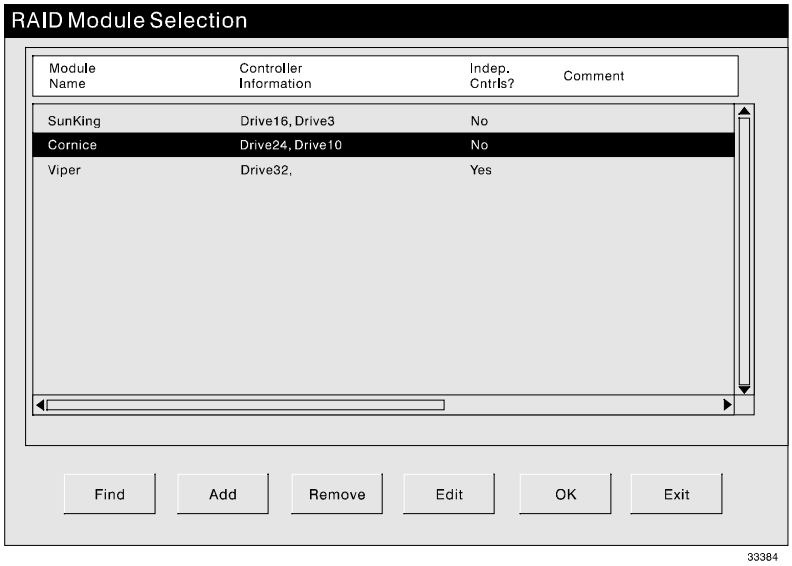

FIGURE 3-3 RAID Module Selection Screen

Note – The display gives the module name, the names of the controllers in the module, indicates whether the controllers have an independent controller configuration (see “Independent Controller Configuration” on page 23), and displays any comments you may have entered with the Edit option.

TABLE 3-3 gives more information on the buttons you can select.

TABLE 3-3 Select Module Screen Elements

Screen Element	Description
Find	Allows you to quickly locate a RAID Module. It is most useful when you have many modules. At the pop-up screen, enter the search term you want to use. Remember that the search item must be contained in one of the fields on this screen.
Add	Applicable for the Networked version only. Allows you to add new modules to your system so that this software can access and monitor them. See the *Sun StorEdge RAID Manager Installation and Support Guide* specific to your operating system for details on adding new modules to your system.
Remove	Allows you to remove RAID Modules from your system. You can only remove one module at a time. This option is grayed out when you select All RAID Modules. **Important:** • Before you use this option, first physically remove the module from the system. Otherwise, the module will be added again when this software detects it. • If you physically remove RAID Modules from your system, but do not use this Remove option, the storage management software will try to contact the controllers in that module. This is especially important in the networked environment because missing modules could cause the software to have long delays or even system hangs while trying to contact the removed module.
Edit	Allows you to assign a name to the RAID Module and to enter comments about it. Use the comments area to provide detailed information about the RAID Module (such as location information, independent controllers, and so on) to help you identify it. On Networked versions, you can also assign names to the RAID Module controllers. You can only edit information for one module at a time. This option is grayed out when you select All RAID Modules.

Locating a Module

The storage management software allows you to physically locate a RAID Module. This is useful in trying to determine the actual unit you are accessing.

Once you locate a module, it is a good idea to label it, so that you will not need to use the software to locate the module again.

If you need to locate individual drive groups/logical units or individual drives, use the List/Locate Drives option in the Configuration application (see "Locating Drive Groups and Drives" on page 45).

Note – For best results, shut down all I/O activity to the module you want to locate, so that you won't be distracted by the normal blinking of the drive activity lights during I/O activity.

To locate a RAID Module, perform the following procedure.

1. **Start the application you want to use.**
2. **Select the RAID Module you want to locate. Do not select All RAID Modules, because this grays out the Locate Module option.**
3. **Click on the Locate Module button or select Module → Locate from the drop-down menus.**
4. **Click Start on the Locate Module screen.**

   The drive activity lights (green lights) on the selected drives in the RAID Module begin to blink.
5. **Locate the module. This is a good time to label it for future reference.**
6. **After you have physically located the module, click Stop and return to the current application.**

   End of Procedure

Locating Drive Groups and Drives

In the Configuration application, you can physically locate drive groups and drives, and identify mirrored pairs in a RAID 1 LUN. Drives are identified by channel number and SCSI ID. See your hardware guides for more information on drive numbering.

Note – Shut down all I/O activity to the module you want to locate so that you won't be distracted by the normal blinking of the drive activity lights during I/O activity.

To locate a drive group or individual drives, perform the following procedure.

1. **Start the Configuration application.**
2. **Select the RAID Module containing the drive group or drives you want to locate.**
3. **Click on the drive group or drives you want to locate in the Module Information area.**
4. **Click on the List/Locate Drives button or select Options → List/Locate Drives.**

   A screen like FIGURE 3-4 **appears.**
5. **Select the drives you want to locate (click Select All if you want to locate all the drives in the group), then select Locate.**
6. **Press Start on the Locate Group screen.**

   The drive activity lights (green lights) on the selected drives begin to blink.
7. **After you have physically located the drives, click Stop and return to the Configuration application.**

   End of Procedure

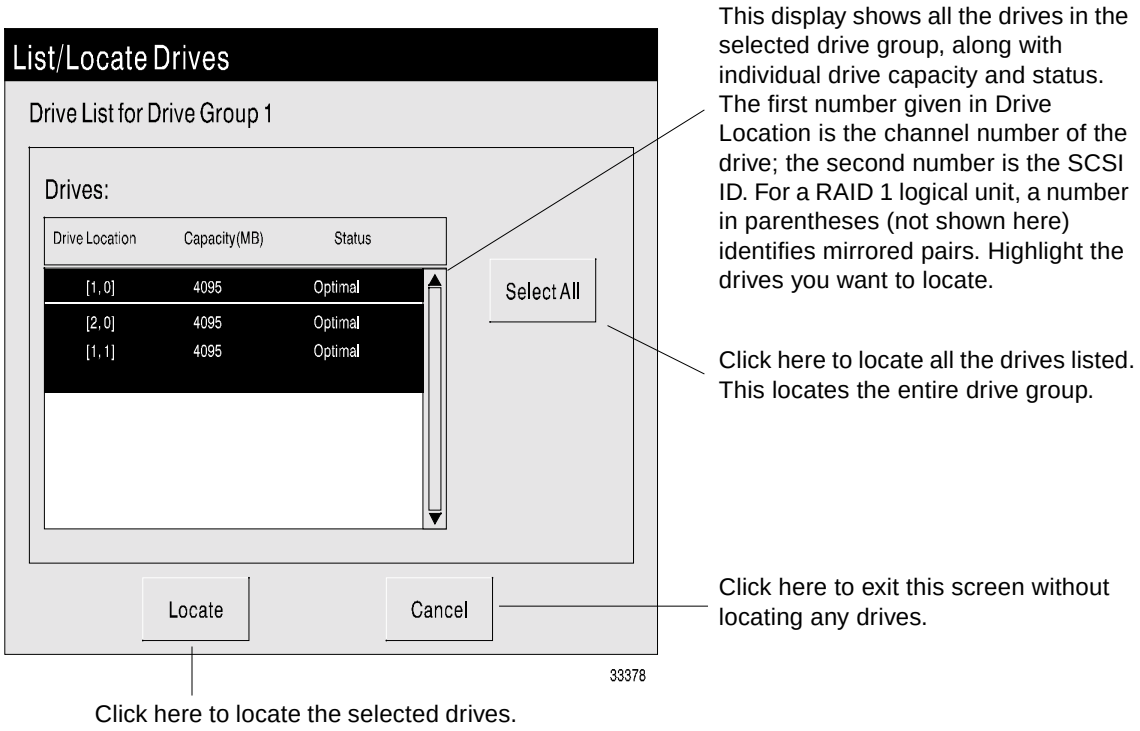

FIGURE 3-4 List/Locate Drives Screen

Viewing a Module Profile

You can view configuration information on a RAID Module using Module Profile. The Module Profile option displays complete information on the controllers, drives, and logical units on the selected RAID Modules.

To display a module profile, perform the following procedure.

1. **Start any application. (Configuration is the recommended application.)**
2. **Select the module you want information on. Do not select All RAID Modules (the Module Profile option is grayed out if you select All RAID Modules).**
3. **Click on the Module Profile button or select Module → Profile from the drop-down menus.**

   The Module Profile screen is displayed (see FIGURE 3-5).
4. **After you have viewed all the profile information you want, click Close on the RAID Module Profile screen to return to the application.**
5. **After exiting Module Profile, you can save the profile information to a file. See “Saving a Module Profile” on page 59.**

End of Procedure

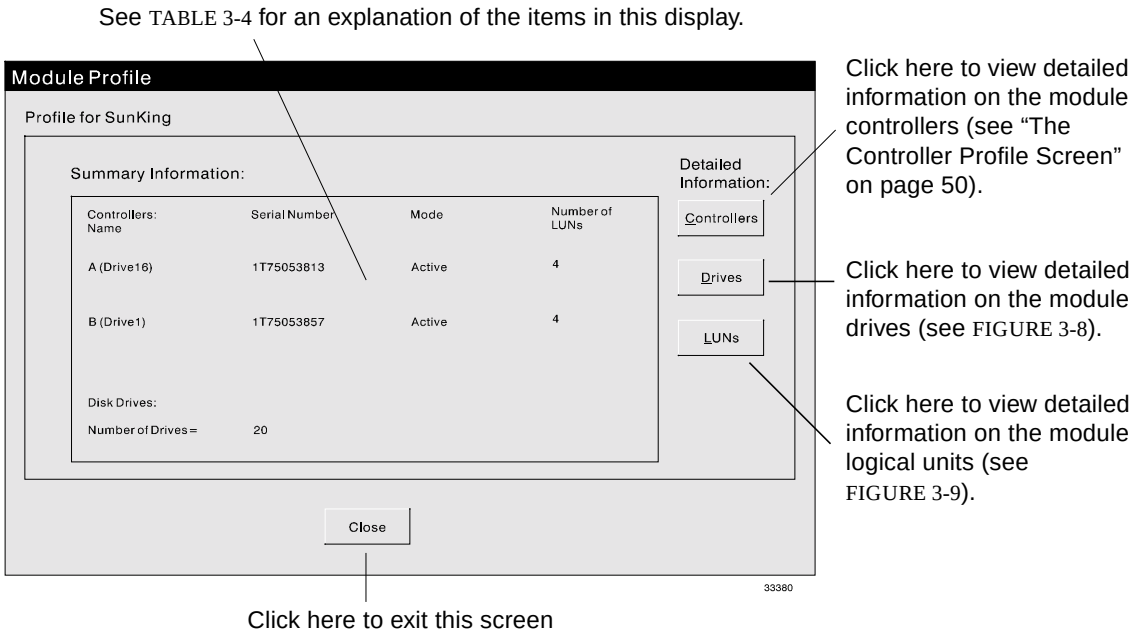

FIGURE 3-5 Module Profile Screen

TABLE 3-4 Module Profile Screen Contents

Screen Element	Explanation
Name	Identifies the controllers in the selected RAID Module as A or B and, where applicable, includes a system device name. The A and B correspond to the hardware controller IDs.
Serial number	Identifies the controller by a number assigned by the manufacturer.
Mode	Identifies the operating state of the controller. Possible modes are Active, Passive, or Offline. You could also see Inaccessible with these statuses if the RAID Module has an independent controller configuration. **Important:** If you do not see Mode information or other information in this screen, there may be a connection problem. Select Recovery Guru and correct any problems indicated.
Number of LUNs	Indicates how many logical units (LUNs) are owned by the particular controller.
Disk drives	Indicates how many drives make up the selected RAID Module.

The Controller Profile Screen

The Controller Profile screen that is displayed depends on whether the host connection is SCSI or Fibre Channel.

- SCSI Controller Profile Screen—page 50
- Fibre Channel Controller Profile Screen—page 52

SCSI Controller Profile Screen

FIGURE 3-6 shows the Controller Profile screen for a SCSI connection. See TABLE 3-5 for an explanation of the items in this display.

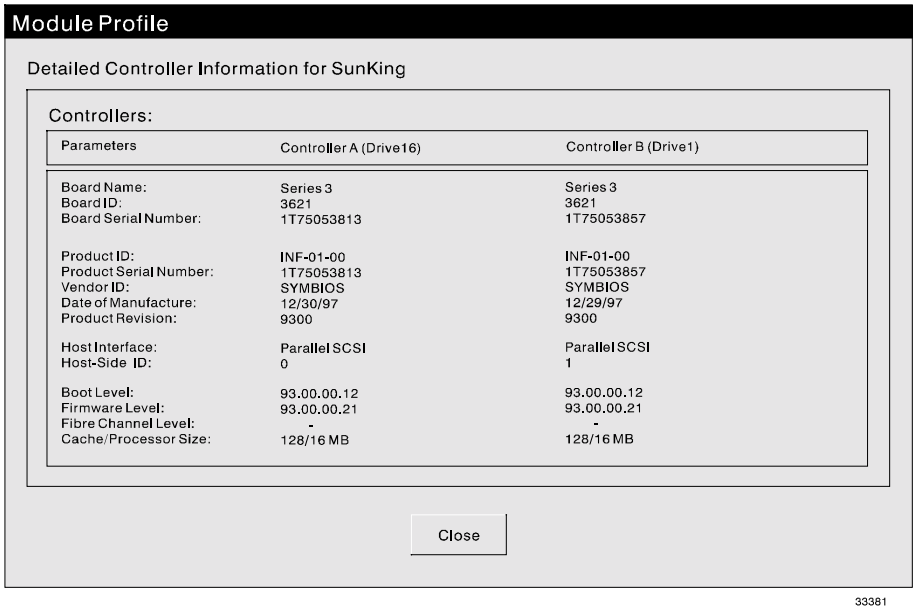

FIGURE 3-6 Controller Profile Screen (SCSI)

TABLE 3-5 Controller Profile Screen Contents (SCSI)

Screen Element	Explanation
Board Names	Controller type designation.
Board ID	Controller model number (3601, 3201, etc.).
Board Serial Number	Unique identification for the controller assigned by the manufacturer.
Product ID	Controller manufacturer's product code.
Product Serial Number	Usually the same as Board Serial Number.
Vendor ID	Controller manufacturer's name.
Date of Manufacture	Date controller was assembled.
Product Revision	A manufacturer’s code giving the revision number of the controller.
Host Interface	The type of connection used by the host (Parallel SCSI).
Host-Side ID	A number used to identify the controller connection between the controller and the host (host-side bus). In this display, it is a SCSI ID. This number is set on the back of the controller using the Host ID switch. (Not applicable for Networked versions.)
Boot Level	Number indicating the release version of controller bootware.
Firmware Level	Number indicating the release version of controller firmware (also referred to as *appware*).
Fibre Channel Level	Number indicating the release version of Fibre Channel controller firmware.
Cache / Processor Size	Amount (in megabytes — MB) of total available cache and processor memories on the controller.

Fibre Channel Controller Profile Screen

FIGURE 3-7 shows the Controller Profile screen for a Fibre Channel connection. See TABLE 3-6 for an explanation of the items in this display.

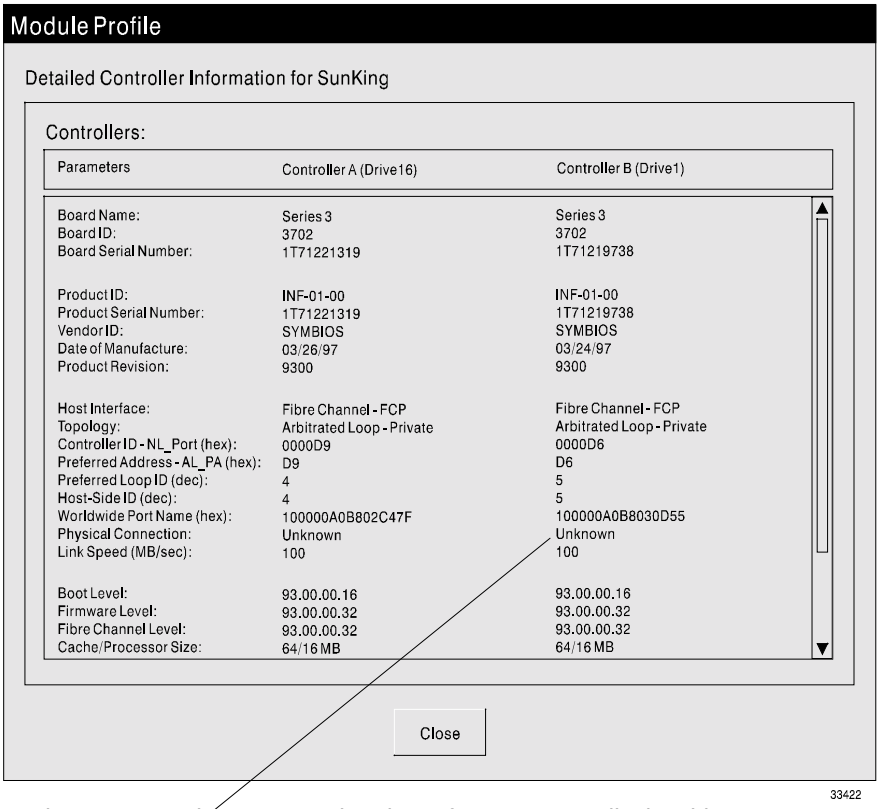

"Unknown" or "_" here means that the software cannot display this information. This is usually a limitation on the Fibre Channel firmware.

FIGURE 3-7 Controller Profile Screen (Fibre Channel)

Note – The information about firmware version levels and cache / processor size displayed for your controller may be different from the information shown in FIGURE 3-7.

TABLE 3-6 Controller Profile Screen Contents (Fibre Channel)

Screen Element	Explanation
Board Name	Controller type designation.
Board ID	Controller model number (3702, 4766, and so on).
Board Serial Number	Unique identification for the controller assigned by the manufacturer.
Product ID	Controller manufacturer's product code.
Product Serial Number	Usually the same as Board Serial Number.
Vendor ID	Controller manufacturer's name.
Date of Manufacture	Date controller was assembled.
Product Revision	A manufacturer's code giving the revision number of the controller.
Host Interface	The type of connection used by the host (Fibre Channel).
Topology	The physical or logical layout of nodes on the Fibre Channel network. There are two topologies, Arbitrated Loop (Public or Private) and Point-to-Point.
Controller ID - NL_Port (hex)	The Fibre Channel Controller's ID in a hexadecimal value. • For Arbitrated Loop, this is a 1-byte Arbitrated Loop Physical Address. • For Fabric and Point-to-Point, this is a 3-byte hexadecimal value.
Preferred Address - AL_PA (hex)	The hard address (AL_PA format) that Fibre Channel controllers attempt to acquire. Displayed only for controllers using the Arbitrated Loop topology. If the hard address is not available, the controller automatically obtains the first address available after the other ports on the loop have obtained their hard addresses.
Preferred Loop ID (dec)	The Fibre Channel Loop ID in NVSRAM used to determine the Preferred AL_PA (hard address). Displayed only for controllers attached using the Arbitrated Loop topology. This ID is *not* the AL_PA, but an index into a table of valid AL_PA values (see Annex K of the FC-AL standard). Loop IDs are assigned an arbitration priority with (host-side IDs) 0 being the lowest and 126 being the highest.
Host-Side ID	A number used to identify the controller connection between the controller and the host (host-side bus). This number is set on the back of the controller using the Host ID switch.
Worldwide Port Name (hex)	An 8-byte hexadecimal value used to uniquely identify Fibre Channel controllers. It consists of fields identifying the naming convention, the node, and its ports.
Physical Connection	The type of Fibre Channel connection, either copper or optical.

TABLE 3-6 Controller Profile Screen Contents (Fibre Channel) *(Continued)*

Screen Element	Explanation
Link Speed (MB/sec)	The speed in megabytes per second (MB/sec) of the Fibre Channel connection.
Boot Level	Number indicating the release version of controller bootware.
Firmware Level	Number indicating the release version of controller firmware (also referred to as appware).
Fibre Channel Level	Number indicating the release version of Fibre Channel controller firmware.
Cache/Processor Size	Amount in megabytes (MB) of total available cache and processor memories on the controller.

"Unknown" or "_" in a column means that the software cannot return the indicated value. This is not necessarily an error; it is usually a limitation on the Fibre Channel firmware.

The Drive Profile Screen

FIGURE 3-8 shows the Drive Profile screen. See TABLE 3-7 for an explanation of the items in this display.

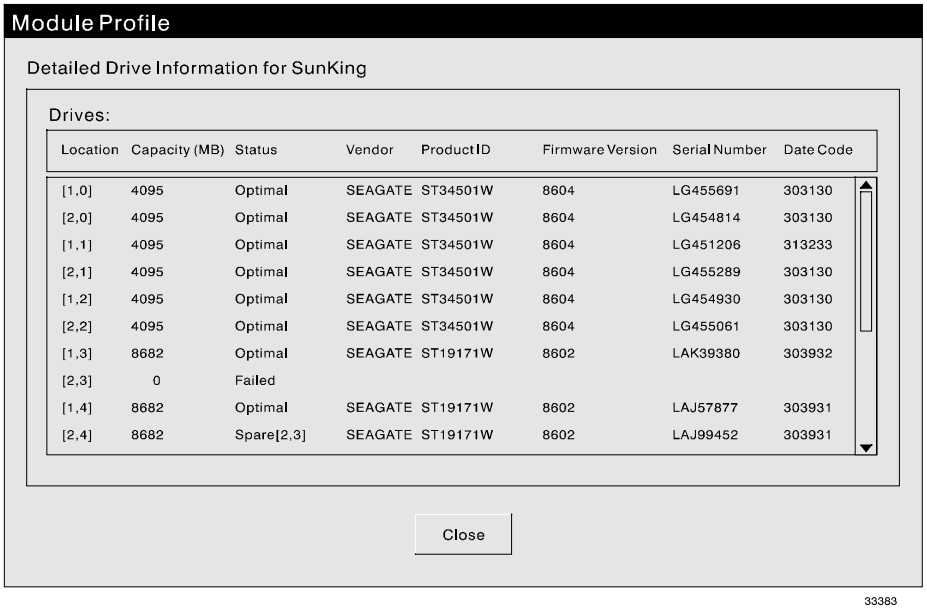

FIGURE 3-8 Drive Profile Screen

TABLE 3-7 Drive Profile Screen Contents

Screen Element	Explanation
Location	Indicates the location of the drive in the selected RAID Module. This information is displayed as [*x*,*y*] on screen where *x* is the channel number and *y* is the drive SCSI ID. For example, [2,1] is to the drive at SCSI Channel 2 and SCSI ID 1.
Capacity (MB)	Amount of storage space on the drive in megabytes (MB).
Status	Operating condition of the drive. If a status other than Optimal appears here, run Recovery Guru to determine what the problem is, and, if necessary, what action to take.
Vendor	Drive manufacturer's name.
Product ID	Drive manufacturer's product code.
Firmware Version	Number indicating the release of drive firmware.
Serial Number	Drive manufacturer's serial number.
Date Code	Date of manufacture.

The LUN Profile Screen

FIGURE 3-9 shows the LUN Profile screen. See TABLE 3-8 for an explanation of the items in this display.

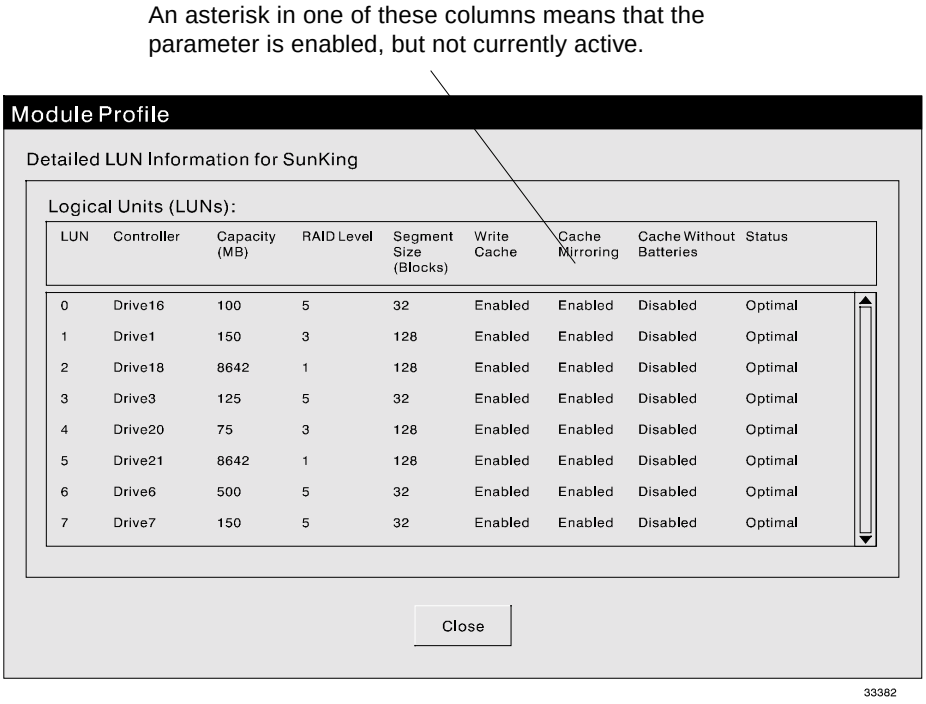

FIGURE 3-9 LUN Profile Screen

TABLE 3-8 LUN Profile Screen Contents

Screen Element	Explanation
LUN	Identifies the number of the logical unit (LUN).
Controller	Identifies the controller that owns the logical unit.
Capacity (MB)	Shows the amount of storage space in megabytes (MB).
RAID Level	Indicates the way the controller reads and writes data and parity on the drives. RAID Levels are 0, 1, 3, and 5.
Segment Size (Blocks)	Indicates the amount of data (in blocks) that the controller writes on one drive in a logical unit before writing data on the next drive.
Write Cache	Indicates whether the write caching parameter has been enabled for a particular LUN.
Cache Mirroring	Indicates whether the cache mirroring parameter has been enabled for a particular LUN.
Cache Without Batteries	Indicates whether the cache without batteries parameter has been enabled for a particular LUN.
Status	Indicates the operating condition of the logical unit. If a status other than Optimal appears here, run Recovery Guru to determine what the problem is, and, if necessary, what action to take.

Note – You might see an asterisk next to the caching parameters column. This indicates that the parameter is enabled, but is currently *not* active. The controller has disabled the parameter for some reason (such as low batteries). If you see this condition, use Recovery Guru to determine the correct action to take, if any.

Saving a Module Profile

You can save module profile information in a text file for future reference. The saved information cannot be used to restore module settings, but can be used as a reference if you need to recreate the module using the standard configuration procedures.

To save a module profile, perform the following procedure.

1. **Start any storage management application.**
2. **Select the RAID Module you want to save profile information for. Do not select All RAID Modules (this option is grayed out if you select All RAID Modules).**
3. **Select File → Save Module Profile from the drop-down menus.**

   A screen like FIGURE 3-10 is displayed.

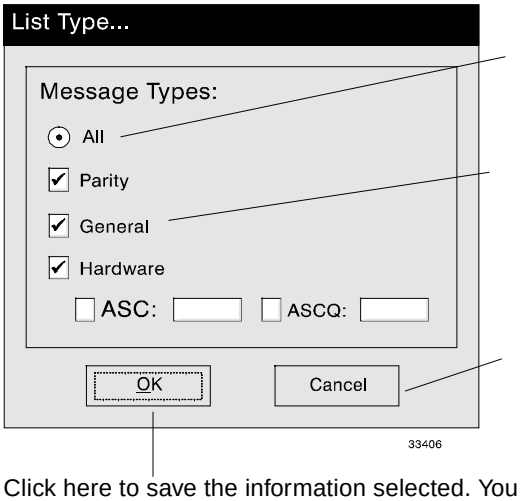

FIGURE 3-10 Save Module Profile Screen

4. **Click on the information you want to save (or click on All to save all information), then click OK.**
5. **You must then assign a name to the file. Enter the name and click OK.**

   End of Procedure

CHAPTER 4

Configuring RAID Modules

This chapter describes how to create and configure RAID Modules and assign hot spares.

Configuration Overview

Tip – Before you attempt to create and configure a RAID Module, make sure you understand the concepts of RAID Module, Drive Group, Logical Unit, and RAID Level. See “Common Definitions” on page 8 for information on these important concepts, which are needed to perform the tasks in this section.

Before you can use the RAID Modules on your system, you must create logical units (LUNs) on the RAID Module. In general, you need to do the following to create and use the RAID Modules:

1. **Create and configure the logical units using the procedures given in this chapter.**
2. **Perform whatever procedures on the logical unit that your operating system requires to bring a new drive on line. For example, you may need to restart the system, create partitions, file systems, and volumes, create mount points, and so on.**

   Refer to the following sections for more information on LUN creation, modification, and formatting:

   - Types of Logical Unit (LUN) Creation—page 63
   - Types of Logical Unit (LUN)/Drive Group Modification—page 63
   - Logical Unit (LUN) Formatting—page 64

Types of Logical Unit (LUN) Creation

There are two types of LUN creation:

- Creating a new LUN from the unassigned drives in the RAID Module. This also creates a new drive group. When you create a LUN from unassigned drives, you assign the RAID level and number of drives. See “Creating a Logical Unit from Unassigned Drives” on page 69.
- Creating a new LUN from remaining capacity in a drive group. When you create a LUN in an existing drive group, you cannot assign a new RAID level or new drives because all of the LUNs in a drive group must use the same RAID level and drives. See “Creating a Logical Unit from Remaining Capacity” on page 72.

Note – Creating LUNS may cause drive groups to be renumbered. See “Drive Group Renumbering” on page 12.

Types of Logical Unit (LUN)/Drive Group Modification

After a LUN is created, you can modify it in the following ways:

- You can change the segment size of the logical unit. This change applies to all LUNs in the drive group. See “Modifying Segment Size” on page 89.
- You can change the caching parameters. See “Changing Cache Parameters” on page 167.
- You can change the reconstruction rate. See “Changing the Reconstruction Rate” on page 158. Note that this change also affects the speed of LUN formatting.

After a Drive Group is created, you can modify it in the following ways:

- You can change the RAID level of the logical units in the drive group. This change applies to all logical units in the drive group. See “Changing RAID Level” on page 85.
- You can increase Drive Group capacity by adding new drives to the drive group. See “Adding Drives” on page 87.

Note – This option increases the size of the drive group. It does not increase the size of the logical units in the drive group.

- You can change the controller assignment of the drive group/LUN. See “Balancing Drive Groups/LUNs” on page 160.

Logical Unit (LUN) Formatting

Whenever you create a logical unit (LUN) of more than 20 MB, the controller formats (writes zeros to) only the first and last 10 MB of the LUN. It then reports the format operation as completed. The LUN status becomes optimal. At this point, you can begin using the logical unit. You may write data to it, create file systems on it, and so on (but see “Restrictions” on page 66).

After the LUN is created, the controller may complete the format operation in the background, depending on the RAID level of the logical unit:

- If the LUN is a RAID 0 logical unit, data is written to the unformatted portion of the LUN normally. The controller performs no further actions to format the logical unit.
- If the LUN is a RAID 1 logical unit, the controller formats the remaining portion of the logical unit in background. Stripe by stripe, the controller rewrites the contents of each mirrored pair so that the contents match. If the controller needs to write user data in the unformatted portion of the LUN, it writes the data normally (so that the contents of each mirrored pair is the same).
- If the LUN is a RAID 3 or 5 logical unit, the controller formats the remaining portion of the logical unit in background. One at a time, the controller rewrites the contents of each stripe so that the parity is correct in each stripe (see Figure 4-1). If the controller needs to write user data in the unformatted portion of the LUN, it writes the data normally (so that the parity in each stripe is correct).

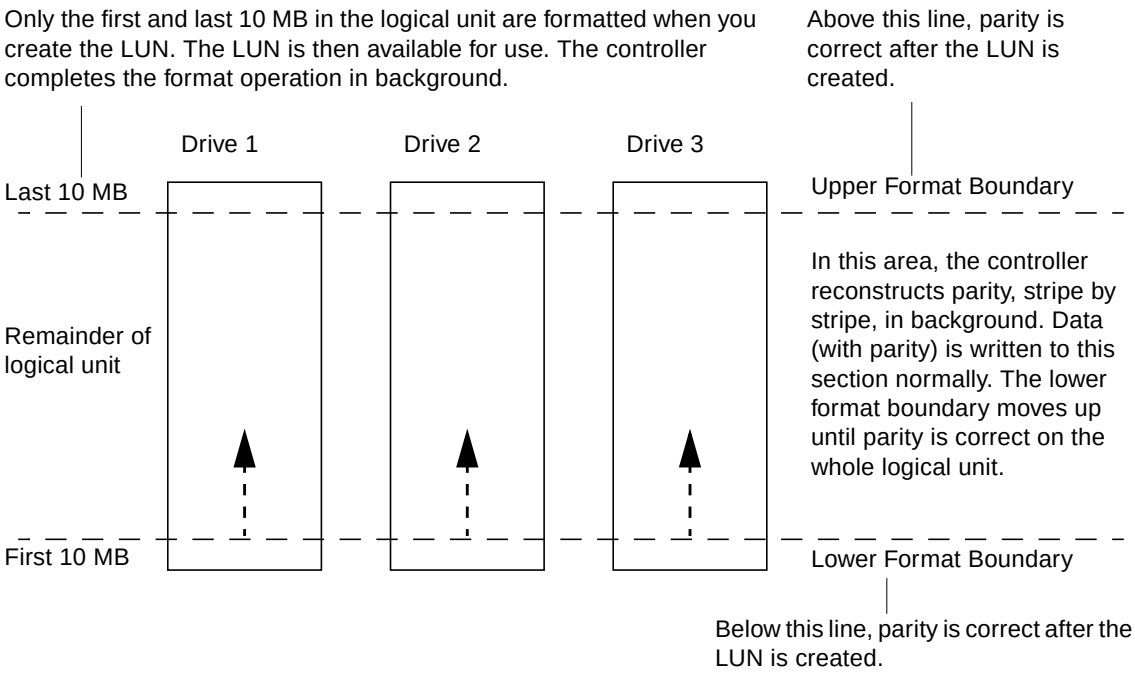

FIGURE 4-1 LUN Creation Example (RAID 5)

While the controller is formatting in background, there is a “boundary” in the LUN. Outside the boundary, the mirroring or parity in the logical unit is correct. Within the boundary, the parity or mirroring is *not* correct. When you perform a parity check on the LUN, only that portion of the LUN that is currently formatted (the section below the format boundary) is checked.

Background formatting is a form of reconstruction, and is affected by the current setting of the Reconstruction Rate parameter (see “Changing the Reconstruction Rate” on page 158). Completing the background format can take a long time (up to 8 hours, depending on the size of the LUN and on how much I/O activity is occurring).

A LUN being formatted in background has an Optimal status. There is no indication in the storage management software that the background format is occurring, aside from the restrictions given below. While background formatting is occurring, however, you can use the LUN normally, except as specified below.

The system may be slower during background formatting (depending on your I/O load and the Reconstruction Rate parameter settings). This is because a reconstruction process is taking place.

If you need to determine whether a LUN is still formatting in the background, try performing a modification option on the affected drive group/LUN (such as changing the segment size as described in "Modifying Segment Size" on page 89). If the LUN is still formatting, you will get a message notifying you that the LUN is still initializing. If the LUN is no longer formatting, you will get the normal Modify screen. Then, you can exit from this screen without modifying the drive group/LUN.

Restrictions

While the LUN is formatting in background, you can perform any operation on it *except*:

- Changing the segment size
- Changing the RAID level
- Adding a drive to the drive group
- Defragmenting the drive group

If the logical unit becomes degraded while it is being formatted in background, the formatting operation stops (in a RAID 1 LUN, the formatting stops only if each mirrored pair has a failed drive — note that in this case the logical unit is still degraded). When you replace the failed drive, it is reconstructed normally (that is, the entire logical unit will be reconstructed before the logical unit becomes optimal).

If the logical unit fails while it is being formatted in background, the format operation is terminated (the logical unit must be reformatted as part of recovery).

The Configuration Application Main Screen

FIGURE 4-2 shows the Configuration application main screen.

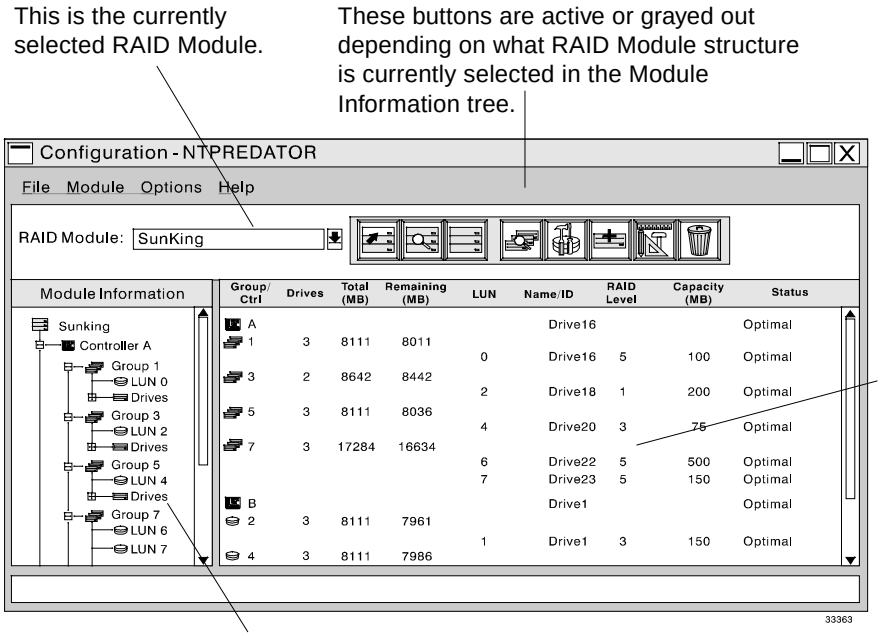

FIGURE 4-2 Configuration Application Main Screen

- Click on the “+” or “-” signs to expand or collapse the Module Information area directory tree.
- Clicking on an element in the Module Information area changes the information displayed in the Detailed Information area on the right of the screen. For example, if you click on a drive group, only information on that drive group is displayed on the right. If you click on the unassigned drive group, information on the drives in the unassigned group (drive ID and capacity) is displayed on the right.

- You cannot click on items in the Detailed Information area.

TABLE 4-1 Configuration Application Main Screen Elements

Column Heading	Explanation
Group / Ctl	The controller (A or B) or drive group.
Drives	The number of drives in the drive group.
Total (MB)	The total capacity of the drive group, after accounting for RAID level.
Remaining (MB)	The total capacity remaining in the drive group.
LUN	The logical units (LUNs) in the drive group.
Name / ID	The name or ID of the controller, logical unit, or drive, depending on the display. • Controller and logical unit names are usually assigned by the operating system, although in the networked version of the storage management software you can assign names to them with the Edit function of the Module Selection screen (see “RAID Module Selection Screen” on page 42). • Drive IDs are the channel number and SCSI ID of the drives. See “Drive Names” on page 13 for more information.
RAID Level	The RAID level of the logical unit.
Capacity (MB)	The capacity of the logical unit in megabytes (MB).
Status	The current status of the logical unit. If a status other than Optimal appears here, run Recovery Guru to determine what the problem is, and, if necessary, what action to take.

Creating Logical Units

You can create a new logical unit (LUN) in two ways:

- By creating a new LUN from unassigned drives (see “Creating a Logical Unit from Unassigned Drives” on page 69).
- By creating a new LUN from the remaining capacity in a drive group (see “Creating a Logical Unit from Remaining Capacity” on page 72).

Creating a Logical Unit from Unassigned Drives

To create a new logical unit (LUN) from unassigned drives, you must have drives remaining in the Unassigned drive group. If there is no Unassigned drive group, this option is grayed out. Creating a new LUN from unassigned drives creates a new drive group.

To create a LUN from unassigned drives, perform the following procedure.

1. **Start the Configuration application.**
2. **Select the RAID Module on which you want to create the new LUNs.**
3. **Select the RAID Module or the Unassigned Drive group and click on the Create LUN button.**

   A screen like FIGURE 4-3 is displayed.

4. **Set the RAID level you want to assign to the new LUN.**

   The RAID level might change the Number of Drives setting, or may be limited by the number of drives available. For example, you need at least three unassigned drives to create a RAID 3 or 5 logical unit, and you must have an even number of drives for a RAID 1.

5. **Set the number of drives you want to include in the new LUN.**

   The drives are automatically assigned unless you specify the drives by selecting Options → Drive Selection in Step 7. Changing the number of drives here may change the RAID level setting. For example, if you select RAID 1 and then specify an odd number of drives, the RAID level will change.

6. **Set the number of LUNs you want to create on the new drive group.**

   This may be limited by the total number of LUNs you can create on your RAID Module. Each LUN created gets an equal amount of the available capacity. For example, if you specify two LUNs, each LUN gets half the available capacity. You can change the capacity by selecting Options → LUN Capacity in Step 7.

7. **If you want to view or change the default settings for drives, capacity, segment size, caching, or LUN assignment, click on Options. See “Using the Options Screens” on page 74 for a description of those screens**

Tip – If your RAID Module contains drives of different sizes, you should use the Options → Drive Selection to make sure that all the drives assigned to the new LUN are of the same size. Otherwise, LUN capacity will be based on the smallest drive size, and capacity will be wasted.

Also, the new drive group is automatically assigned to a controller if the RAID Module has redundant controllers. The first drive group you create in a configuration session is assigned to controller A, the second to controller B, the third to controller A again and so on. If you repeatedly create one drive group per session, therefore, all the created drive groups are assigned to controller A. You can change controller assignment by selecting Options → LUN Assignment in this step.

---

8. **After you have set all of the LUN parameters you want, start LUN creation by clicking Create in the Create LUN screen.**

Tip – If you change any parameters in the Create LUN screen after making changes in the Options screens, any Options changes are lost.

The new LUNs are created. Wait for the LUN status to change from Formatting to Optimal before continuing. Note that if this is a RAID 1, 3, or 5 LUN, formatting may continue in background (see “Logical Unit (LUN) Formatting” on page 64).

Note – If you are creating the *first* LUN on the module, wait for the create/format process to finish before creating any LUNs on other drive groups. In addition, Module Profile information may not be displayed correctly until the first LUN is created.

---

9. **After the LUNs are created, exit the application and do whatever is necessary on your operating system to make the LUNs available for use (using standard OS utilities). This may involve restarting the system, creating partitions, file systems, and volumes, and so on. Check your operating system documentation and the *Sun StorEdge RAID Manager Installation and Support Guide* for more information.**

End of Procedure

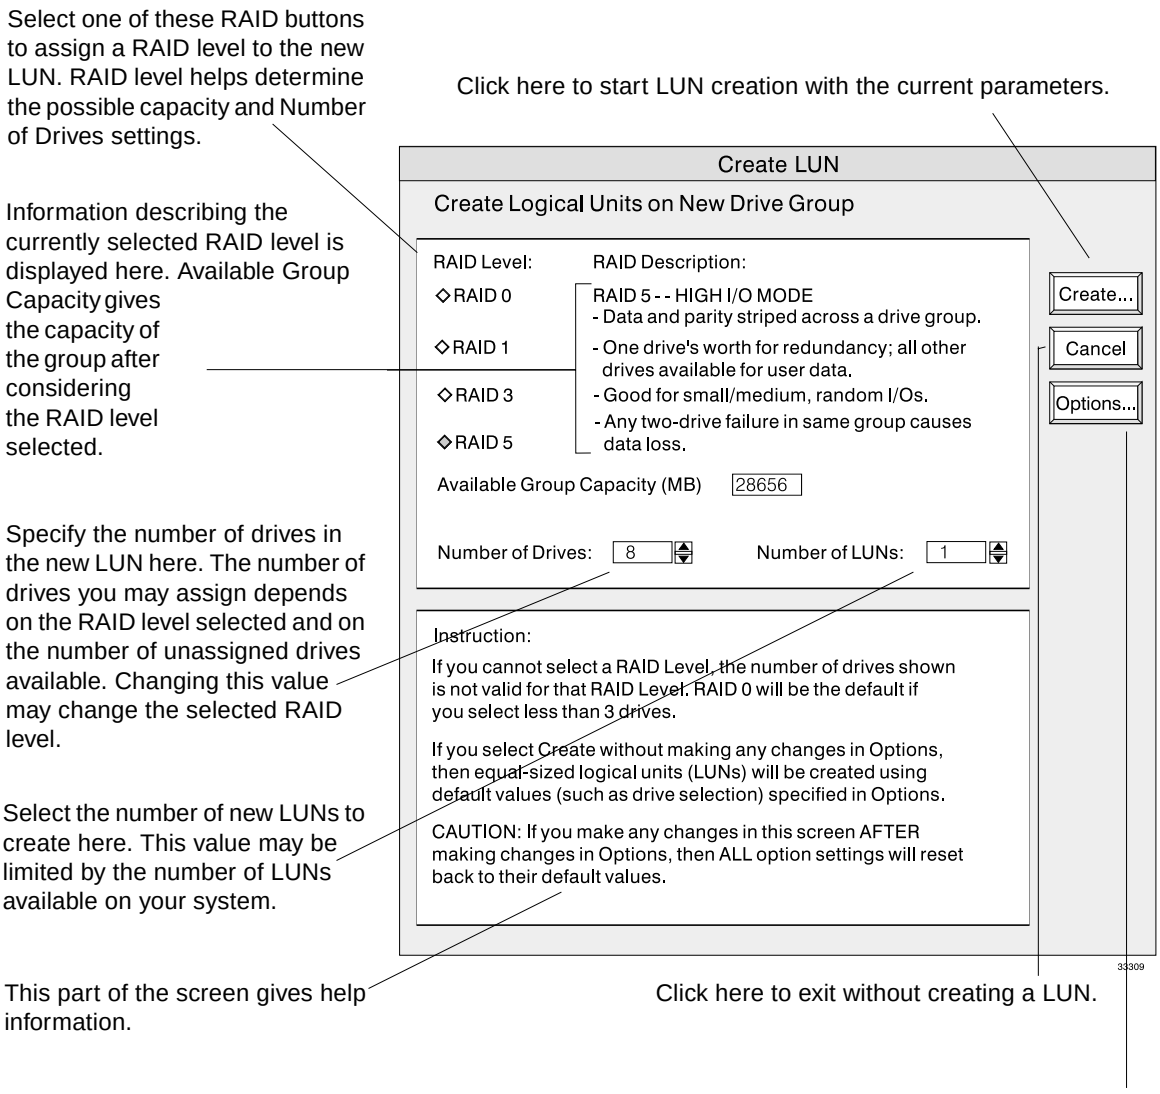

FIGURE 4-3 Create LUN Main Screen (Unassigned Drives)

Creating a Logical Unit from Remaining Capacity

To create a new logical unit (LUN) from remaining capacity in an existing drive group, you must have drive groups with unused capacity. The new LUN will have the same RAID level, drives, and controller assignment as the current drive group.

To create a LUN from remaining capacity in a drive group, perform the following procedure.

1. **Start the Configuration application.**
2. **Select the RAID Module on which you want to create the new LUNs.**
3. **Select the Drive Group with remaining capacity in the Module Information area and click on the Create LUN button.**

   A screen like Figure 4-4 on page 73 is displayed.

4. **Set the number of LUNs you want to create from the unassigned capacity.**

   This may be limited by the total number of LUNs you can create on the RAID Module. By default, each LUN created receives an equal amount of the displayed capacity.

5. **If you want to change the default settings for capacity, segment size, or caching, click on Options (see** "Using the Options Screens" on page 74**).**
6. **After you have set all the LUN parameters you want, start LUN creation by clicking Create in the Create LUN screen.**

Note – If you change any parameters in the Create LUN screen after making changes in the Options screens, any Options changes are lost.

The new LUNs are created. Wait for the LUN status to change from Formatting to Optimal before continuing. Note that if this is a RAID 1, 3, or 5 LUN, formatting may continue in background (see "Logical Unit (LUN) Formatting" on page 64).

7. **After the LUNs are created, exit the application and do whatever is necessary on your operating system to make the LUNs available for use (using standard OS utilities). This may involve restarting the system, creating partitions, file systems, and volumes, and so on. Check your operating system documentation and the *Sun StorEdge RAID Manager Installation and Support Guide* for more information.**

   End of Procedure

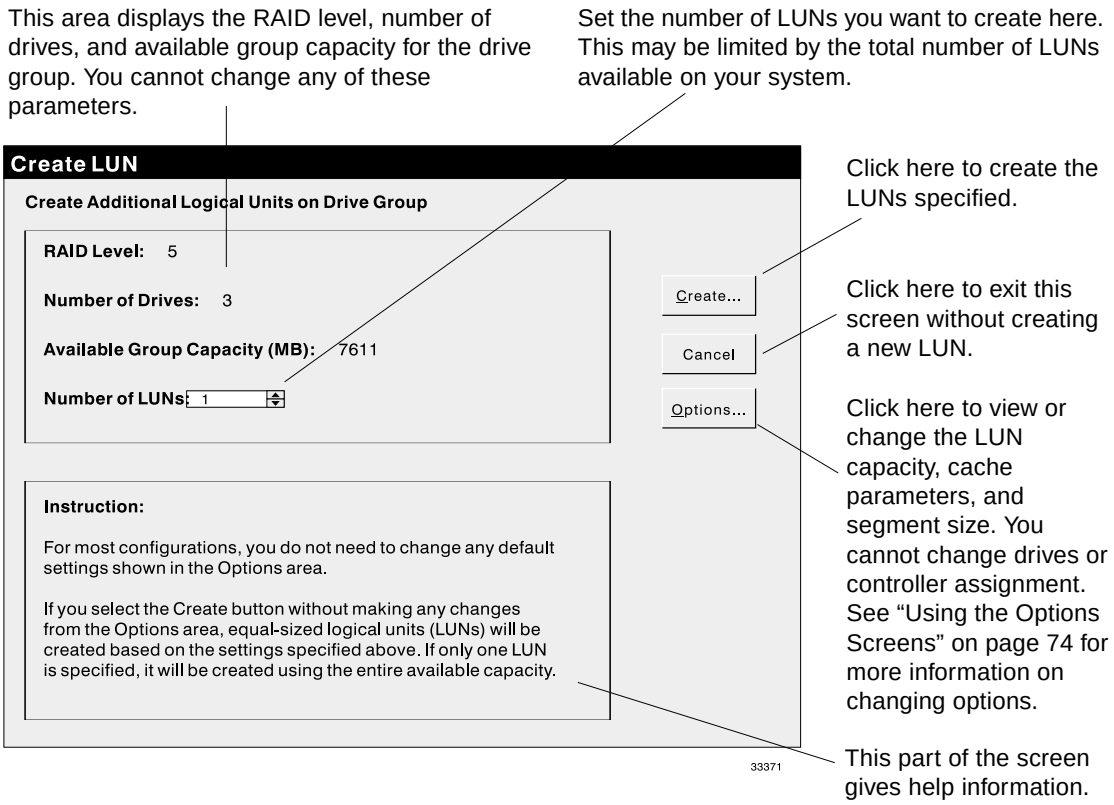

FIGURE 4-4 Create LUN Main Screen (Existing Drive Group)

Using the Options Screens

The Options screens allow you to view or change:

- The capacity of the LUNs to be created (see FIGURE 4-5).
- The drives in the LUN (see FIGURE 4-6).
- The caching parameters (see FIGURE 4-7).
- The segment size (see FIGURE 4-8).
- The controller that “owns” the new LUN. This option is allowed only if your RAID Module has active/active redundant controllers (see FIGURE 4-9).

Note – When you move from screen to screen, the Configuration application automatically saves the values displayed in the current screen.

LUN Capacity Option

Use this option to set the capacity of the new LUNs. By default, the new LUNs are assigned an equal proportion of the total capacity of the drive group. For example, if you create two logical units, each is assigned a half of the total capacity.

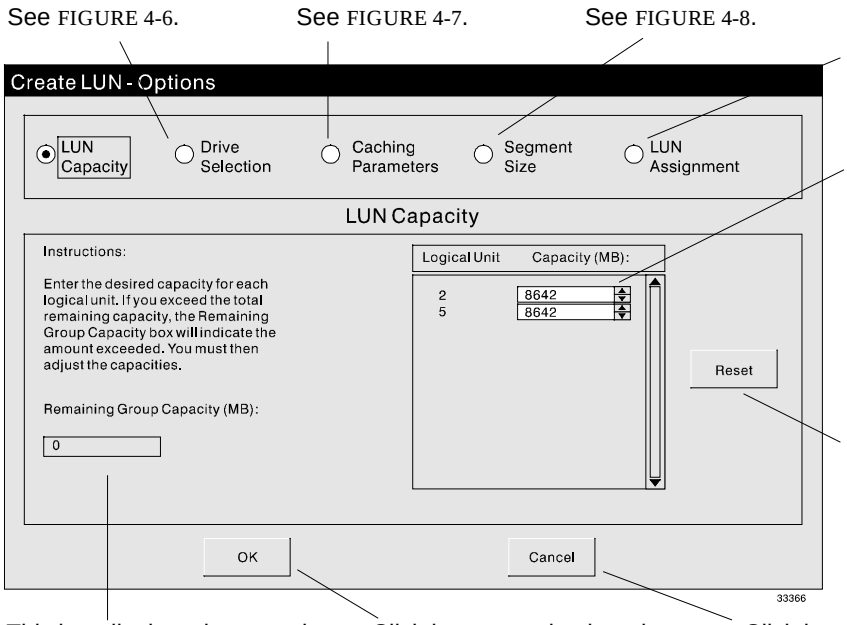

FIGURE 4-5 LUN Capacity Screen

Note – Normally, you will want to use the maximum capacity setting.

Drive Selection Option

Use this option to specify the drives assigned to the new drive group/LUN. In particular, you should make sure that all the drives in the new drive group/LUN have the same capacity. You should also attempt to make sure that all drives are on different channels (the first number in the drive ID), although this is not always possible. Assigning drives on different channels prevents a channel failure from disabling multiple drives.

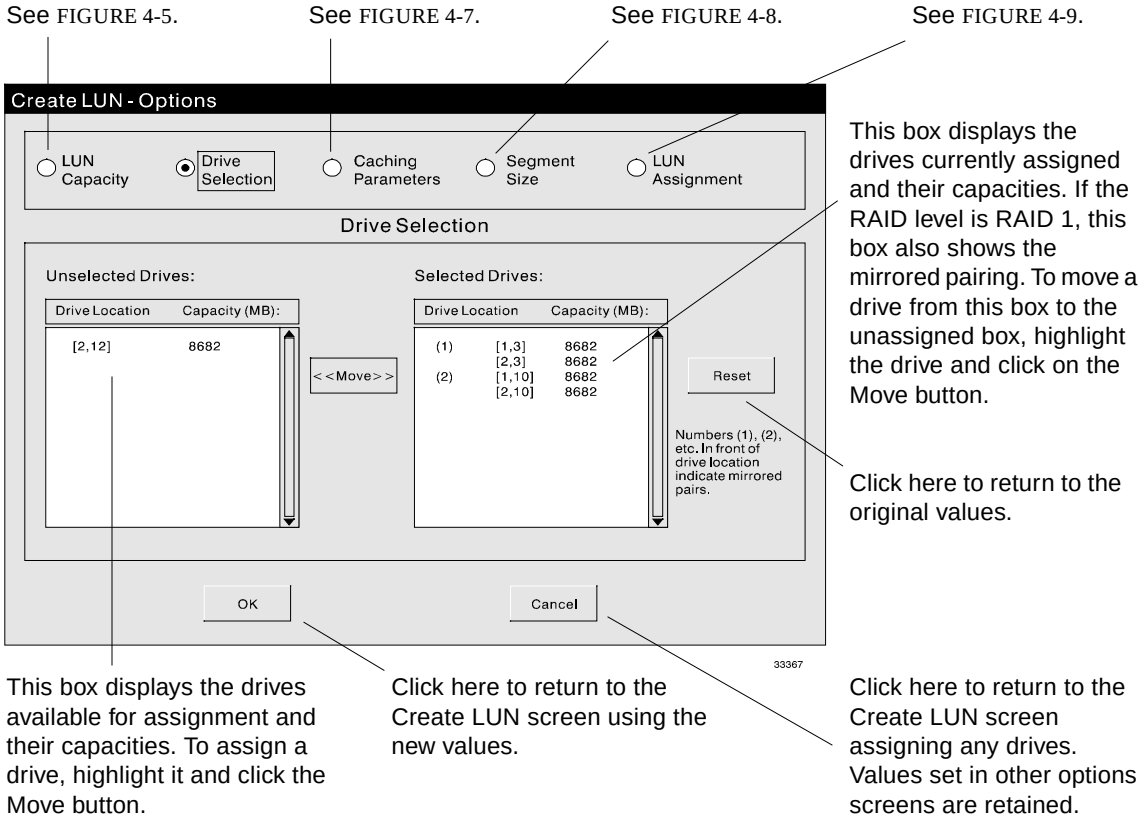

FIGURE 4-6 Drive Selection Screen

Note – You cannot exit from this screen if you have specified a different number of drives than you entered in the Create LUN Main screen (see FIGURE 4-3).

Cache Parameters Option

Use this option to change the cache parameters assigned to the new LUNs. Parameters are assigned on a LUN basis. See “Cache Memory” on page 9 and “Changing Cache Parameters” on page 167 for more information on the cache parameters.

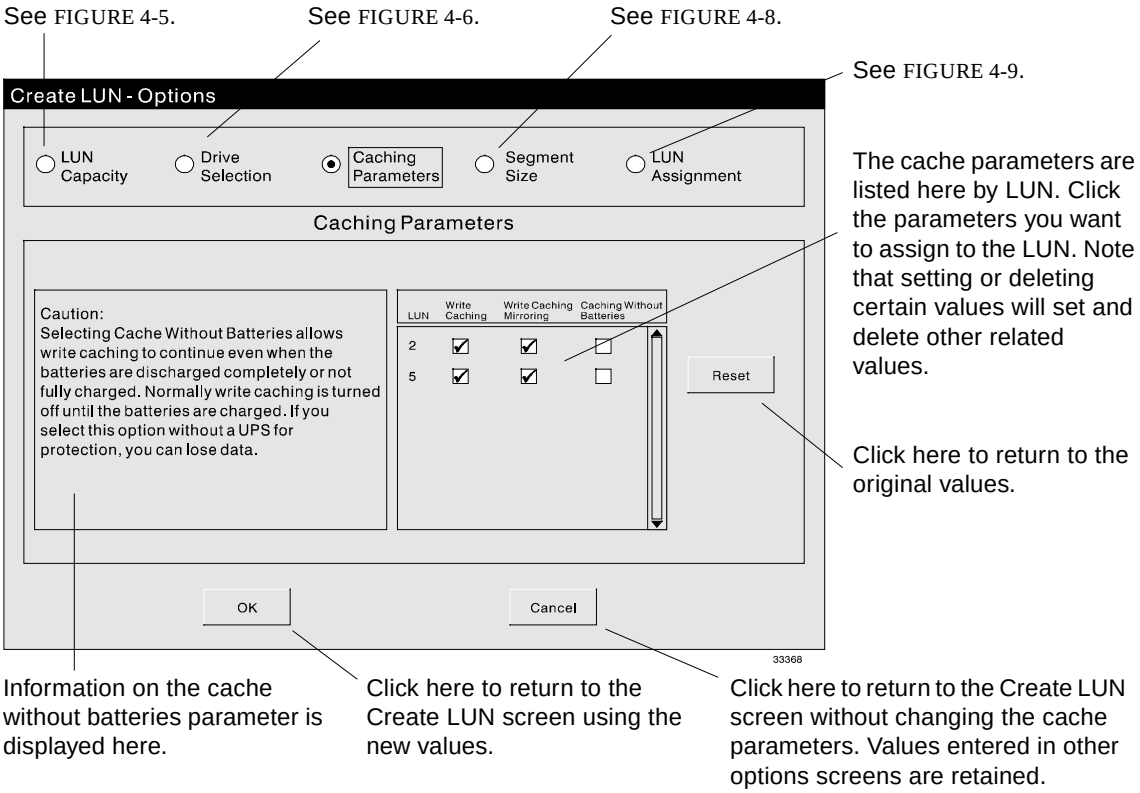

FIGURE 4-7 Cache Parameters Screen

- Write Cache Mirroring is only effective for RAID Modules with redundant controller pairs (active / active or active / passive) that have the same size cache.
- Cache parameters have the following interdependencies:
  - If you select Write Caching, Write Cache Mirroring is also enabled (if possible).
  - If you select Write Cache Mirroring or Cache Without Batteries, Write Caching is also enabled.
  - If you deselect Write Caching, Write Cache Mirroring and Cache Without Batteries are also deselected (if currently enabled).

Segment Size Option

Use this option to change the segment size used by the new LUNs.

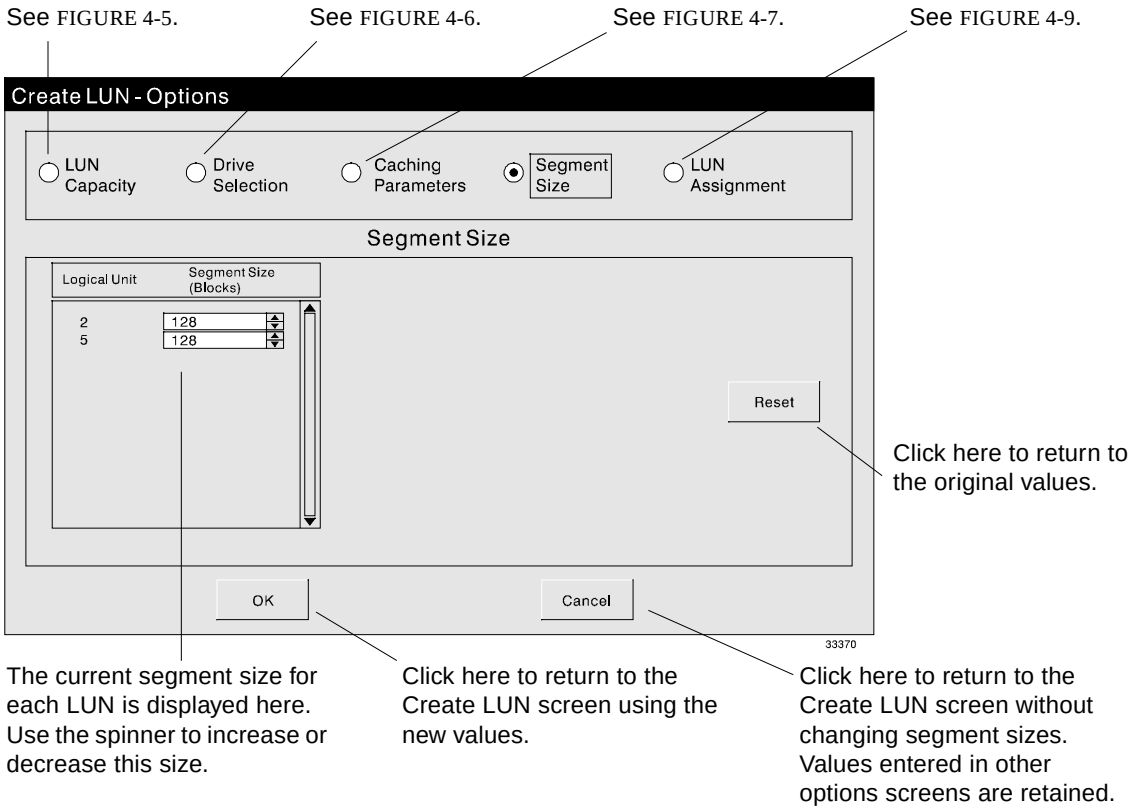

FIGURE 4-8 Segment Size Screen

- You can set the segment size for individual LUNs.
- LUNs in a drive group do not need to have the same segment size value.
- You can change the segment size of a LUN after you create it; see “Modifying Segment Size” on page 89.

LUN Assignment Option

Use this option to change the controller assignment of the new LUNs.

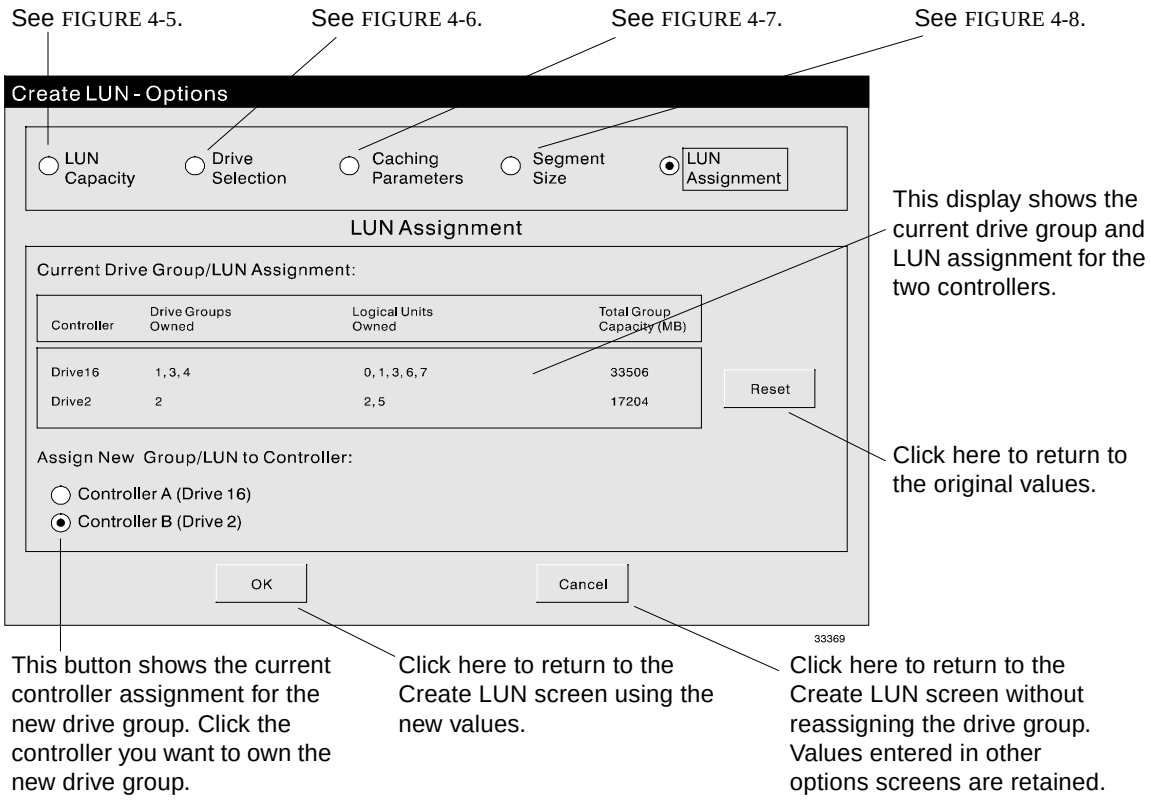

FIGURE 4-9 Set LUN Assignment Screen

Modifying Drive Groups and Logical Units

Once a drive group / logical unit is created, you can modify it in the following ways:

- Defragment the logical units on the drive group – This may increase total available capacity. See "The Defragment Option" on page 81 for a description of defragmenting and when it applies.
- Change the RAID level – You can change the RAID level of all the logical units in a drive group. All logical units in the drive group must still have the same RAID level. You may be limited in the new RAID level by the number of drives in the drive group; for example, you cannot change a 5-drive RAID 5 to RAID 1 because RAID 1 must have an even number of drives. See "Defragmenting the Drive Group" on page 84.
- Add Capacity – You can increase the capacity of a *drive group* by adding drives. This does *not* increase the capacity of the LUNs in the drive group. The amount of capacity increased depends on the number of drives added and the RAID level (see TABLE 4-2). See "Adding Drives" on page 87.
- Modify the segment size – You can change the segment size of individual logical units (you don't have to change all the logical units in the drive group). This may improve performance. See "Modifying Segment Size" on page 89.
- Change the cache parameters – You can change the Write Caching, Write Cache Mirroring, and Cache Without Batteries parameters using the Maintenance and Tuning application. See "Changing Cache Parameters" on page 167. You can change other cache parameters using the `rdacutil` utility. For more information about the `rdacutil` utility, refer to Appendix A.
- Change the reconstruction rate – you can change the reconstruction rate assigned to individual logical units. See "Changing the Reconstruction Rate" on page 158.

You *cannot* make the following changes without deleting a logical unit and recreating it.

- Change the capacity of an individual LUN.
- Decrease the number of drives in the drive group.

Note – You cannot change the RAID level, capacity, or segment size or defragment the drive group until background formatting is completed (see "Logical Unit (LUN) Formatting" on page 64).

Also, make sure you have closed all other storage management applications before starting a LUN modification operation.

The Defragment Option

The Defragment option is used to reclaim space on your drive group that might be inaccessible after LUN deletion. This is a result of the way data is written to the logical units.

For example, FIGURE 4-10 shows a three-drive logical unit. Each logical unit has a capacity of 600 MB, with 200 MB written on each disk (ignore RAID level effects on capacity for purposes of this example). Remaining capacity in the drive group is 900 MB.

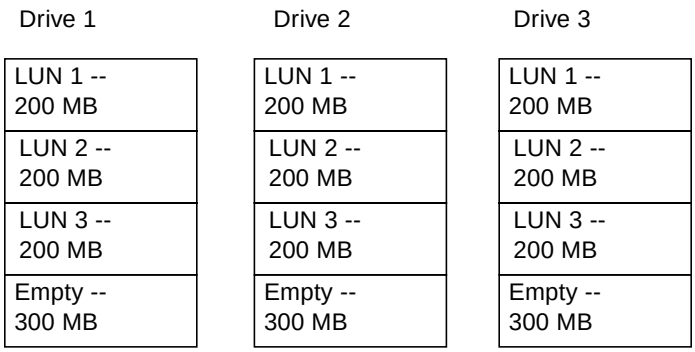

FIGURE 4-10 Defragment Logical Units #1

Now assume LUN 2 (the middle LUN) is deleted. This is shown in FIGURE 4-11.

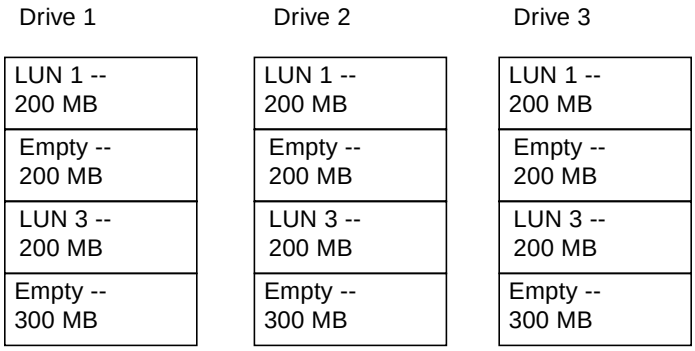

FIGURE 4-11 Defragment Logical Units #2

Even though there is actually 1500 MB of empty space in this drive group, the largest logical unit that can be created is 900 MB, because a logical unit must be created in a contiguous area. The Configuration application shows the remaining capacity as 900 MB, ignoring the additional 600 MB.

Next, assume you create a new LUN 2 of 600 MB. The new logical unit is written in the 300 MB remaining capacity, as shown in FIGURE 4-12.

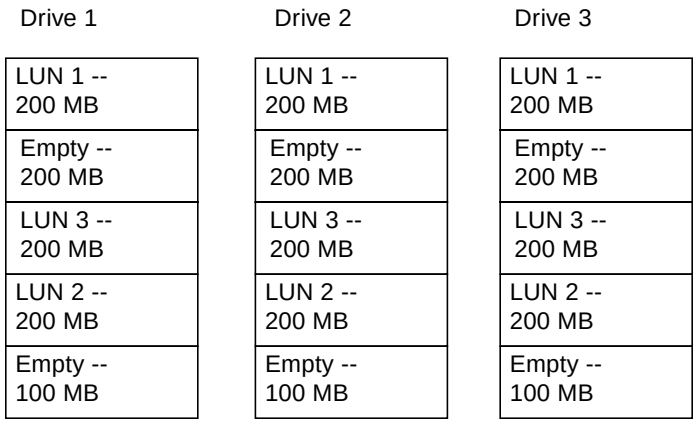

FIGURE 4-12 Defragment Logical Units #3

Total remaining capacity is 900 MB, but the configuration utility shows a capacity of 600, because that is the biggest single area in which a logical unit can be created. The remaining 300 MB is unavailable until you create a LUN larger than 300 MB in the free area. Configuration always shows the largest contiguous area available.

The Defragment option rewrites the data on the drive group so that all the free space is contiguous. For example, if the Defragment option is performed on the drive group in FIGURE 4-12, the resulting drive group looks like FIGURE 4-13.

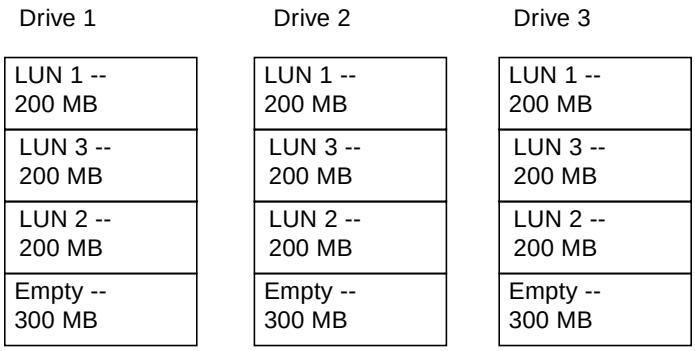

FIGURE 4-13 Defragment Logical Units #4

Note that the logical units were “moved down” to fill in the gap left by the deletion of the logical unit. The remaining capacity for this drive group is shown as 900 MB again, because all the free space is now contiguous.

You may need to run the Defragment option after deleting a logical unit from a drive group if there is more than one logical unit on that drive group. The option tells you if there is any space that needs to be reclaimed on the drive group. See “Defragmenting the Drive Group” on page 84 for procedural information.

Defragmenting the Drive Group

Defragmenting the drive group can free up space after the deletion of a logical unit from the drive group. See “The Defragment Option” on page 81 for more information.

To defragment the drive group, perform the following procedure.

Note – You cannot defragment the drive group while background formatting is in progress (see “Logical Unit (LUN) Formatting” on page 64). If the selected drive group contains a LUN that is still formatting in the background, you will get a message notifying you that a LUN is still initializing when you select this option.

---

1. **Close all storage management applications.**
2. **Start the Configuration application.**
3. **Select the RAID Module containing the drive group you want to defragment.**
4. **Select the drive group you want to defragment in the Module Information area.**
5. **Click on the Modify LUNs button, then select Defragment, or select the option from the drop-down menus.**
6. **A message is displayed, stating that *either*:**
   - There is no space to recover in the drive group you selected. Click OK to return to the Main Configuration screen. You have completed this procedure.
   - There is space to be recovered on the selected drive group. The message indicates how much space can be recovered. Click OK to recover the space, or click Cancel to return to the Main Configuration screen.
7. **If there is space to recover and you clicked OK, the software begins to defragment the drive group.**

   As the controller defragments the drive group, the configuration display shows a status of “Modifying” for the drive group and for the individual logical units. The status gives a percentage, indicating the percentage of the individual logical unit modified and the percentage of the overall drive group modified. When the status of all elements changes to Optimal, the operation is complete.

   End of Procedure

Changing RAID Level

You can change the RAID level of all logical units in the drive group. Keep in mind that you may be limited by the number of drives in the group. For example, if the drive group has an odd number of drives, you cannot change it to RAID 1 (you could, however, add a drive to the drive group using the procedure in “Adding Drives” on page 87 and then change to RAID level 1).

To change the RAID level of all the logical units in a drive group, perform the following procedure.

Note – You cannot change the RAID level while background formatting is in progress (see “Logical Unit (LUN) Formatting” on page 64). If the selected drive group contains a LUN that is still formatting in the background, you will get a message notifying you that a LUN is still initializing when you select this option.

---

1. **Close all storage management applications.**
2. **Start the Configuration application.**
3. **Select the RAID Module containing the drive group/logical units you want to modify.**
4. **Select the drive group you want to modify in the Module Information area.**
5. **Click on the Modify Group/LUNs button, then select Modify RAID Level, or select the option from the drop-down menus.**

   A screen like FIGURE 4-14 is displayed.

6. **Select the RAID level you want for the drive group.**
7. **After you have selected the new RAID level, click OK to make the change.**

   As the controller changes the RAID level, the configuration display shows a status of “Modifying” for the drive group and for the individual logical units. The status gives a percentage, indicating the percentage of the individual logical unit modified and the percentage of the overall drive group modified. When the status of all elements changes to Optimal, the operation is complete.

   End of Procedure

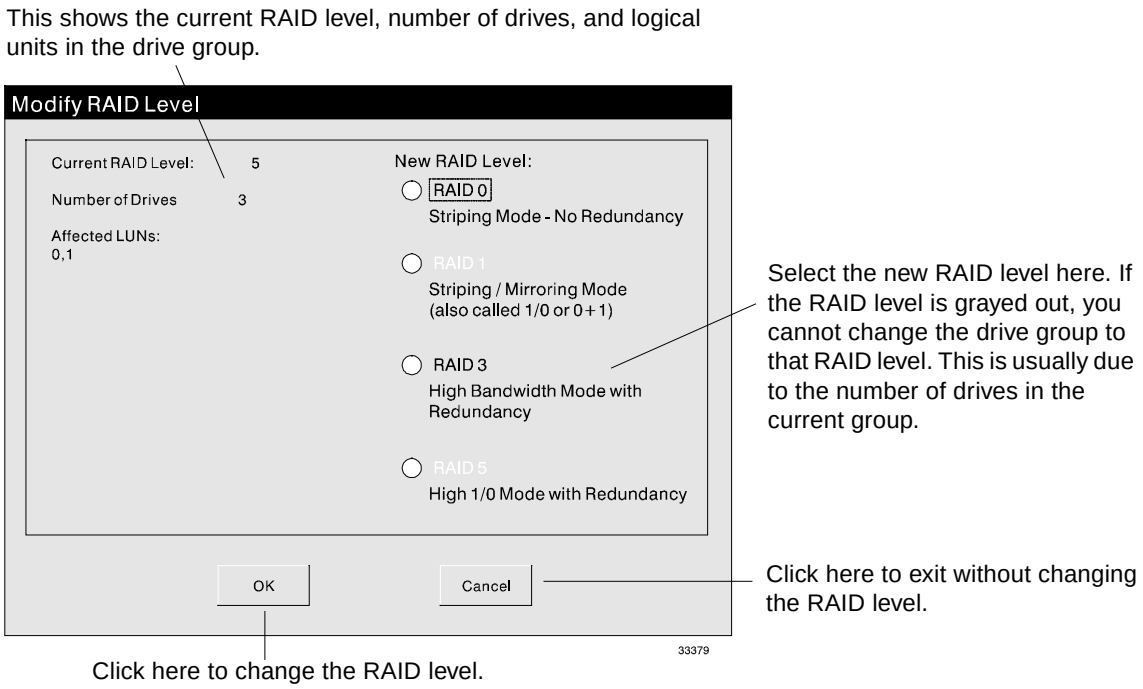

FIGURE 4-14 Modify RAID Level Screen

Adding Drives

You can increase the capacity of a drive group by adding drives to the group. TABLE 4-2 gives the formulas for calculating how much space is added to the drive group.

TABLE 4-2 Adding Capacity Formulas

RAID Level	The Drive Group Capacity Is *Increased* by:
RAID 1	One half the Number of Drives Added x Capacity of Drives
RAID 0, 3, or 5	Number of Drives Added x Capacity of Drives

For example, assume a five-drive group containing RAID 5 logical units. Each drive is 18 GB, giving a current group capacity of 72 GB (one drive's worth of capacity is required for parity). Adding two drives to the drive group would increase the capacity of the drive group by 2 x 18, or 36 GB. The drive group would then have a capacity of 72 + 36, or 108 GB.

Note – Adding drives *does not* increase the size of the LUNs in the drive group. The LUNs are spread out over the added drives. However, after adding drives, the drive group will have capacity for new LUN creation.

To add drives to the drive group, perform the following procedure.

Note – You cannot add drives to the drive group while background formatting is in progress (see "Logical Unit (LUN) Formatting" on page 64). If the selected drive group contains a LUN that is still formatting in the background, you will get a message notifying you that a LUN is still initializing when you select this option.

---

1. **Close all storage management applications.**
2. **Start the Configuration application.**
3. **Select the RAID Module containing the drive group you want to modify.**
4. **Select the drive group you want to modify in the Module Information area.**
5. **Click on the Modify Group/LUNs button, then select Group Capacity, or select the option from the drop-down menus.**

   A screen like FIGURE 4-15 is displayed.

6. **Select the drives you want to add to the drive group.**

   Choose drives of the same size as the current drives in the group.

7. **After you have selected the drives, click OK to make the change.**

8. **The drives are added to the drive group.**

   As the controller adds the drives, the configuration display shows a status of "Modifying" for the drive group and for the individual logical units.

   The status gives a percentage, indicating the percentage of the individual logical unit modified and the percentage of the overall drive group modified. When the status of all elements changes to Optimal, the operation is completed.

   End of Procedure

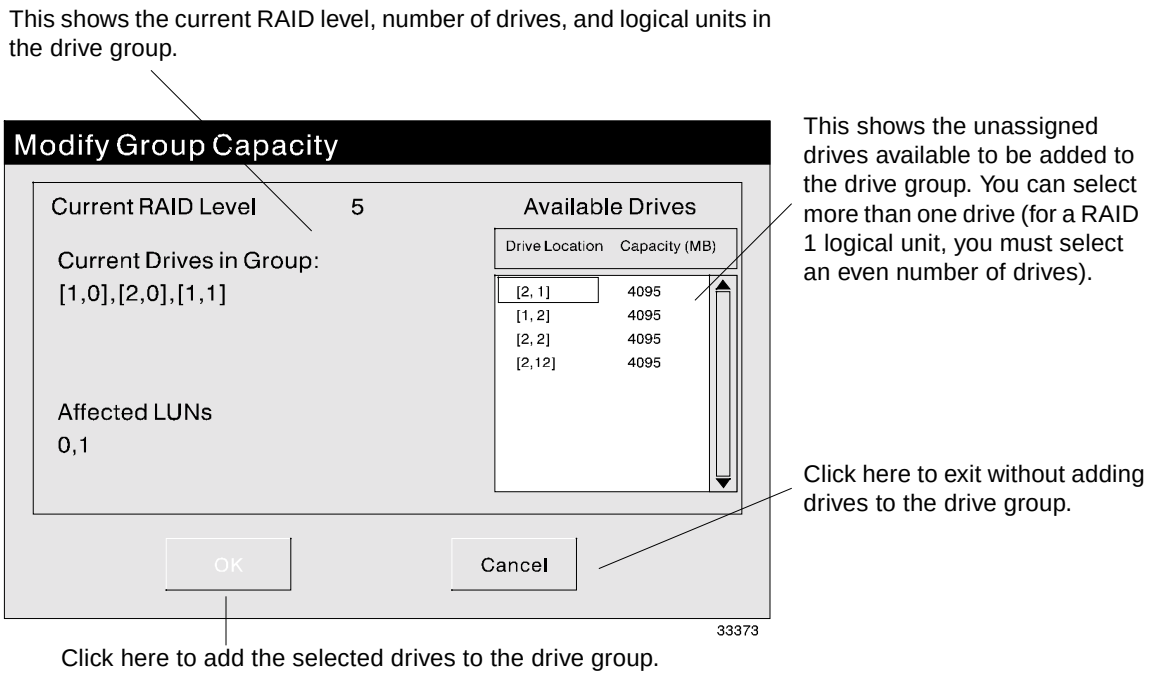

FIGURE 4-15 Modify Group Capacity Screen

Modifying Segment Size

You can modify the segment size of a logical unit. You do not need to change the segment size of all the logical units in a drive group.

To change the segment size of a logical unit, perform the following procedure.

Note – You cannot change the segment size while background formatting is in progress (see “Logical Unit (LUN) Formatting” on page 64). If the selected LUN is still formatting in the background, you will get a message notifying you that it is still initializing when you select this option.

---

1. **Close all storage management applications.**
2. **Start the Configuration application.**
3. **Select the RAID Module containing the logical units you want to modify.**
4. **Select the drive group or logical unit you want to modify in the Module Information area.**
5. **Click on the Modify Group/LUNs button, then select Modify Segment Size, or select the option from the drop-down menus.**

   A screen like FIGURE 4-16 is displayed.
6. **One at a time, select the logical units you want to modify, then pick the new segment size.**
7. **After you have changed the segment size, click OK to make the change.**
8. **The controller changes the segment size.**

   As the segment size changes, the configuration display shows a status of “Modifying” for the drive group and for the individual logical units.

   The status gives a percentage, indicating the percentage of the individual logical unit modified and the percentage of the overall drive group modified. When the status of all elements changes to Optimal, the operation is completed.

   End of Procedure

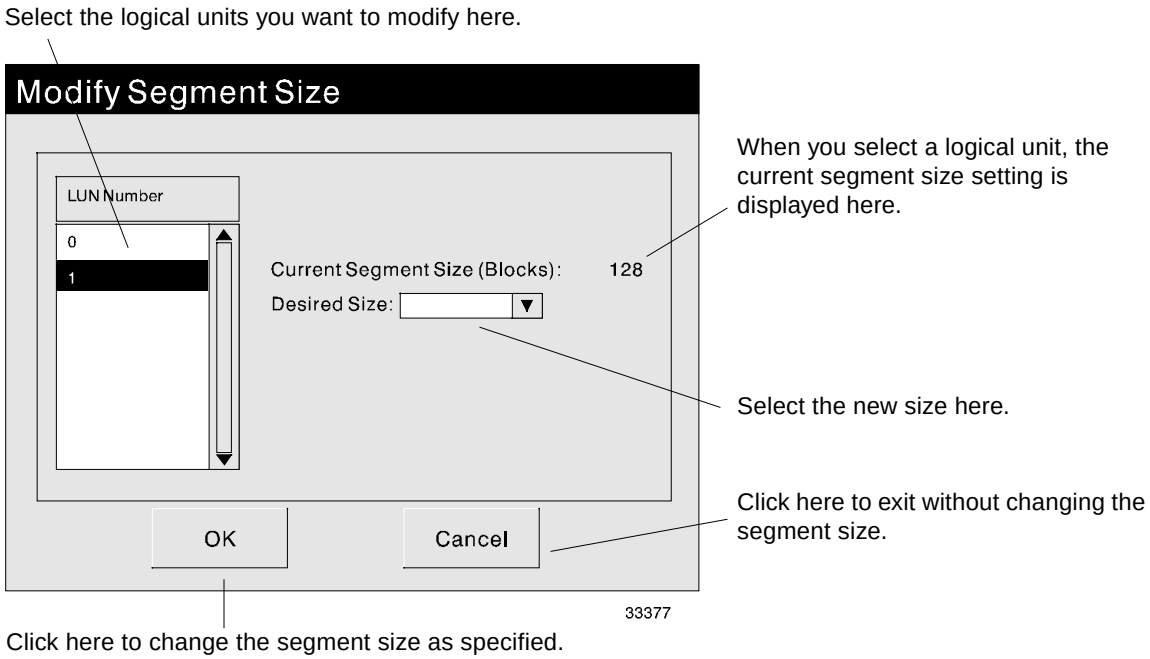

FIGURE 4-16 Modify Segment Size Screen

Deleting Logical Units

When you delete a logical unit (LUN), any data on the LUN is lost. If you delete all of the LUNs in a drive group, the drives in the drive group return to the unassigned drive group.

Caution – Do not delete all the LUNs in the RAID Module and exit configuration without creating at least one LUN. Doing so could prevent you from accessing the RAID Module.

Before you delete a LUN, check the *Sun StorEdge RAID Manager Installation and Support Guide* to see if there are any special restrictions or special requirements on LUN deletion. For example, you may be required to first delete partitions, de-assign drives, or unmount file systems before deleting. After you delete a LUN, you may need to perform procedures on your operating system before it recognizes that the LUN is gone. For example, you may need to restart your system before the operating system records the LUN deletion. Also see the special consideration "Deleting LUNs in the Networked Configuration" on page 92.

To delete a LUN, perform the following procedure.

1. **Perform whatever steps are necessary on your operating system to prepare for LUN deletion.**

   See your operating system documentation and the *Sun StorEdge RAID Manager Installation and Support Guide* for information.

2. **Start the Configuration application.**
3. **Select the RAID Module containing the LUN you want to delete.**
4. **Select the drive group containing the LUN you want to delete in the Module Information area.**
5. **Press the Delete key, select Options → Delete from the drop-down menus, or right click on the drive group containing the LUN, then select Delete.**

   A screen like FIGURE 4-17 appears.

6. **Select the LUNs you want to delete, or click Select All to delete all the LUNs in the drive group.**
7. **When the LUNs you want to delete are highlighted, click Delete.**

Note – Deleting LUNs may cause drive groups to be renumbered. See "Drive Group Renumbering" on page 12.

8. **After the LUNs are deleted, do whatever is necessary on your operating system to delete a drive. This might involve restarting the system, clearing file tables, and so on.**

   See your operating system documentation and the *Sun StorEdge RAID Manager Installation and Support Guide* for more information.

   End of Procedure

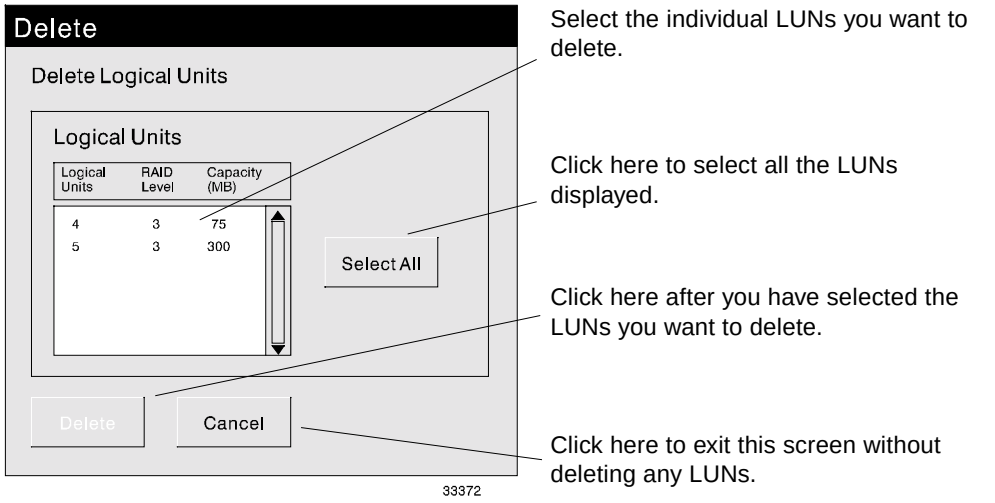

FIGURE 4-17 Delete LUNs Screen

Deleting LUNs in the Networked Configuration

If you are using the Networked version of the storage management software or a host-RAID Module configuration with more than one host (see "RAID Module Configurations" on page 19), it is possible to open a Configuration application and begin the Delete operation from more than one host/station. However, you must be sure that you are using only one configuration session at a time or the drive group/LUN you delete could be in use and cause problems for one of the hosts.

Creating Hot Spares

Hot spares add another level of protection to the data on your RAID Modules. Hot spares automatically replace failed drives in your logical units. You can only create hot spares if you have unassigned drives.

Tip – A hot spare cannot replace a drive with a larger capacity. If you have drives of two different capacities in your RAID Modules, you should select a drive of the *larger* capacity for your hot spare drives.

To create one or more hot spare drives, perform the following procedure.

1. **Start the Configuration application.**
2. **Select the RAID Module you want to create hot spares for.**
3. **Highlight the Unassigned drive group or the Hot Spare drive group.**
4. **Click on the Hot Spare button or select Options → Create Hot Spare from the drop-down menus.**

   A screen like FIGURE 4-18 is displayed.

5. **Enter the number of hot spares you want to create.**

   You can create as many hot spares as there are drive channels on your RAID Module (this may be further limited by the number of unassigned drives remaining in the RAID module).

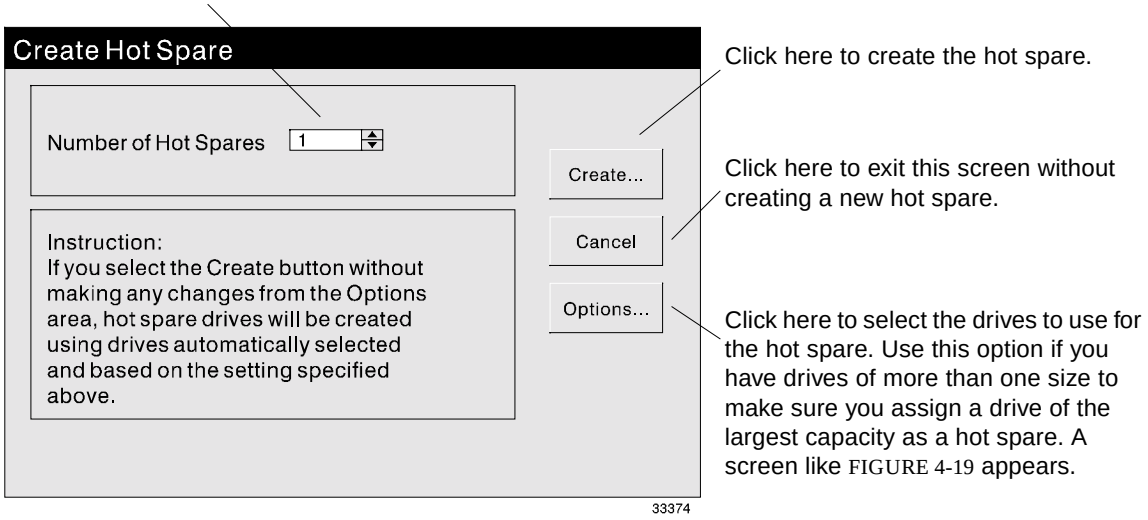

FIGURE 4-18 Create Hot Spare Screen

6. **If you want, select Options to physically select the drive you want to assign as a hot spare.**

   A screen like FIGURE 4-19 is displayed.

   In most cases, you should accept the default choice instead of selecting drives manually. It may be necessary to use this option if you have unassigned drives of more than one size and need to make sure that you have selected the highest capacity drive as the hot spare.

7. **Click OK to create the hot spare drives.**

   End of Procedure

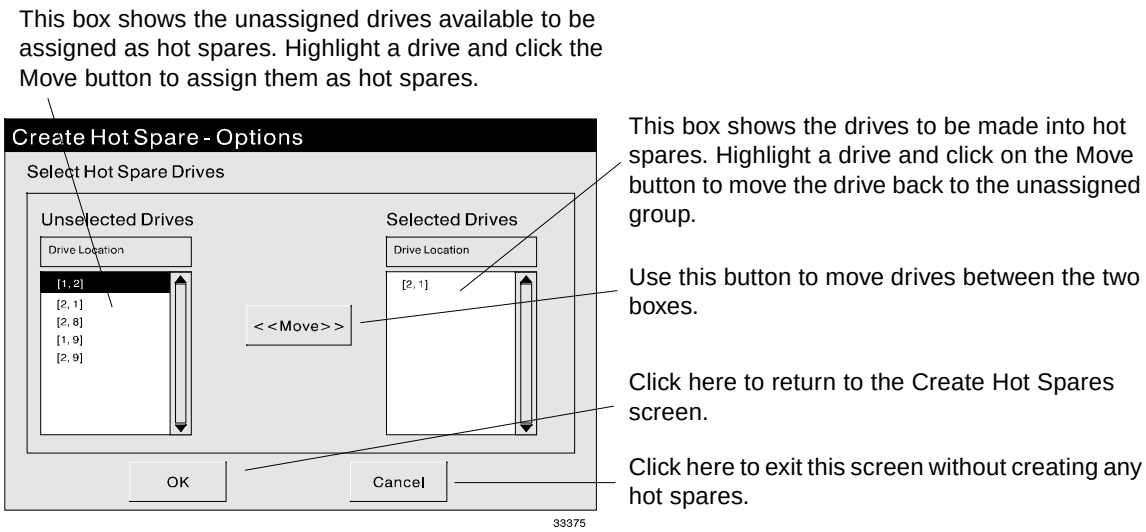

FIGURE 4-19 Select Hot Spare Screen

- You can create as many hot spares on the RAID Module as there are drive channels.
- Always use the largest capacity drives for hot spares.

Deleting a Hot Spare

Note – You cannot delete a hot spare that is currently assigned to a drive group (taking the place of a failed drive).

When you delete a hot spare, the drive returns to the unassigned drive group.

To delete a hot spare, perform the following procedure.

1. **Start the Configuration application.**
2. **Select the RAID Module containing the hot spare you want to delete.**
3. **Select the Hot Spare drive group.**
4. **Press the Delete key or select Options → Delete from the drop-down menus.**

   A screen like FIGURE 4-20 is displayed, listing all the drives currently assigned as hot spares.
5. **Select the hot spares you want to delete, then click Delete.**
6. **The selected hot spares are deleted and returned to the Unassigned drive group.**

   End of Procedure

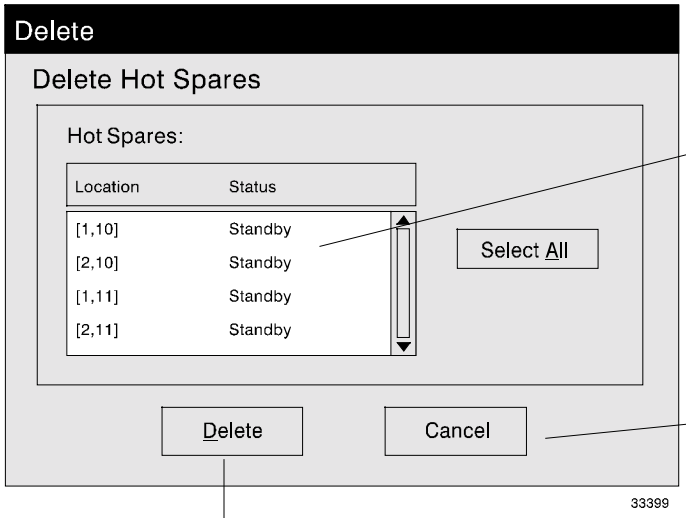

FIGURE 4-20 Delete Hot Spares Screen

Resetting the Configuration

Caution – Use this option only as a last resort if your configuration is totally inaccessible or you want to completely redo your configuration.

When you reset the configuration, the selected RAID Module (*all* drive groups and their logical units) is reset back to a default configuration based on the NVSRAM settings specified in the controller.

- Because deleting LUNs causes data loss, back up data on *all* the drive group / LUNs in the RAID Module. This operation also deletes any file systems mounted on the LUNs.
- You must first stop I/Os to the affected RAID Module and you should ensure no other users are on the system.
- You cannot reset the configuration on a RAID Module with an independent controller configuration. For a description of this configuration, see "Independent Controller Configuration" on page 23.
- You cannot reset the configuration if the controllers are currently active / passive. You must change them to active / active before resetting. (See "Changing Controller Mode" on page 165 for information on how to change controller mode to active / active.)
- This operation could also fail if the storage management software cannot gain exclusive access to the drive groups / logical units (for example, if file systems are mounted).
- You must restart your system so the new configuration is recognized by the operating system.

To reset the configuration on the RAID Module, perform the following procedure.

1. **From the top menu, select File → Reset Configuration.**
2. **Select OK to confirm that you want to reset your configuration.**

Caution – Resetting the configuration destroys *all* data on *all* the logical units in your RAID Module! Step 3 is your last chance to Cancel!

3. **Select OK to confirm again that you want to reset your configuration.**

   A default configuration appears in the main Configuration screen.

Note – This does *not* necessarily mean that you have all unassigned drives; you may have a small LUN configured.

4. **Select OK at the "Reset Was Successful" confirmation screen.**
5. **Restart your system.**
6. **You will need to redefine all of your logical units and drive groups using Create LUN.**

Note – Your operating system may have additional requirements to complete the configuration process so that it can recognize the new LUNs. Refer to the *Sun StorEdge RAID Manager Installation and Support Guide* specific to your operating system (see the sections describing restrictions and troubleshooting) and to the appropriate system documentation for specific details.

End of Procedure

CHAPTER 5

Checking Status

This chapter describes how to check RAID Module status and performance.

Status Overview

The storage management software monitors the status of the RAID Modules in the background and logs any status changes, or “events,” in a log file. These events may be component failures, parity check / repair operations, or configuration changes. You use the Status application to view and maintain this log file.

Background Polling

The storage management software polls all RAID Modules periodically to find any status changes or other “events.” When the software finds a status change, it writes detailed information about the event into the log file.

You can set the polling interval for this background check (see “Changing Log File Settings” on page 108). The default time is 5 minutes. You can set any time from 1 to 59 minutes. However, setting too short of a time interval may slow system performance because of the time required to poll, and setting too large a time may delay notification of serious failures.

The Message Log File

The message log is the file the application uses to store event messages. The default file is defined by your operating system (see the *Sun StorEdge RAID Manager Installation and Support Guide*), but can be changed to any file (see “Changing Log File Settings” on page 108).

The log file has a maximum size limit, which is set by default to 40K (this limit can be changed; see “Changing Log File Settings” on page 108). When you open any storage management application after this limit has been reached, the log file is automatically copied to a backup file and all messages are erased (see“Clearing the Message File” on page 110 for a description of this procedure). Until you open an application and reset the log file, however, messages continue to be written to the file, even after the limit has been reached.

You can save the contents of the log file by selecting File → Save Log As from the drop-down menus. This does not delete the original file or change the default file. Saving the log saves the entire contents of the log file, regardless of which RAID Module or message types are currently selected.

You can open a saved file later by selecting File → Open Log from the drop-down menus. This does not change the default log file.

Message Types

The software has three categories of events as shown in TABLE 5-1:

TABLE 5-1 Message Types

Type	Description
Parity	Messages related to parity check / repair (either manual or automatic). The message gives the location of any bad blocks found, the affected logical unit, and the number of blocks repaired.
Hardware	Messages related to component failures. The messages give the affected component, the affected logical unit (if any), the probable cause, and the action to take (usually you will be advised to run Recovery Guru / Health Check). The messages may also give ASC (Additional Sense Code) and ASCQ (ASC Qualifier) information. This is SCSI information providing further details about the failure. This information is most useful for troubleshooting by Customer Services personnel.
General	Messages related to log file backups, I / O errors, configuration changes, and other event types. The messages try to give further descriptions of the events. In many cases, they will require no action.

The Status Log and Health Check

Note that, in some cases, you may discover a problem when running Recovery Guru / Health Check that is not found in the Message Log. This is due to the delay in the polling interval. If the error occurred after the last polling cycle, but before the current one, Health Check finds the error even though it is not yet in the log.

Event Notification

Depending on the operating system, you can set up SNMP or other scriptable notification to perform some other action when an event is found. See the *Sun StorEdge RAID Manager Installation and Support Guide* for details.

Viewing the Status Log

To view the status log, perform the following procedure.

1. **Open the Status application.**
2. **Select the RAID Module you want to view status information for (or select All RAID Modules).**
3. **Click on the Message Log button or select Options → Message Log from the drop-down menus.**

   The contents of the currently selected message log is displayed (if you didn't select another message log file, this is the default message log). FIGURE 5-2 shows an example of the screen displayed.
4. **If you want to view a log other than the current default message log, go to File → Open Log and specify the log file you want to view.**
5. **If desired, click on List Type to filter the types of messages displayed (see** FIGURE 5-1**).**
6. **Highlight a displayed message or messages and click Show Details to get more information about the message.**

   A screen like FIGURE 5-3 is displayed.
7. **View the details and take the appropriate action.**

8. **You can also perform the following actions:**
   - Save the message log file to another file – Select File Save → Log As from the drop-down menus. See "The Message Log File" on page 102.
   - Open another, previously saved, log file – Select File → Open Log from the drop-down menus.
   - Copy message text to the clipboard – Highlight text in the Message Details screen and select Edit → Copy To Clipboard. You can only perform this option from the Message Details screen (FIGURE 5-3).
   - Refresh the screen – Select Options → Refresh All from the drop-down menus. This updates the message log with any events that occurred since you opened the log. You can only perform this option from the Message Log screen (FIGURE 5-2).
9. **When you are finished viewing the messages, close the Status application or select another option.**

End of Procedure

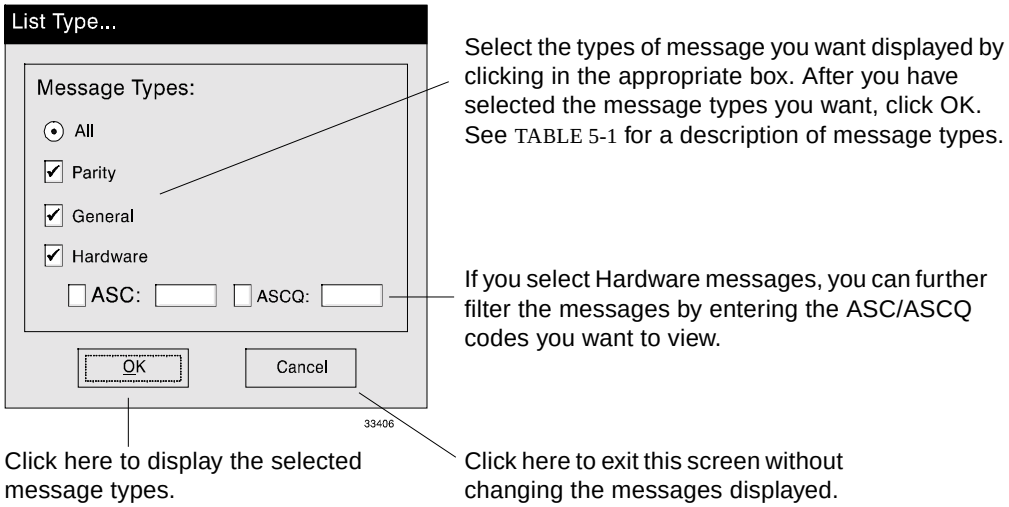

FIGURE 5-1 Selecting Message Type

The messages shown in the display are for the selected RAID Module.

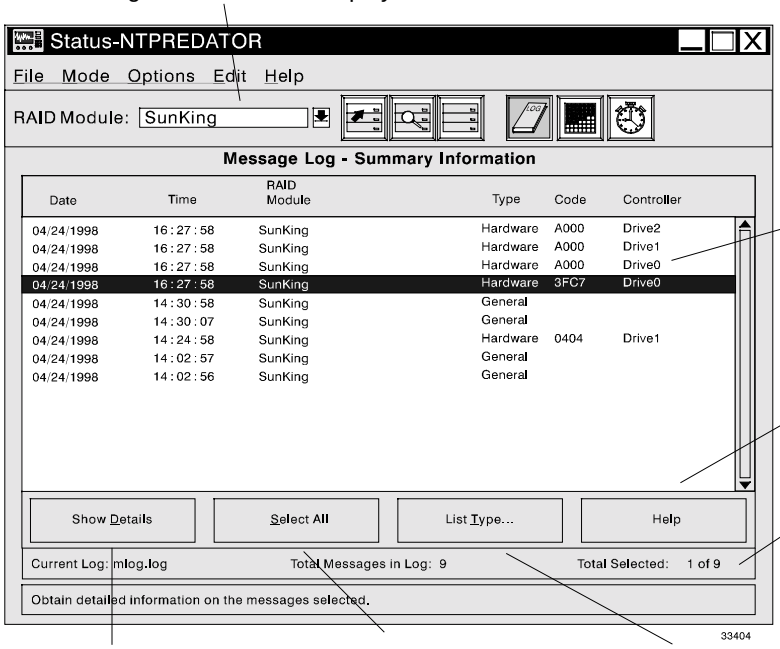

This display gives summary information on the messages in the currently selected message log. TABLE 5-2 explains the entries shown.
Highlight the messages you want to view, then click Show Details.

Click here for Help.

This bar shows the selected message log, the number of messages in the log, and the number of messages selected.

Click here to display detailed information on the selected messages (see FIGURE 5-3).

Click here to select all the messages, then click Show Details.

Click here to specify the types of messages you want to display (see FIGURE 5-1).

FIGURE 5-2 Message Log Screen

TABLE 5-2 Message Log Explanation

Column Heading	Explanation
Date & Time	Indicates when the detected event was logged. • If the event is a component failure detected by the background monitor, this is the time the error was detected, not necessarily when it occurred. • If the event is a parity error, this is the time the error was detected and corrected by Parity Check/Repair. • If the event is a general status change, this is the time the change was detected.
RAID Module	The RAID Module where the event occurred.
Type	The type of RAID Module event; either Hardware, Parity, or General (see "Message Types" on page 103).
Code	The ASC/ASCQ code for the hardware message, if any. Help has a summary of these codes.
Controller	The system device name of the affected controller.

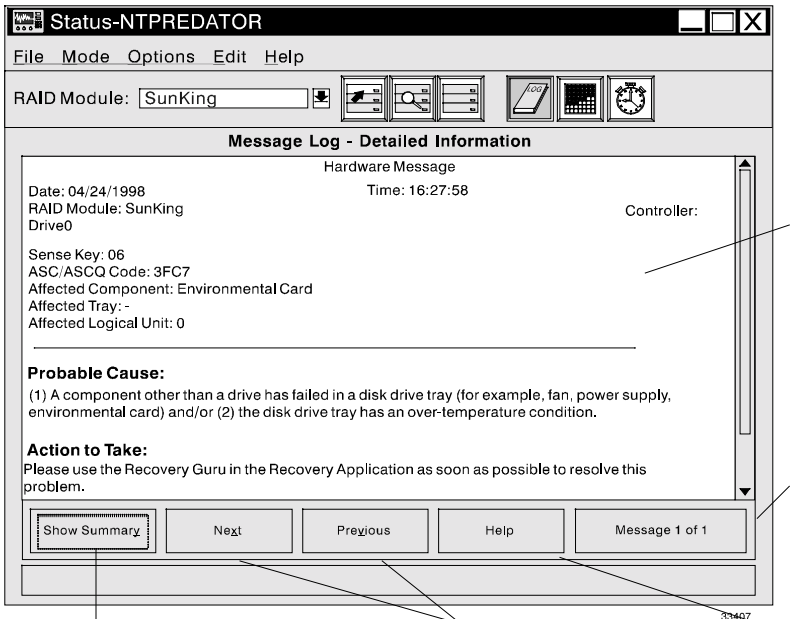

This displays detailed information on a selected message. The message includes information on the affected component, the probable cause, and the action to take. Usually, you should run Recovery Guru for a complete description of the action to take.

This shows the message you are viewing and the total number of messages selected.

Click here to return to the message summary screen (see FIGURE 5-2).

If you have selected more than one message, click Next or Previous to move through the selected messages.

Click here for help.

FIGURE 5-3 Message Details Screen

Changing Log File Settings

You can change the following log settings (see “The Message Log File” on page 102):

- Default log file – This is the file where RAID events are written to.
- Maximum Message Log Size – This is the file size that must be exceeded before the file is backed-up and cleared (see “Clearing the Message File” on page 110). Note that setting a size does not limit the size of the file. The file grows beyond this limit until you open a storage management application and reset it.
- Polling interval – This is the frequency in minutes that the background monitor checks the RAID Modules. Setting too small a value may affect system performance. Setting too large a value may delay notification of serious errors.

To change the log file settings, perform the following procedure.

1. **Start the Status application.**
2. **Select *any* RAID Module (this change affects all RAID Modules).**
3. **Select Options → Log Settings from the drop-down menus.**

   A screen like FIGURE 5-4 is displayed.

4. **Enter the values you want, then click Save (or view the current values and click Cancel to exit without changing values).**

   End of Procedure

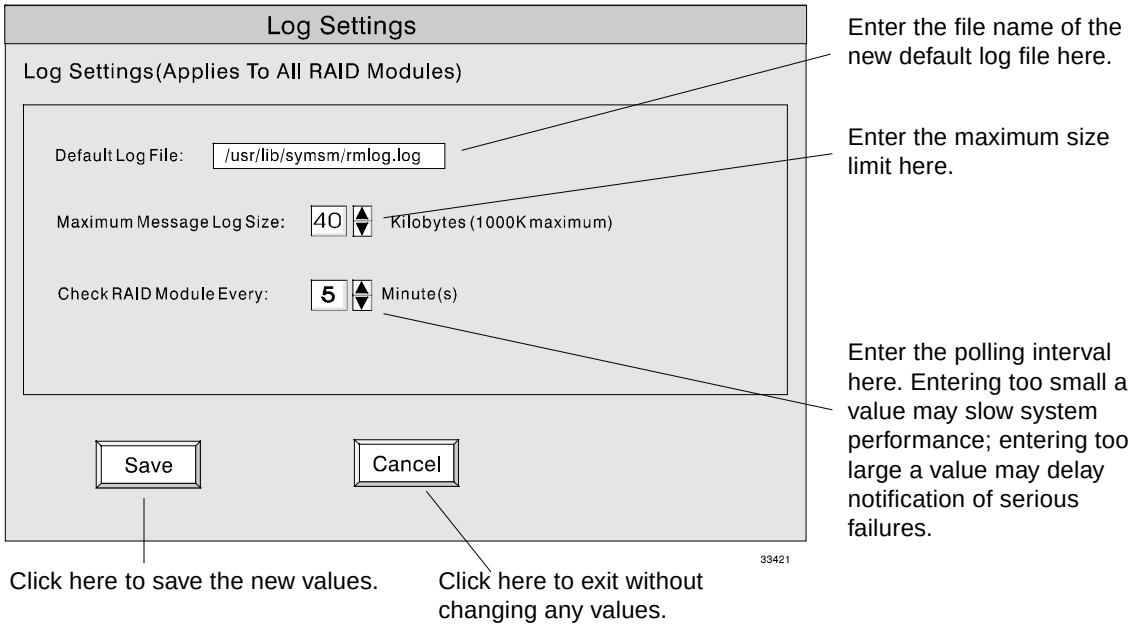

FIGURE 5-4 Log Settings Screen

Clearing the Message File

If you open a storage management application and the message log file currently exceeds the defined maximum size limit (see “Changing Log File Settings” on page 108), the message log file is saved and cleared. The following sequence occurs:

1. **When you open the storage management application, a warning message like** FIGURE 5-5 **is displayed.**
2. **Click OK.**
   - If the `rmlog.<hostname>_1` file does not exist, the message file is backed up to the `rmlog.<hostname>_1` file and cleared.
   - If the file `rmlog.<hostname>_1` already exists, a screen like FIGURE 5-6 is displayed.
3. **Click Yes if you want to write over the current `rmlog.<hostname>_1` file. Click No if you want to create the next file in the sequence (_2, _3, _4, and so on).**
4. **The first message in the new log file will notify you that this backup occurred. Highlight the message and click Show Details to find the name of the backup file used.**

End of Procedure

Saved Log File Location

The saved log file is placed in a default directory, named in the message. You can change this directory by changing the System LogFileDefaultBackupPath parameter in the `rmparams` file. For more information on the `rmparams` file, refer to the *Sun StorEdge RAID Manager Installation and Support Guide*.

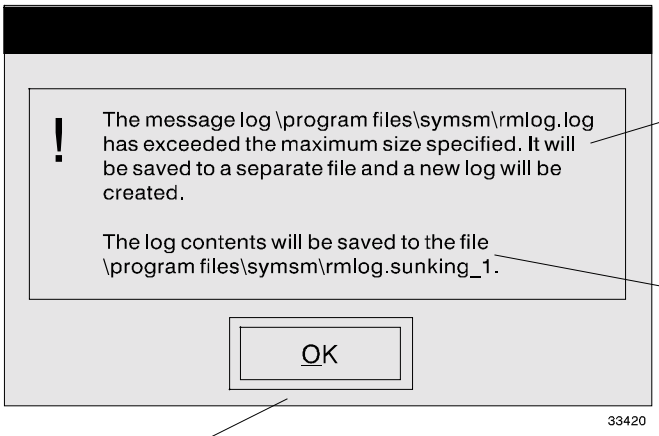

This screen indicates that the default message file has reached or exceeded its defined maximum size limit.

This is the name of the backup file to be created. The default save name is always `rmlog.<hostname>_1`.

Click OK to continue. If the `rmlog.<hostname>_1` file does not exist, the current contents of the file is written to the backup file and the message log file is cleared. If the `rmlog.<hostname>_1` file already exists, a screen like FIGURE 5-6 is displayed when you click OK.

FIGURE 5-5 Message Limit Warning Message

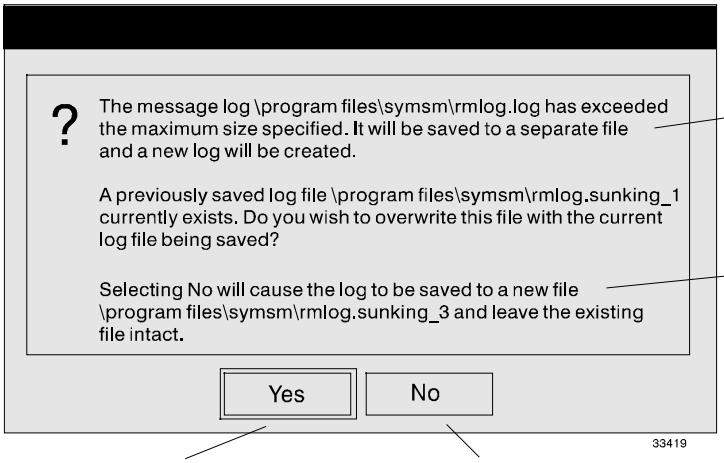

This screen tells you that a file named `rmlog.<hostname>_1` already exists.

If you click No, the next file in the sequence, given here, will be created.

Click Yes to write over the current `rmlog.<hostname>_1` file.

Click No to copy the message log to a new file.

FIGURE 5-6 Message Log Save Message

Using the Performance Monitor

Note – Before you start the Performance Monitor, make sure that I/O is being processed on the RAID Modules you want to monitor.

To run the Performance Monitor, perform the following procedure.

1. **Start the Status application.**
2. **Select the RAID Module you want to monitor. (Do not select All RAID Modules.)**
3. **Click on the Performance Monitor button or select Options → Performance Monitor from the drop-down menus.**

   A screen like FIGURE 5-7 is displayed.

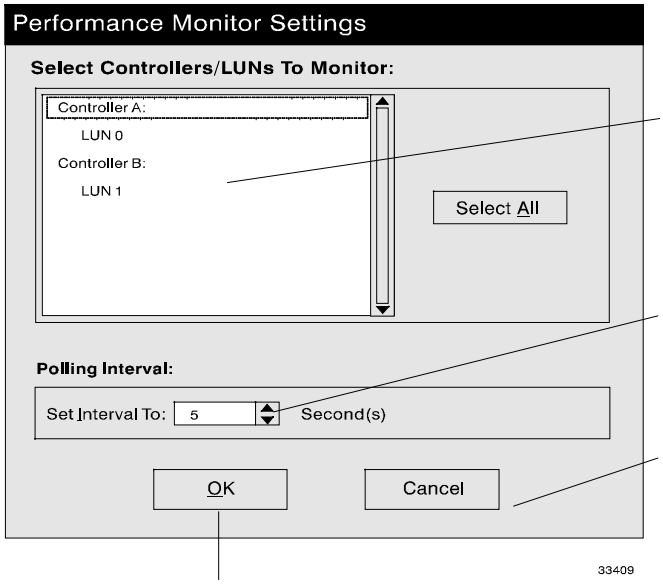

Select the controllers or LUNs you want to monitor here. Click Select All to highlight all the elements in the RAID Module.

Set the Polling interval here. This is the frequency which the Performance Monitor uses to check I/O activity on the RAID Module.

Click here to exit without starting the Performance Monitor.

Click here to start the Performance Monitor using the current settings.

FIGURE 5-7 Performance Monitor Settings Screen

4. **Select the Controllers or LUNs you want to monitor and click OK.**

   A screen like FIGURE 5-8 is displayed.

5. **Click Start to start the Performance Monitor.**

   A screen like FIGURE 5-9 is displayed.

6. **You can save the displayed performance data as a text file by selecting File → Save Performance Data from the drop-down menus.**

7. **After you have gathered all the data you want, click Stop to end performance monitoring.**

   End of Procedure

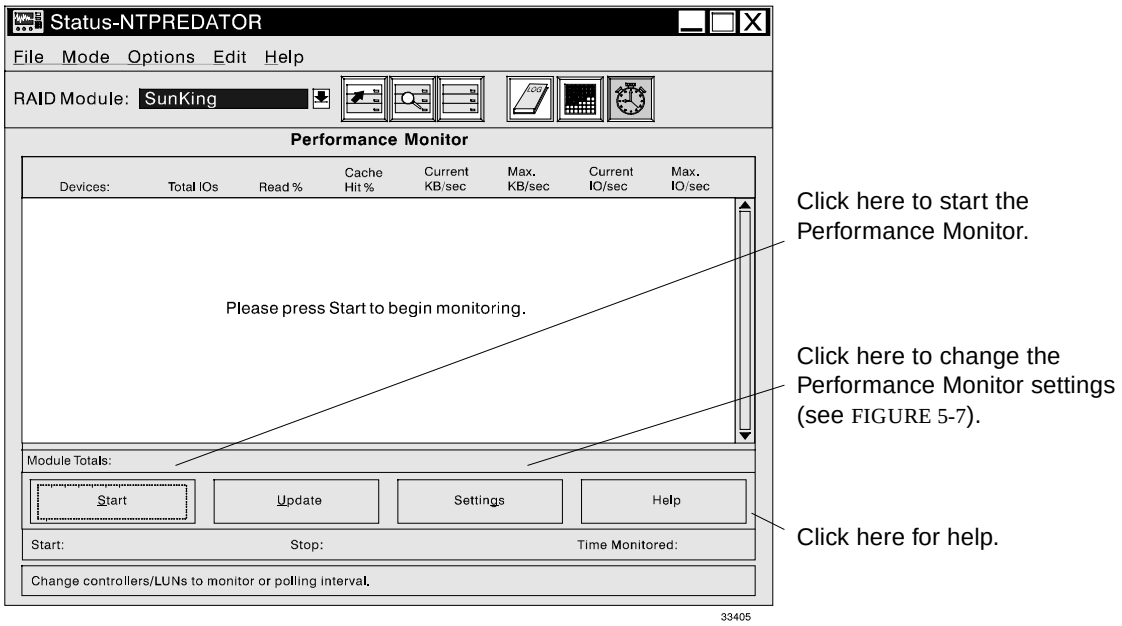

FIGURE 5-8 Performance Monitor Start-Up Screen

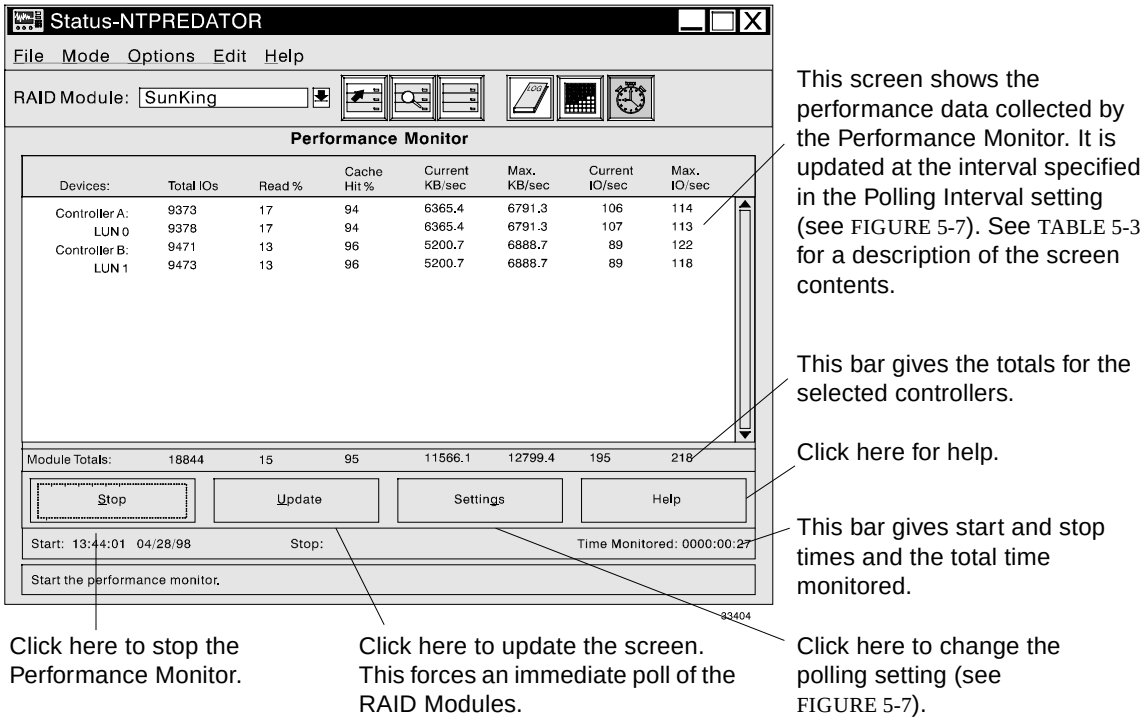

This screen shows the performance data collected by the Performance Monitor. It is updated at the interval specified in the Polling Interval setting (see FIGURE 5-7). See TABLE 5-3 for a description of the screen contents.

This bar gives the totals for the selected controllers.

Click here for help.

This bar gives start and stop times and the total time monitored.

Click here to change the polling setting (see FIGURE 5-7).

Click here to stop the Performance Monitor.

Click here to update the screen. This forces an immediate poll of the RAID Modules.

FIGURE 5-9 Performance Monitor Data Screen

TABLE 5-3 Performance Monitor Screen Contents

Column	Explanation
Devices	Controller name (A or B) or logical units (LUN 0, LUN 1, and so on). The LUNs are listed in numerical order following the controller that owns them.
Total IOs	The number of total I/Os performed by this device since the "Start" button of the performance monitor was last clicked. When this number exceeds 99,999, it is displayed as 100K until the number reaches 9999K, at which time it is displayed as "M" units. For amounts less than 100M, the value is displayed in tenths (for example, 12.3M).
Read%	The percentage of Total I/Os (previous column) that are read operations for this device. Write percentage can be calculated as 100 minus this value.
Cache Hit%	The percentage of reads that are fulfilled by data from the cache rather than requiring an actual read from disk.
Current KB/sec	The current KB per second. *Current* means the number of KB/second since the last time the polling interval elapsed, causing an update to occur, or since the update button was pressed.
Max KB/sec	The highest value achieved by the current KB/sec column for this polling session (since the "Start" button was pressed).
Current IO/sec	The current I/Os per second. Current means the number of I/Os/second since the last time the polling interval elapsed, causing an update to occur, or since the update button was pressed.
Max IO/sec	The highest value achieved by the current I/Os/second column for this polling session (since the "Start" button was last pressed).

Controller values represent the totals for all LUNs owned by the controller, not just the LUNs presently selected for monitoring. Module totals represent totals for both controllers in the controller pair, regardless if one, both, or neither are selected for monitoring.

Using Performance Monitor Data

The following considerations apply in analyzing the data gathered by the Performance Monitor.

Total I/Os – This field is useful for monitoring the I/O traffic to a specific controller and a specific LUN. This is needed to identify possible I/O "hot spots." It is also useful to be able to identify actual I/O patterns to the individual LUNs and compare those with the expectations based on the application. If a particular LUN has considerably more I/O activity, consider moving the LUN to the other controller in the RAID Module. If a particular array controller has considerably more I/O activity, consider moving it to another SCSI bus or system bus. The system administrator should also consider moving the storage set to another array controller or distributing the storage set across multiple arrays and LUNs.

Read% – Knowing the read/write ratio of an application gives more insight into the actual application behavior. With redundant RAID levels, write performance is typically slower than read performance, and RAID 5 is slower than RAID 1. If a high amount of write activity is occurring, the user might want to consider changing the RAID level from RAID 5 to RAID 1.

Cache Hit% – Indicates the percentage of read operations that were satisfied from the array cache (did not require a disk access). A higher percentage is better. Increasing the amount of installed cache increases the possibility of a cache hit, but there is no guarantee that more cache will improve cache hit percentage. This field can also be used to determine the effectiveness of read request prefetch. By default, the controller does not perform any read prefetching. If read prefetch is enabled, you can measure the effectiveness by noticing changes in the cache hit percentage.

KB/Sec Rates – The I/O size and the I/O rate determine the transfer rates of the controller. In general, a small host I/O request size results in lower throughput rates (but provides faster I/O rates and shorter response times). With larger host I/O request sizes, higher throughput rates are possible. Wide Ultra SCSI supports sustained data transfer rates of up to 40MB/second with large host I/O sizes. The maximum transfer rate with relatively small I/O sizes, 4096 bytes, is approximately 17MB/second. So, if the transfer rate is low, it may not be a problem.

IO/Sec Rates – Indicates current and maximum I/O rates. This can be useful in determining if the array is performing up to expectations based on access pattern, I/O size, RAID level, and number of drives in the array. Sequential access to the array results in higher I/O rates. Higher write I/O rates are experienced with write-back caching enabled compared to disabled. The higher the cache hit rate, the higher I/O rates will be.

CHAPTER 6

Recovery

This chapter describes how to recover from component failures.

Recovery Overview

Recovery involves replacing a failed RAID Module component, such as a drive, controller, fan, or power supply. Because most RAID Modules use redundant components, a single failure does not stop the RAID Module from operating. For example, if a single drive in a RAID 5 logical unit fails, the controller uses the data and parity information from the remaining good drives to continue delivering data.

Check the hardware guides accompanying your RAID Module for the redundant features on that module. In most cases, the module has redundant controllers, fans, and power supplies.

Tip – You may need to refer to the hardware guides accompanying your RAID Modules for information on how to physically replace failed components. The information in this *User's Guide* provides information on what you must do with the software to enable component replacement.

If a failure occurs in a redundant component, in most cases you can replace it without shutting down the RAID Module or stopping I/O to it. For example, if your RAID Module has redundant fans or power supplies, you can replace one of these components if it fails without shutting down the RAID Module.

For a discussion of the general principles involved in recovering from various kinds of failures, see:

- Viewing Recovery Procedures—page 118
- General Recovery Steps—page 119
- Recovering from Drive Failures—page 120
- Recovering from Controller/Connection Failures—page 125
- Recovering from Fan or Power Supply Failures—page 127
- Recovering from Battery Failures—page 127

Viewing Recovery Procedures

You can view the recovery procedures in the Recovery Guru by selecting Help → Recovery Guru Procedures in the Recovery application.

General Recovery Steps

The recovery process involves these general steps:

1. **You are notified of a component failure through one of the following:**
   - Software error messages
   - An error message in the Message Log
   - Fault lights or audible alarms on your RAID Module
   - SNMP notification
   - Console messages from your operating system
   - Results of a Recovery Guru / Health Check (or `healthck` command-line utility)
   - A non-optimal status appearing for any RAID Module component in the Configuration screen or in Module Profile
2. **You run Recovery Guru/Health Check to get information on the failure (see "Running Recovery Guru/Health Check" on page 128).**
3. **You should view the recommended procedure in Recovery Guru/Health Check for detailed information on the steps you need to take to recover from the failure (see** FIGURE 6-2**).**
4. **You take the steps indicated, which often involve replacing the failed component.**

   In some cases, you may need to refer to your hardware guides for information on how to physically replace failed components. You may also need to perform the functions given in "Manual Recovery Options" on page 131, but do not do so unless instructed by Recovery Guru.

Note – Unlike previous versions of the storage management software, in this version Recovery Guru does not initiate any commands or perform any operations automatically. You have complete control over the recovery process. Some operations (such as automatic reconstruction when you replace a failed drive in a redundant LUN), still occur automatically, but they are not initiated until you perform the tasks necessary to start them.

---

5. **You perform the functions that are necessary on your operating system to recover from the failure (see your operating system documentation and the *Sun StorEdge RAID Manager Installation and Support Guide*).**

   End of Procedure

Recovering from Drive Failures

Caution – This section gives the general principles involved in replacing failed drives. However, you should *always* start with the information given in the procedure screen of Recovery Guru. The information given there is specific to the type of failure that occurred and should be followed as much as possible.

Degraded and Dead Modes

Drive failures leave the logical unit in one of two general states, degraded or dead.

- Degraded Mode – Degraded mode occurs when a RAID 1, 3, or 5 logical unit suffers a single drive failure (or when multiple drives in a RAID 1 logical unit fail, but none of the drives make up a mirrored pair). In degraded mode, the logical units continue to operate normally.

  It is important to note that your operating system will *not* tell you when logical units enter degraded mode. This is because, the operating system, the drive is still functioning normally.

- Dead Mode – Dead mode occurs when any drives in a RAID 0 logical unit fail, or when multiple drives in a RAID 1, 3, or 5 logical unit fail. Your operating system handles a dead logical unit in the same way that it handles any failed drive. You will see the same error messages, for example, that you would see if a standard hard drive failed, and will need to follow the same recovery procedures that you take if a standard drive fails. See your operating system documentation for information on how to handle logical unit failures at the operating system level.

Recovery and RAID Level

The recovery procedure you need to perform depends on the RAID level of the logical units on that drive and on the number of drive failures. TABLE 6-1 summarizes the recovery procedures required for each RAID level. Note that in the case of single drive failures on RAID 1, 3, or 5 logical units, you do not need to shut down the RAID Module to recover from the failure. TABLE 6-1 gives general information only; *always* run Recovery Guru / Health Check before replacing drives.

TABLE 6-1 Drive Failures and Recovery (Summary)

RAID Level	Number of Failed Drives	Explanation
0	1 or more	A single drive failure in a RAID 0 logical unit results in the loss of all data on the logical unit, because RAID 0 does not provide data redundancy. To recover from a drive failure in a RAID 0 logical unit, you need to replace the failed drive, reformat the logical unit, and restore data from backup.
1, 3, or 5	1	A single drive failure in a RAID 1, 3, or 5 logical unit does not result in the loss of data. The logical unit continues to operate in degraded mode. To recover from a single drive failure in a RAID 1, 3, or 5 logical unit, you need to replace the failed drive. The controller then reconstructs the data on the replaced drive, and the logical unit returns to optimal.
1	2 or more	In a RAID 1 logical unit, multiple drives can fail without data loss so long as none of the failed drives make up a mirrored pair. • If two mirrored drives fail, all data in the logical unit is lost. To recover from this kind of multiple failure, you need to replace the drives, reformat the logical unit, and restore data from backup. • If none of the failed drives are in the same mirrored pair, data is not lost. To recover, you need to replace the failed drives. The controller then reconstructs the data on the replaced drives, and the logical unit returns to optimal.
3 or 5	2 or more	Two or more drive failures in a RAID 3 or 5 logical unit results in the loss of all data on the logical unit. To recover from multiple drive failures in a RAID 3 or 5 logical unit, you need to replace the failed drives, reformat the logical unit, and restore data from backup.

The information in this table ignores the effect of hot spare drives. Hot spare drives increase the number of drives that must fail in a RAID 1, 3, or 5 logical unit before data is lost (assuming sufficient time between failures). They do this by automatically replacing the failed drive. For example, if your RAID Module has 1 hot spare drive, *three* drives must fail before data is lost in a RAID 5 logical unit. See “Hot Spares” on page 13 and “Protecting Yourself From Component Failures” on page 25 for more information.

Multiple Drive Failures and Channel Failures

If multiple drives on the same drive channel fail simultaneously, the problem may be a channel failure rather than a series of drive failures. For example, if drives [0,1], [0,3], and [0,7] fail at the same time, the problem may be in channel 0 rather than in the three drives.

To prevent channel failures from causing multiple drive failures in the same logical unit, you should try to create logical units so that drives are not on the same channel, although this is not always possible.

If the drives fail as the result of a channel failure, data on the drives may not be lost. You need to fix the channel and then try to revive the logical unit using manual recovery (see “Manual Recovery Options” on page 131). As always, however, follow the steps given in the Recovery Guru procedure screen before attempting any kind of manual recovery.

Reviving Drives

Under limited circumstances, you can “revive” a failed drive and return it to an optimal status using the Revive Drive function in Manual Recovery. In general, you can only revive a drive if you accidentally removed it or failed it, not if the controller failed it. See “Manual Recovery / Drives” on page 132 for a description of the Revive Drive procedure.

Revive a drive only if:

- You are instructed to do so by Recovery Guru.
- You accidentally failed an *optimal* drive in an *optimal* RAID 0 logical unit (with Manual Recovery → Fail Drive) and have not performed any writes to it in the meantime. In this case, the logical unit should return to an optimal status after you revive it and the data is not affected.
- You accidentally failed an *optimal* drive in an *optimal* RAID 1, 3, or 5 logical unit (with Manual Recovery → Fail Drive) and have not performed any writes to it in the meantime. In this case, the logical unit should return to an optimal status after you revive it and the data is not affected. If you are not sure whether any writes occurred, reconstruct the drive instead of reviving it.
- You accidentally removed or failed an *optimal* drive in an *optimal* RAID 0 logical unit and you were writing to it at the time (the logical unit status changes to Dead, and the operating system sends you messages indicating the failure). In this case, reviving the drive may allow you to recover some of your data intact (any files that were not being written when you removed or failed the drive are probably OK).
- You accidentally removed or failed an *optimal* drive in an *degraded* RAID 1, 3, or 5 logical unit and you were writing to it at the time (the logical unit status changes to Dead, and the operating system sends you messages indicating the failure). In

this case, reviving the drive may allow you to recover some of your data intact (any files that were not being written when you removed or failed the drive are probably OK).

You do not need to revive a drive if:

- You accidentally remove an *optimal* drive from an *optimal* logical unit and the logical unit status remains optimal (that is, you were not writing to it at the time you removed it). The drive automatically returns to an optimal status after you reinsert it.
- You accidentally remove an *optimal* drive from an *optimal* RAID 1, 3, or 5 logical unit and the logical unit status changes to degraded (that is, you were writing to the logical unit at the time you removed it). The drive is automatically reconstructed after you reinsert it. (If reconstruction does not start automatically, use Manual Recovery → Drives → Reconstruct to start reconstruction.)

Tip – Do *not* revive a drive if the controller failed the drive.

Monitoring Reconstruction Progress

After you replace a single drive in a RAID 1, 3, or 5 logical unit, the controller starts reconstruction to rebuild the data on the replaced drive. See “Reconstruction” on page 17 for more information on reconstruction.

You can monitor the progress of reconstruction and change the reconstruction rate by using the Status application. See “Monitoring and Changing Reconstruction” on page 138 for more information.

Using Configured Drives as Spares

Note – This section applies to RAID Modules with firmware levels 3.0 or greater.

Because of the way controllers with firmware levels of 3.0 or greater write logical unit (LUN) information on the drives, you may run into trouble if you use previously configured drives as spares. For example, if you take a drive from one RAID module and use it as a spare in another RAID module, and that drive was part of a LUN when you removed it, the drive will have a status of Offline after you install it in the new module.

If you want to use configured drives as spares, always delete the LUNs on them before removing them.

If this is not possible, you can used the configured drives as spares, but you should insert them into the new module one at a time. When the drive status goes to Offline, you can then revive the drive using the procedure in "Reviving Drives" on page 122 *before* inserting the next drive. Note that this procedure is possible only if the drives were part of a LUN that contained at least two drives.

SMART PFA Reporting

RAID Modules with controller firmware levels of 3.01 or later support SMART (Self Monitoring And Reporting Technology). The controller compiles statistics from the drives and, based on a set of conditions, marks a drive if the results indicate that the drive is likely to fail. This is called setting a PFA (Predictive Failure Analysis) flag on the drive.

If a drive is marked with a PFA flag, a notification is sent to the message log *only* (as a hardware message type). The status of the drive does not change, the drive fault light does not come on, and Recovery Guru/Health Check does not pick up the condition. This means that you will be notified directly only if you have set up an SNMP Trap or some other means of notification.

If you find a drive marked with a PFA flag, see the Recovery Guru Browser for the procedure to use to fix it. In general:

- If a hot spare is currently being reconstructed on the logical unit affected by the problem drive, wait for reconstruction to finish before performing any recovery action.
- If the drive is in an optimal RAID 1, 3, or 5 LUN, you can wait for the drive to fail and then fix it, with no risk of data loss.
- If the drive is in an optimal RAID 0 LUN, you should repair the problem immediately, because all data will be lost if the drive fails. This will require you to stop all I/O to the LUN and copy data from backup.
- If the drive is in a degraded RAID 1, determine if the problem drive is the mirrored pair of the drive that already failed (use the List/Locate Drives option; see "Locating Drive Groups and Drives" on page 45).
  - If the drive is not the mirrored pair of a failed drive, you can wait for the drive to fail and then fix it, at no risk of losing data.
  - If the drive is the mirrored pair of the failed drive, you should repair the problem immediately, because all data will be lost if the drive fails. This requires you to stop all I/O to the LUN and copy data from backup.
- If the drive is in a degraded RAID 3 or 5 LUN, you should repair the problem immediately, because all data will be lost of the drive fails. This requires you to stop all I/O to the LUN and copy data from backup.

Recovering from Controller / Connection Failures

Caution – This section gives the general principles involved in replacing failed controllers and recovering from connection failures. However, you should *always* start with the information given in the procedure screen of Recovery Guru. The information given there is specific to the type of failure that occurred, and should be followed as much as possible.

A hardware failure in a controller or along the connection from the host to the controller usually does not result in the loss of data. However, it will result in the RAID Module becoming unusable unless *all* of the following conditions are met:

- The RAID Module has redundant controllers (see “Redundant Controllers” on page 17).
- The RAID Module is configured in the single host configuration using two host busses (see “RAID Module Configurations” on page 19).
- The RAID Module has RDAC protection (see “RDAC Driver” on page 16).

Always consult your hardware guides for instructions on how to physically replace the controller, as well as any possible instructions on additional steps you must take.

Recovery and Fail-Over Protection

TABLE 6-2 summarizes the steps you need to take to recover from a failed controller depending on whether or not your RAID Module has fail-over protection.

TABLE 6-2 Recovering from Failed Controllers (Summary)

Fail-Over Protection?	Recovery Procedure
Yes	If a controller fails, the other controller in the pair takes over for it. LUNs assigned to the failed controller are automatically assigned to the remaining good controller. Data continues to flow along the connection. If you are using write caching but not write cache mirroring, any data in the cache of the failed controller is lost. **Recovery**: You will probably be able to replace the failed controller without shutting down the RAID Module. Follow the instructions in the Recovery Guru procedure screen. In general, you must manually fail the failed controller, remove it, and insert a new one.
No	If the RAID Module does not have fail-over protection, data will not flow along the failed connection, and any data currently in the cache of the failed controller is lost. Data on the drives is not necessarily lost, but is inaccessible. **Recovery**: You will need to shut down the RAID Module to replace the controller. Follow the instructions on the procedure screen.

Controller Replacement and Firmware (Autosynch)

Note – At the time this manual was published, this feature was not supported by Sun.

Note – This feature applies only to controllers with firmware levels of 3.0 or greater, and only if certain NVSRAM bits are set. See the *Sun StorEdge RAID Manager Installation and Support Guide* for information about these bit settings on your system.

When you replace a failed controller, the firmware on the replaced controller is automatically synchronized with the firmware on the good controller. This ensures that the controller firmware level is the same on both controllers after a replacement. You do not need to download new controller firmware after you replace the controller.

For example, assume you have two controllers (controllers A and B) with a firmware level of 3.0. One controller (controller A) fails. You replace it with a new controller with a firmware level of 3.1. After you bring the replaced controller online, the firmware on the replaced controller (controller A) is replaced with the 3.0 controller firmware from controller B.

If the *replacement* controller contains firmware of 3.1 or higher, the RAID Module controller fault light blinks while the firmware is being synchronized (even if the controller is being synchronized to firmware level 3.0). The LEDs on the controller blink in a moving pattern indicating a serial download. Do *not* remove the controller while the autosynch is occurring.

Note that if either controller has a firmware level less than 3.0, the firmware will not be copied, and you will need to download new firmware to the new controller to make it match the firmware on the good controller.

Two NVSRAM bits must be set for this feature to work. See the *Sun StorEdge RAID Manager Installation and Support Guide* for the settings on your system.

Controller Replacement and Cache Parameters

After you replace a failed controller, it may take as much as 15 to 20 minutes before the controller re-enables any cache parameters. You can use the RAID Module during this time, but without caching. You can tell if cache parameters are enabled by looking for asterisks in the Change Cache Parameters screen (see “Changing Cache Parameters” on page 167).

Recovering from Fan or Power Supply Failures

Caution – This section gives the general principles involved in replacing failed RAID Module fans and power supplies. However, you should *always* start with the information given in the procedure screen of Recovery Guru. The information given there is specific to the type of failure that occurred, and should be followed as much as possible.

In most cases, the RAID Module contains redundant fans and power supplies. A single failure, therefore, does not shut down the RAID Module. You can then replace the failed component while the RAID Module is running. The replaced component is automatically returned to service in the RAID Module.

If necessary, see your hardware guides for information on how to physically replace failed fans and power supplies, as well as information on how to locate the particular failed fan or power supply.

Recovering from Battery Failures

Most RAID Modules have a battery unit to preserve cache memory in the event of a power failure. This battery has a limited life-span (usually two years). Recovery Guru / Health Check generates an error message if the battery has expired or if it is nearing its expiration date. See your hardware guides for more information about the battery unit.

To replace the battery, use the recovery procedure in Recovery Guru. After you replace the battery, you must reset the battery date in the storage management software (otherwise, you will continue to get error messages).

To reset the battery date, type the following from the command line:

```
raidutil -c<controller_id> -R
```

where <*controller_id*> is the name of the controller contained in the RAID Module you put the new battery in. If there are two controllers in the RAID Module, you must perform this command on each of them.

See the Help files for the `raidutil` command if you need more information.

Running Recovery Guru / Health Check

Run Recovery Guru / Health Check when you have some indication that there is a failure in some array component or if you want to check the status of your modules. Health Check locates the failure and gives detailed information on the steps you need to take to correct it.

To run Recovery Guru / Health Check, perform the following procedure.

1. **Start the Recovery application.**
2. **Select the RAID Modules you want to check, or select All RAID Modules to check all the modules on your system.**

   A screen like FIGURE 6-1 is displayed.

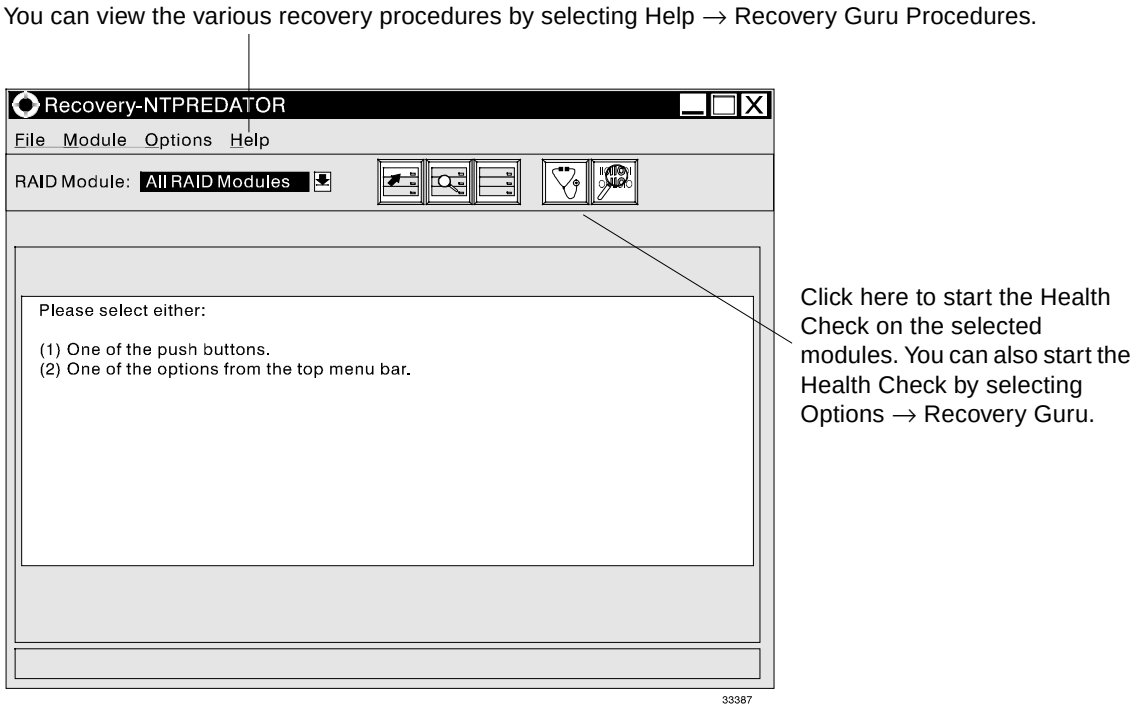

FIGURE 6-1 Recovery Main Screen

3. **Click on the Recovery Guru/Health Check button or select Options → Recovery Guru from the drop-down menus.**

   The software checks all the selected modules for non-optimal statuses. If a non-optimal status is found, a screen like FIGURE 6-2 is displayed.

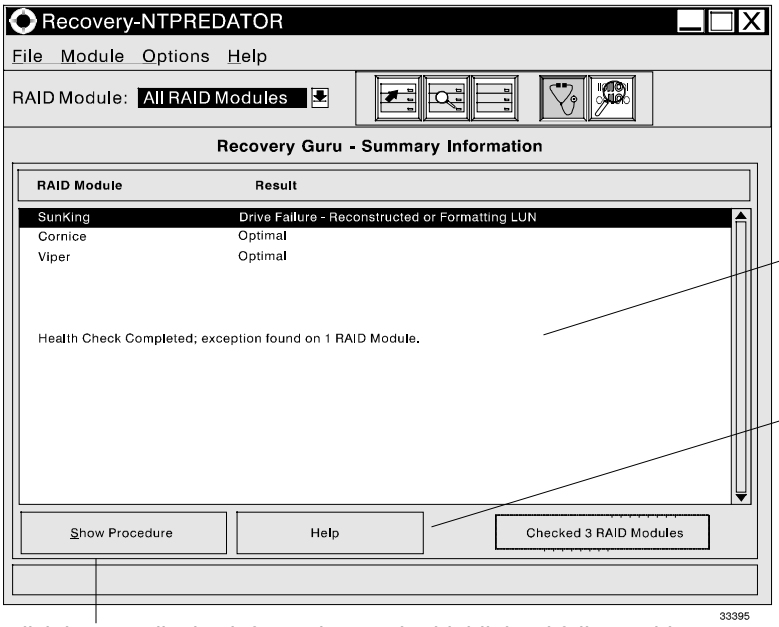

This area displays the results of the check. If a non-optimal status is found, highlight the line and click Show Procedure for recovery information.

Click here for help.

Click here to display information on the highlighted failure. This information will include recovery information. See FIGURE 6-3 for a sample screen.

FIGURE 6-2 Recovery Screen Showing Exceptions Found

4. **If exceptions are found, highlight the exception and click on Show Procedures for recommendations on how to fix the problem (see** FIGURE 6-3**).**

5. **Print out the procedure, if desired, using File → Print Topic.**

6. **Follow the procedure given to recover from the component failure. You may need to use the procedures given in "Manual Recovery Options" on page 131.**

End of Procedure

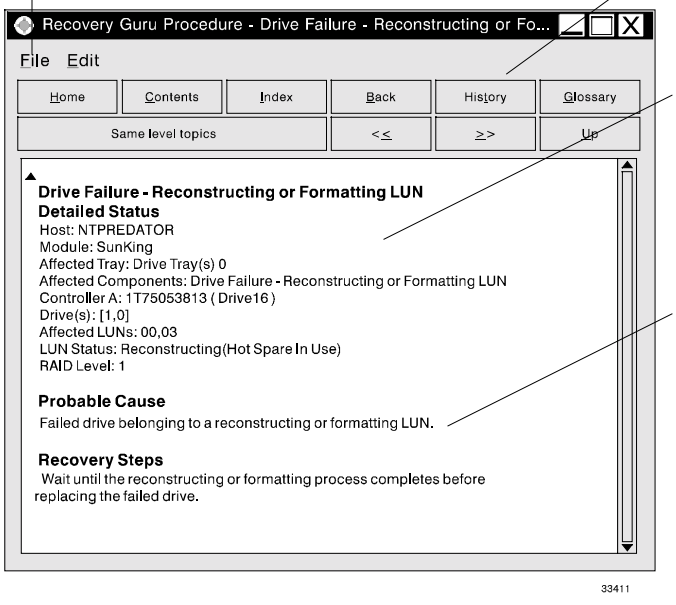

FIGURE 6-3 Recovery Procedure Screen

Manual Recovery Options

Caution – Take care when using any of these options. In general, do not use these options unless specifically directed by the procedure screen of Recovery Guru or by your local solution center. Using this option under other circumstances could result in the loss or corruption of data. Do not attempt manual recovery without understanding the circumstances of a failure.

The Manual Recovery options provide a means to perform recovery operations on drives, logical units, and controllers.

Use the Manual Recovery options if:

- You are instructed to in the procedure screen of Recovery Guru.
- You need to format a logical unit after replacing the drives in a dead logical unit.
- Reconstruction does not start automatically after you replace a drive.
- You accidentally remove a good drive from a logical unit and no write operation was performed on the logical unit while the drive was removed (see “Reviving Drives” on page 122).
- You want to force a drive failure rather than wait for the controller to fail it.
- You want to remove a controller from the RAID Module (place the controller offline before you remove it and then place it online after you replace it).
- A channel failure caused multiple drives in a logical unit to fail.

There are three Manual Recovery options:

- Manual recovery of drives – see “Manual Recovery / Drives” on page 132.
- Manual recovery of logical units – see “Manual Recovery / Logical Units” on page 134.
- Manual recovery of controllers – see “Manual Recovery / Controllers” on page 136.

Manual Recovery / Drives

Caution – Take care when using manual recovery. In general, do not use these options unless specifically directed by the procedure screen of Recovery Guru or by your local solution center. Using this option under other circumstances could result in the loss or corruption of data. Do not attempt to manually recover without understanding the circumstances of the failure.

Using manual recovery on drives, you can:

- Fail a drive – you may need to do this if the drive is causing problems but the controller won't fail it.
- Reconstruct data on a drive – you may need to do this if reconstruction doesn't occur automatically when you replace a failed drive.
- Revive a drive – you may need to do this if you accidentally fail or remove a drive.

Caution – Before attempting to revive a drive, see "Reviving Drives" on page 122. Reviving a drive under the wrong circumstances can result in corrupted data or system problems. Reviving a drive is not a substitute for reconstructing the drive.

To manually recover drives, perform the following procedure.

1. **Start the Recovery application.**
2. **Select the RAID Module containing the drives you want to recover.**
3. **Select Options → Manual Recovery → Drives from the drop-down menus.**

   A screen like FIGURE 6-4 is displayed.
4. **Select the individual drive you want to fail, reconstruct, or revive, and then click on the appropriate button to perform the operation.**

   End of Procedure

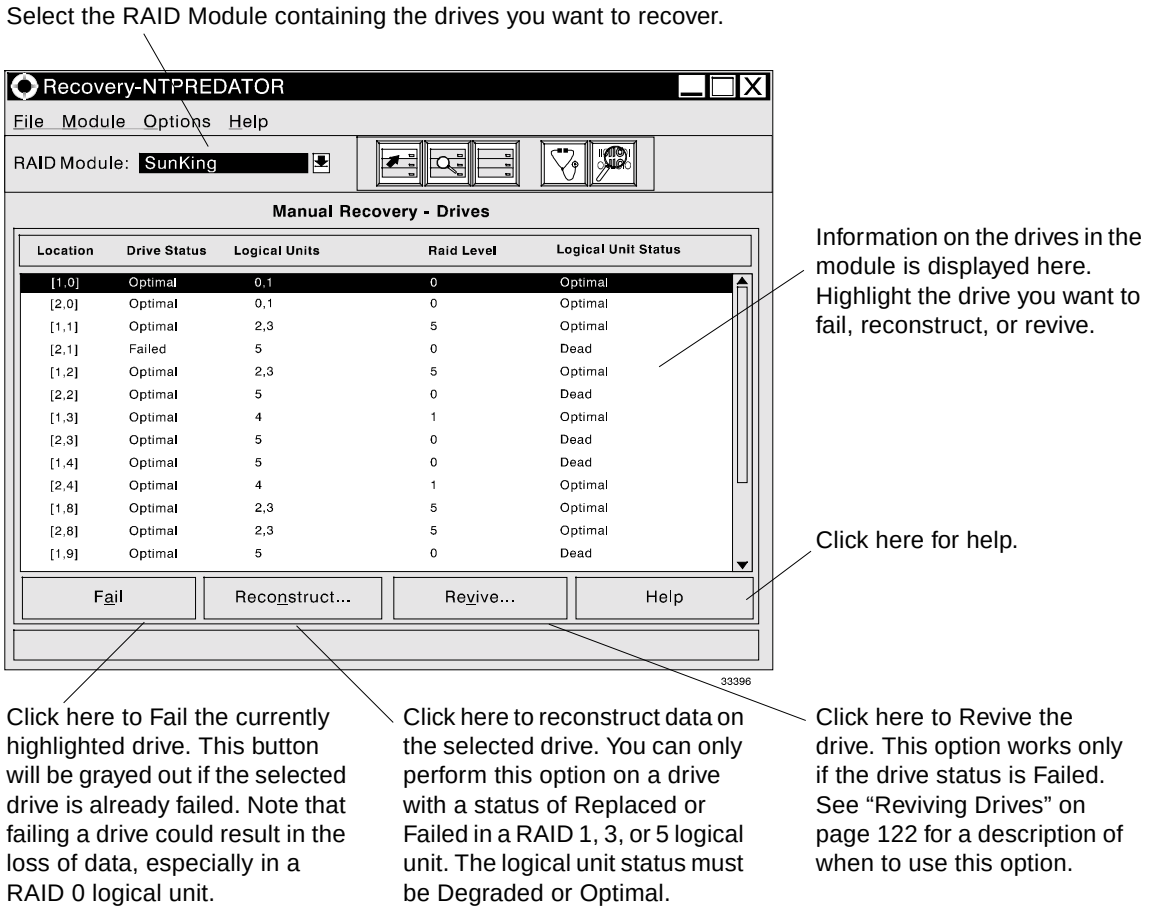

FIGURE 6-4 Manual Recovery / Drives Screen

Failing a drive in a RAID 0 logical unit results in the loss of data on the logical unit.

Failing a drive in a degraded RAID 3 or 5 logical unit results in the loss of data.

Failing a drive in a degraded RAID 1 logical unit results in the loss of data if the drive you fail is the mirror drive of a failed drive.

Reconstruction should only be performed on a drive after you have replaced it.

Reviving a drive should be done only under the circumstances described in “Reviving Drives” on page 122.

Manual Recovery / Logical Units

Caution – Take care when using manual recovery. In general, do not use these options unless specifically directed by the procedure screen of Recovery Guru or by your local solution center. Using this option under other circumstances could result in the loss or corruption of data. Do not attempt to manually recover without understanding the circumstances of the failure.

Using Manual Recovery on logical units, you can:

- Format the logical unit – you need to do this to recover from a dead logical unit after you have replaced the failed drives.
- Revive a logical unit – you may need to do this if a logical unit failed because of a channel failure. A channel failure can fail multiple drives in the logical unit but not cause the data on those drives to be lost. Do not use this option unless instructed to by Recovery Guru.

To manually recover logical units, perform the following procedure.

1. **Start the Recovery application.**
2. **Select the RAID Module containing the logical units you want to format or revive.**
3. **Select Options → Manual Recovery → Logical Units from the drop-down menus.**

   A screen like FIGURE 6-5 is displayed.
4. **Select the logical units you want to format or revive, then click on the appropriate button to perform the operation.**

   End of Procedure

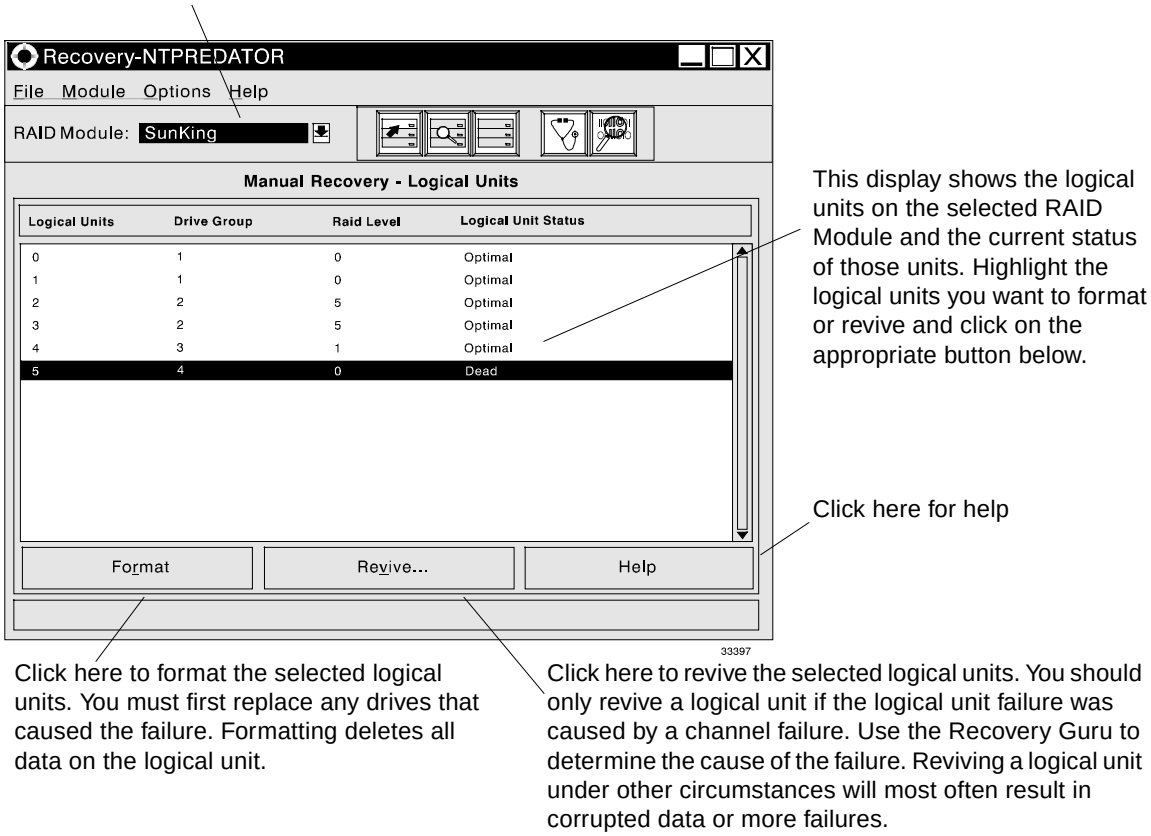

FIGURE 6-5 Manual Recovery / Logical Units Screen

Do not format a logical unit until you have replaced any failed drives that caused the failure.

Formatting deletes all data on the logical unit (it is probably already lost anyway, assuming the logical unit status is Dead).

Reviving a logical unit should be performed only if the logical unit failure was caused by a channel failure. Do not attempt to revive a logical unit unless instructed to by Recovery Guru. Doing otherwise can result in corrupted data or another failure. See “Multiple Drive Failures and Channel Failures” on page 122 for more information.

Manual Recovery / Controllers

Caution – Take care when using manual recovery. In general, do not use these options unless specifically directed by the procedure screen of Recovery Guru or by your local solution center. Using this option under other circumstances could result in the loss or corruption of data. Do not attempt to manually recover without understanding the circumstances of the failure.

Using Manual Recovery on controllers, you can:

- Place a controller offline – Use this option if you want to remove the controller in order to replace it.
- Place a controller online – Use this option if you have replaced a failed controller and it did not go back online automatically.

To manually recover controllers, perform the following procedure.

1. **Start the Recovery application.**
2. **Select the RAID Module containing the controllers you want to place online or offline.**
3. **Select Options → Manual Recovery → Controller Pairs from the drop-down menus.**

   A screen like FIGURE 6-6 is displayed.
4. **Select the controller you want to place online or offline and then click on the appropriate button to perform the operation.**

   End of Procedure

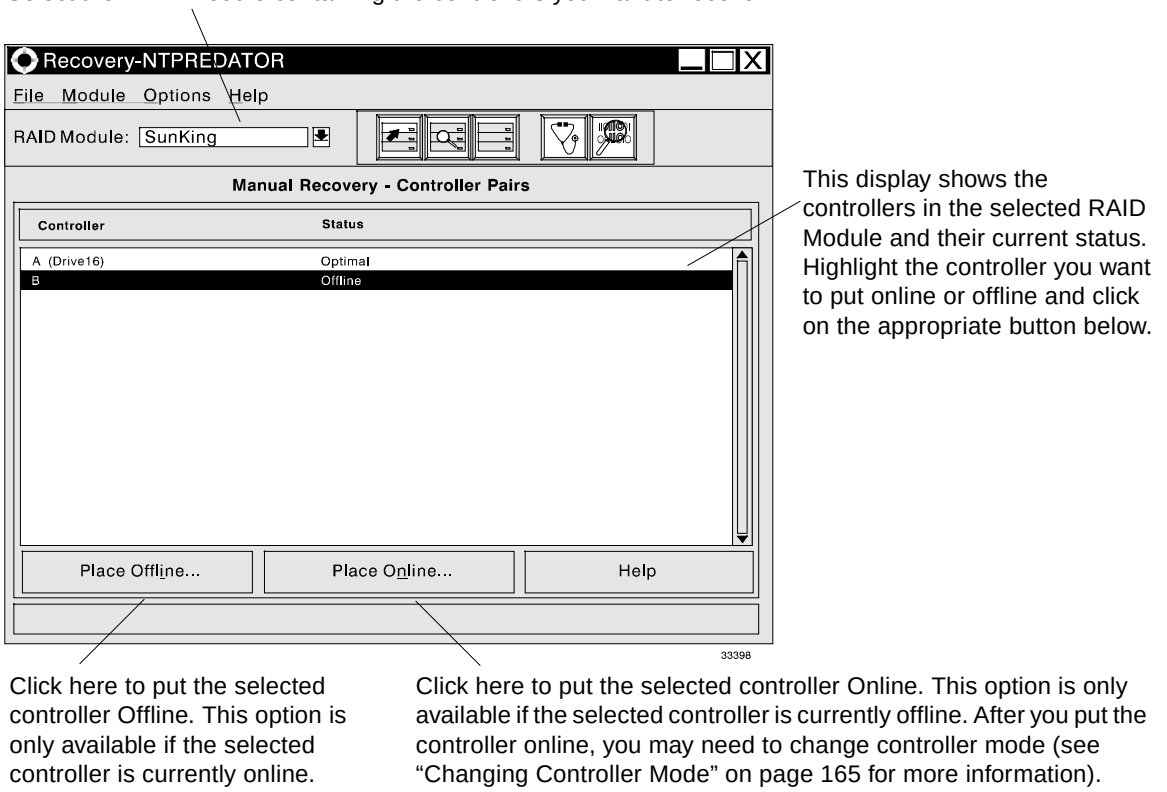

FIGURE 6-6 Manual Recovery / Controllers Screen

You can only place a controller offline if there are two controllers in the RAID Modules.

If the controller has RDAC failover support, the other controller takes over for the offline controller (see “Redundant Controllers” on page 17).

If you place an offline controller online, you may need to go into the Maintenance and Tuning application to reset the controller modes and assign LUNs to the online controller (see “Changing Controller Mode” on page 165).

Monitoring and Changing Reconstruction

After you replace a failed drive in a degraded RAID 1, 3, or 5 logical unit, the controller begins rebuilding data on the replaced drive. You can use this option to view the progress of this reconstruction, and to change the reconstruction rate.

You can only change the rate for logical units currently being reconstructed. To change the reconstruction rate of logical units that are not currently being reconstructed, see “Changing the Reconstruction Rate” on page 158.

To view reconstruction progress and change the rate of reconstruction, perform the following procedure.

1. **Start the Status application.**
2. **Select the RAID Module containing the logical units currently being reconstructed. Do not select All RAID Modules.**
3. **Click on the LUN Reconstruction button or select Options → LUN Reconstruction from the drop-down menus.**

   A screen like FIGURE 6-7 is displayed. The screen shows the progress of reconstruction on each logical unit.
4. **If you want to change the reconstruction rate, move the slider bar for that logical unit.**

Tip – Reconstruction can take a long time (several hours or longer) if the logical units involved are very large and there is a lot of I/O going to the LUNs being reconstructed. To some extent, you can speed up this process by adjusting the reconstruction rate with the slider bars.

In addition, if you replace a failed drive while the hot spare that replaced the failed drive is being reconstructed, reconstruction is completed on the hot spare before it begins on the replaced drive. This may increase the time required to complete the recovery procedure.

End of Procedure

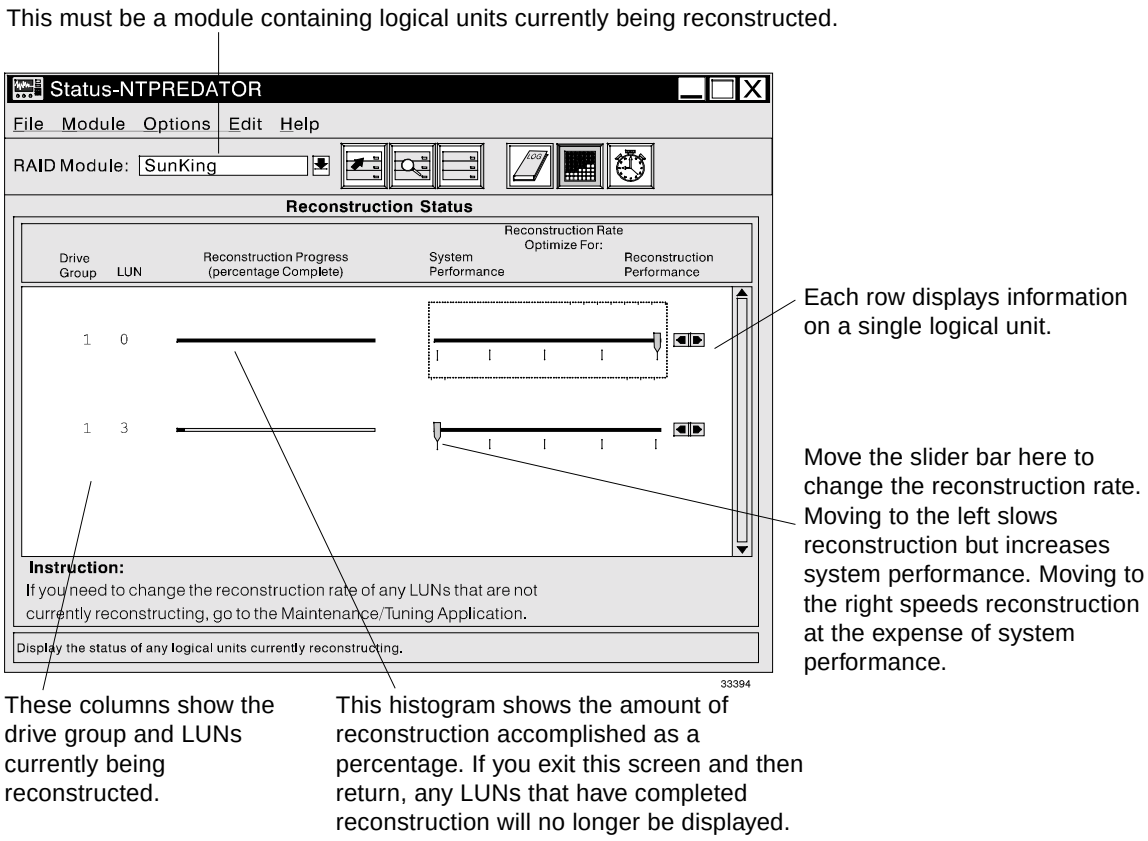

FIGURE 6-7 Viewing / Changing Reconstruction Screen

If you exit this screen and then return to it, any logical units that have completed reconstruction are no longer be displayed.

The reconstruction rate is changed immediately when you move the slider bar, although you may notice a slight delay if you have many or large logical units reconstructing.

From left to right, the points on the slider bar indicate the following reconstruction rates (blocks / seconds delay). Note that each 1024 block = 512 KB.

- Slow – 1024 blocks / 0.8 second delay
- Medium slow – 1024 blocks / 0.4 second delay
- Medium – 1024 blocks / 0.2 second delay
- Medium fast – 1024 blocks / 0.1 second delay
- Fast – Reconstruct with no delays

CHAPTER 7

Maintenance and Tuning

This chapter describes the maintenance and tuning procedures needed to maintain the RAID Modules.

Maintenance and Tuning Overview

The Maintenance and Tuning application performs functions to ensure that your RAID Modules operate efficiently. With the Maintenance and Tuning application, you can:

- Download new controller firmware and NVSRAM files (see “Upgrading Controller Firmware” on page 143).
- Manually run Parity Check (see “Running Parity Check” on page 153).
- Set the time for Automatic Parity Check (see “Running Parity Check” on page 153).
- Change the Reconstruction Rate for logical units (see “Changing the Reconstruction Rate” on page 158).
- Balance drive groups / logical units across active / active redundant controller pairs (see “Balancing Drive Groups / LUNs” on page 160).
- Swap active / passive controllers or change them to active / active (see “Changing Controller Mode” on page 165).
- Change the write cache, write cache mirroring, and cache without batteries parameters for logical units (see “Changing Cache Parameters” on page 167).

Upgrading Controller Firmware

Controller firmware resides on the controller in the RAID Module. It contains the commands necessary to operate the RAID Modules. Normally, you do not need to load new firmware when you install a RAID Module.You will receive new firmware files from your local solution center.

Note – Controller firmware is different from the drive firmware. If you need to upgrade drive firmware, contact your local Sun solution center or Sun service provider.

See the *Sun StorEdge RAID Manager Release Notes* and the *Sun StorEdge RAID Manager Installation and Support Guide* for information about the firmware supported by your RAID Modules and any special requirements for downloading new firmware.

Controller Firmware Files

Note – It is recommended that you contact your local Sun solution center or Sun service provider to find out if your controllers have pre-installed controller firmware files. For example, on some controllers the NVSRAM file may already be installed.

Upgrading controller firmware involves up to five files:

- The NVSRAM file – The NVSRAM file specifies default settings for the controller. If present, you must download this file *before* loading the other firmware files.
- The Boot Level (also called *bootware*) file – This is a firmware file that controls controller operation.
- The Firmware Level (also called *appware*) file – This is a firmware file that controls controller operation.

- The `fwcompat.def` file – This file enables the software to compare the firmware files for compatibility during the upgrade process, providing you with a list of compatible files to select for downloading.

You can determine current boot, fibre channel, and firmware version levels through the Controller Profile screen; see "The Controller Profile Screen" on page 50.

General Procedure

In general, use the following steps to upgrade new controller firmware.

1. **Refer to your *Sun StorEdge RAID Manager Installation and Support Guide* for important additional information on upgrading controller firmware on your RAID Modules.**
2. **Determine whether you can download the new firmware files using the Online or Offline method, and take any necessary steps required (see the next section).**

Note – If you do not copy the `fwcompat.def` file to the correct subdirectory, the software is unable to check the files for compatibility. Although you can still enter firmware file names, the software is unable to provide you with a list of compatible files to select for downloading.

---

3. **Copy the new NVSRAM file, the firmware files, and the `fwcompat.def` file to the default subdirectory in the installation directory of your system (see the *Sun StorEdge RAID Manager Installation and Support Guide* for information on which directory to use).**
4. **Download the new NVSRAM file using the Firmware Upgrade option in the Maintenance and Tuning application.**
5. **Run `nvutil -vf` from the command line to ensure that NVSRAM settings are correct for the storage management software.**
6. **Turn off the power to the RAID module and then turn it on again to establish the new NVSRAM settings.**
7. **Download the new firmware files to your system.**

Read through the following sections and then go to "Download Procedure" on page 147.

Online and Offline Method

You can download controller firmware using one of two methods: Online or Offline. The method you need to use depends on your controllers and configuration.

Note – Because you must turn the RAID Module off and then on again after upgrading NVSRAM, you must use the Offline method if a new NVSRAM file must be downloaded.

- Select Online to download firmware while the selected RAID Module receives I/O. You may only use the Online method if:
  - The controllers have the (RDAC) driver installed for redundant controller support.
  - All LUNs on the RAID Modules have a status of Optimal.
- Select Offline to download firmware when the selected RAID Module is not receiving I/O. You must use the offline method if:
  - You need to download a new NVSRAM file (you can load an NVSRAM file using the Online method, but the changes in the file do not take effect until after you turn the RAID Module off and then on again).
  - The RAID Module has an independent controller configuration (in this case, you must download the firmware to each controller from its host).
  - The RAID Modules have single controllers.
  - The Offline option also requires exclusive access to the logical units in the selected RAID Modules; that is, no other operations can be running on the RAID Module.

The software notifies you if you can perform a firmware upgrade using the selected method.

Selecting Controllers

Whether or not you can select specific controllers for downloading NVSRAM or upgrading controller firmware depends on the RAID Module you select:

- When you select `All RAID Modules` → `Firmware Upgrade`, you will download NVSRAM files or upgrade controller firmware files to *every* controller in those modules. You *cannot* select individual controllers in this case.
- If you select one RAID Module that has only one controller, you must use the Offline method. The controller is automatically selected in this case.
- If you select a RAID Module with an independent controller configuration, you must use the Offline method. In addition, you must download the firmware to each controller from each host.
- If you select one RAID Module with a pair of redundant controllers and the Offline method, you need to select the controllers on which you want to upgrade firmware in addition to highlighting the version level you want to download.

Caution – Remember that both controllers in a redundant pair must have the same version of controller firmware installed. You should therefore select both controllers to ensure that they have compatible versions of NVSRAM and controller firmware *unless* you are replacing a failed controller and the replacement controller has an earlier firmware version than the original pair was using. Note that this may be affected by the autosynch feature (see “Controller Replacement and Firmware (Autosynch)” on page 126).

Download Procedure

To download new controller firmware and NVSRAM files, perform the following procedure.

1. **Refer to the *Sun StorEdge RAID Manager Installation and Support Guide* for important additional information on downloading controller firmware to your RAID Modules.**
2. **Copy the NVSRAM, firmware, and the `fwcompat.def` files to the correct subdirectory in the installation directory of your system (see your *Sun StorEdge RAID Manager Installation and Support Guide*).**
3. **Determine whether you can download the firmware using the Online or Offline method (see "Online and Offline Method" on page 145).**

Tip – If you must use the Offline method, make sure that you stop all I/O to the RAID Modules you are going to upgrade.

---

4. **Start the Maintenance and Tuning application.**
5. **Select the RAID Module containing the controllers that you want to upgrade.**

   Select All RAID Modules to download firmware to all controllers in all RAID modules.
6. **Click on the Firmware Upgrade button or select Options → Firmware Upgrade from the drop-down menus.**
7. **Read the Important Notes screen and click OK.**

   A screen appears prompting you to select the Online or Offline procedure.
8. **Select either**
   - Online to download firmware while the RAID module processes I/O,

     or
   - Offline to download firmware while the RAID module is *not* processing I/O.

   After you select Online or Offline, the screen displays "Verifying the controller state" while the software checks the selected RAID Modules for restrictions based on the type of firmware download you selected. If there are no restrictions, a screen like FIGURE 7-1 is displayed.

This display shows the controllers in the selected RAID Modules and their current firmware levels. Make sure both controllers are highlighted here if you want to download new firmware to both controllers. Note that for firmware levels 3.0 and higher, the Fibre Channel level is the same as the firmware level because the Fibre Channel firmware is contained in the firmware file.

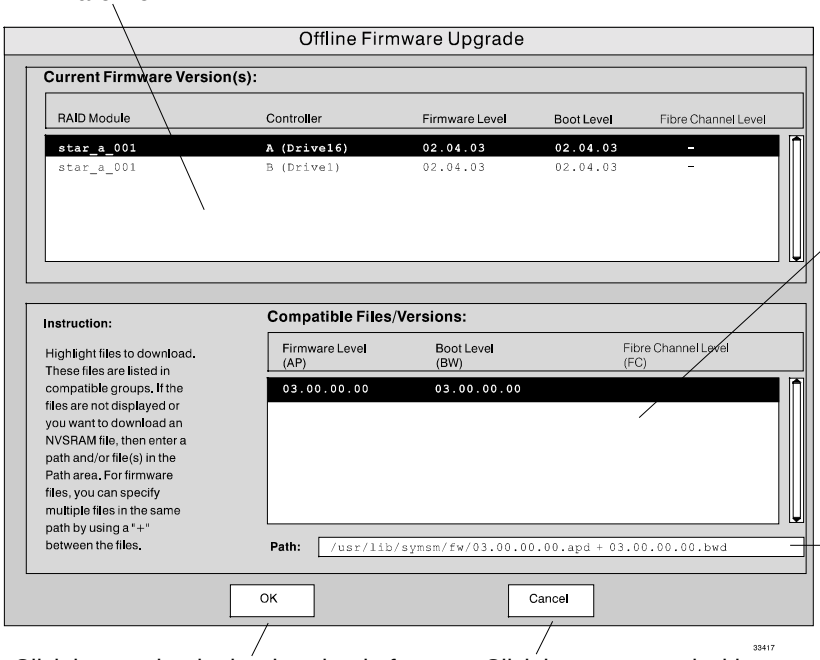

This screen displays the compatible firmware files found using the `fwcompat.def` file. Highlight the version level you want to download.

Note that for firmware versions 3.0 or higher, the firmware (*appware*) file contains the Fibre Channel firmware, so no Fibre Channel file is listed here.

The path is updated to show the files currently selected in the Compatibility display. You can also enter a file name here to download that file (this is how to download an NVSRAM file).

Click here to begin the download after you have selected the firmware files to load or entered the name of the NVSRAM file in the path box.

Click here to cancel without downloading any new files.

FIGURE 7-1 Firmware Upgrade Screen

9. **Do one of the following:**

Note – Remember, you need to download any NVSRAM files first, before downloading any new firmware files.

- To download an NVSRAM file, type its complete path information in the path box and select OK.

  You see an Upgrade Restriction screen indicating that the software does not recognize the file as a firmware file. Verify that the NVSRAM file name is entered correctly and select OK.
- To download firmware, highlight the version level you want and select OK.

After you select OK, you either receive notification that a problem occurred, or you have a final confirmation that the download process is about to begin.

10. **Select OK to start the download.**

Do not select any other options or exit from the Maintenance and Tuning application until the upgrade process is complete. You can, however, monitor the upgrade progress. A histogram for the selected RAID module indicates the download progress of the NVSRAM or firmware files. This graphic shows the amount of progress as a percentage and starts over at 0% for each file if you have more than one. If you selected All RAID Modules, the module number is updated as each module begins its upgrade process.

When the NVSRAM download or the firmware upgrade is finished, you see a summary report indicating whether the upgrade is Successful or Failed. Table 7-1 on page 150 shows the information this screen displays.

Note – If you selected All RAID Modules, it is possible that the upgrade was successful for some modules, but not for others. The summary report should indicate which modules were not successful and give an appropriate cause. For more information, see “Confirming the Download” on page 151.

TABLE 7-1 Firmware Confirmation Screen Elements

Screen Element	Description
Summary Report for Files	Lists the files used to upgrade the firmware. These are the files loaded in the Path line when you selected files on the Compatible Files / Versions screen (FIGURE 7-1).
RAID Module	Identifies the specific RAID Module.
Download Status	Indicates whether the download process was completed successfully. You see either "Successful" or "Failed" with a reason why the upgrade was unsuccessful. See TABLE 7-2 if you see any Failed download statuses.

11. **After the download is completed, select OK.**

    You return to the Maintenance and Tuning screen.

12. **Did you download an NVSRAM file or new firmware?**

New Firmware	NVSRAM File
You are finished with this procedure.	Continue with Step 13.

13. **At the command line, type:**

    `nvutil -vf`

    The `nvutil` utility checks and corrects any settings on all controllers in your RAID Modules to ensure that certain settings in the NVSRAM are set up correctly for this software. Consult the Help file or man page specific to your operating system for a description of the `nvutil` utility.

14. **Turn off the power to the RAID module controllers and then on again to implement the new NVSRAM settings.**

    Note – On SCSI RAID modules, you can reboot the system instead. Be sure to stop all I/O to the system before doing so.

15. **Go back to Step 9 on page 149 and download any new firmware files.**

    End of Procedure

Confirming the Download

At the final summary report of the NVSRAM or firmware download procedure, you will see if the upgrade was Successful or Failed for each of the selected RAID Modules.

- If you see that the upgrade was Successful, you should still verify that all of the logical units (LUNs) are *not* assigned to only one controller. See "Balancing Drive Groups/LUNs" on page 160 for the procedures to balance the LUNs across both controllers.
- If you see "Failed" for any module, you should fix the specified failure and try the firmware upgrade procedure again. See TABLE 7-2 for possible actions to take to correct a failed upgrade.

TABLE 7-2 Controller Firmware Upgrade: Failed Statuses

Reasons for Failed Status	Actions
The selected module had I/O activity occurring or file systems mounted (Offline download).	Stop I/O to that module and be sure file systems are unmounted, then try to upgrade the firmware again.
At least one of the selected firmware files had bad file contents.	Copy the firmware files to the correct subdirectory in the installation directory again. If you see this message a second time, one or more of your files is most likely corrupt. Obtain a new copy of the firmware upgrade files.
The SCSI command write buffer failed.	Try to perform the upgrade again for this module. If it fails a second time, contact your local solution center.
The software was unable to reset the controller.	Try to upgrade the firmware again.
One or more logical units for the selected module were not Optimal.	Use Recovery Guru/Health Check in the Recovery application to restore the LUNs to an Optimal status, then try to upgrade the firmware again. See Chapter 6.
Upgrading to the selected firmware version requires that you use the Offline method.	Try to upgrade the firmware again and this time be sure to select Offline.
The current firmware version is unable to upgrade to the files you selected.	Most likely, you need to upgrade to an intermediate version of firmware. Try to upgrade to a version earlier than the one you selected. If that upgrade is successful, perform a second upgrade for this latest firmware version. See your *Sun StorEdge RAID Manager Installation and Support Guide* for additional information on intermediate upgrades.

TABLE 7-2 Controller Firmware Upgrade: Failed Statuses *(Continued)*

Reasons for Failed Status	Actions
The files you selected are not compatible with the current firmware version on the selected module's controllers.	Most likely, the current directory does not contain *all* the necessary firmware files. Copy the firmware files and the `fwcompat.def` file to the correct subdirectory in the installation directory and try again. Be sure the version you select has both Firmware Level and Bootware Level versions specified. If the upgrade fails a second time, obtain a new copy of the firmware upgrade files.
The software was unable to access the controllers during the upgrade process.	Use Recovery Guru / Health Check in the Recovery application to determine if the module has a failure. See Chapter 6. • If a failure is indicated, fix it and try to upgrade the firmware again. • If Recovery Guru / Health Check does not indicate a failure, try to upgrade the firmware again.
You tried to load a pre-2.04 firmware version, which is not supported by this software or the redundant controller configuration.	Do *not* try to load this firmware version again.
The selected firmware files are not compatible with your controller model.	Use Module Profile → Controller details to check your controller type and model (see FIGURE 3-6), and obtain the correct firmware version files. See your *Sun StorEdge RAID Manager Installation and Support Guide* for additional information on downloading controller firmware.
The online upgrade cannot be performed because either the selected module has only one controller or one of the controllers in the pair is not accessible.	Use Module Profile → Controller details to determine how many controllers the module has (see FIGURE 3-6). • If there is only one controller, try to upgrade the firmware again and be sure to select Offline. • If you have two controllers, use the Recovery application to select Recovery Guru / Health Check and follow the recommended procedure to fix the controller problem before attempting to upgrade the firmware again (see Chapter 6).
An unknown failure occurred.	Use the Status application to select Message Log for component information (see “Viewing the Status Log” on page 104).

Running Parity Check

The RAID Modules use parity information in RAID 3 and 5 logical units to enable data redundancy. When a single drive fails in a RAID 3 or 5 logical unit, the controller can reconstruct the data on the missing drive by using the parity information stored on the drives. RAID 1 logical units do not use parity, but you can still run Parity Check/Repair on RAID 1 logical units. In this case, the check compares the data on the mirrored drives.

RAID 0 logical units do not have parity and therefore cannot be checked. A non-optimal RAID 1, 3, or 5 logical unit (such as a degraded units) also cannot be checked.

When you run Parity Check/Repair, you can specify whether you want the controller to automatically fix any parity errors found. If the error is found on a RAID 3 or 5 logical unit, the controller changes the parity information so that the parity is correct. If an error is found on a RAID 1 logical unit, the controller changes the data stored on the mirrored drive so that it is the same as that on the other drive.

Note that repair only guarantees that the parity information, not the data, is correct. The data in the logical units may be corrupted even though the parity information is correct. If you get a parity error, you may need to locate the file containing the error and restore the file from back-up.

To ensure that a redundant logical unit will be able to recover after a failure, you should set your RAID Modules to perform Automatic Parity Check/Repair on a regular basis. You can set it to run either daily or once a week. At times, you may need to run Parity Check/Repair manually as part of the recovery process. The two following sections describe how to set the time for Automatic Parity Check and how to run a Manual Parity Check.

- Setting Automatic Parity Time—page 154
- Running Manual Parity Check—page 156

Setting Automatic Parity Time

To ensure data integrity on your RAID Modules, you should set your software to run Automatic Parity Check/Repair on a regular basis (either daily or once a week).

Note – If you have very large RAID Modules, it may take longer than 24 hours to run Automatic Parity Check/Repair. This can result in multiple parity checks running at the same time. If this occurs, set Parity Check/Repair to run once a week.

To set the time for Automatic Parity Check/Repair, perform the following procedure.

1. **Start the Maintenance and Tuning application.**
2. **Select any RAID Module or All RAID Modules.**

   The automatic time always applies to all RAID 1, 3, and 5 logical units on all the RAID Modules on your system, even if you select an individual module.
3. **Select Option → Auto Parity Settings from the drop-down menus.**

   A screen like FIGURE 7-2 is displayed.
4. **Enter the settings you want, then click Save to make the settings effective, or press Cancel to return to the application without changing the settings.**

Note – Automatic parity check will not be performed twice in the same day. If it has already run today, and you specify a later time for it to run today, it will not run until that time the next day. For example, if automatic parity check was performed at 2:00 a.m., and at 1:00 p.m. you set it to run at 11:00 p.m., it will not run until 11:00 p.m. tomorrow. However, this does not affect your ability to run a manual parity check.

End of Procedure

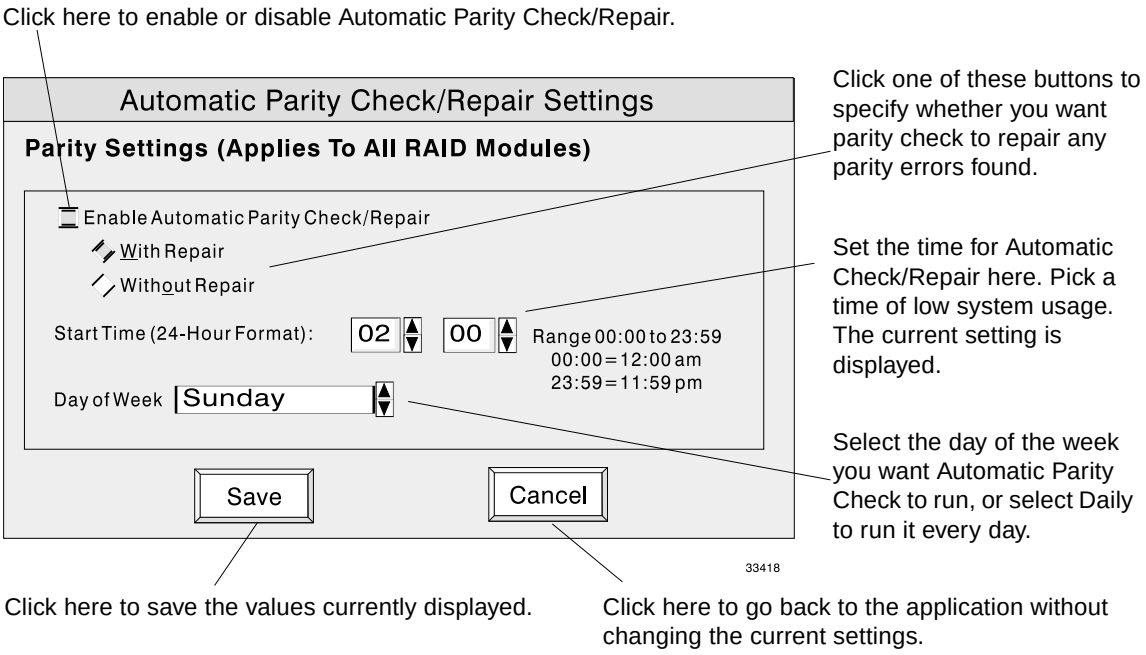

FIGURE 7-2 Setting the Time for Automatic Parity Check/Repair

Running Manual Parity Check

Normally, you should set your RAID Modules so that Parity Check is run automatically on a regular basis (see “Setting Automatic Parity Time” on page 154 for information on how to set the time for Automatic Parity Check). However, during recovery, or if you suspect a problem, you may need to run Manual Parity Check/Repair to find parity errors immediately.

To run Parity Check/Repair manually, perform the following procedure.

1. **Start the Recovery application.**
2. **Select the RAID Module containing the logical units (LUNs) you want to check (or select All RAID Modules).**
3. **Click the Manual Parity Check/Repair button or select Options → Manual Parity Check/Repair from the drop-down menus.**

   A screen like FIGURE 7-3 is displayed.
4. **Select the LUNs you want to check.**
5. **After you have selected the LUNs you want to check, click the Start Parity Check/Repair button.**

   Click any of the other buttons to exit this screen without performing the check.
6. **After you have selected the LUNs you want to check, click Start Parity Check/Repair. Click any of the other buttons to exit this screen without performing the check.**

   A screen like FIGURE 7-4 is displayed.
7. **Click the repair option you want, then click OK.**

   As each LUN is checked, a histogram bar appears on the screen indicating the Parity Check/Repair progress on that LUN.

   When Parity Check/Repair is completed, you see a message indicating if any errors were found and repaired.

   End of Procedure

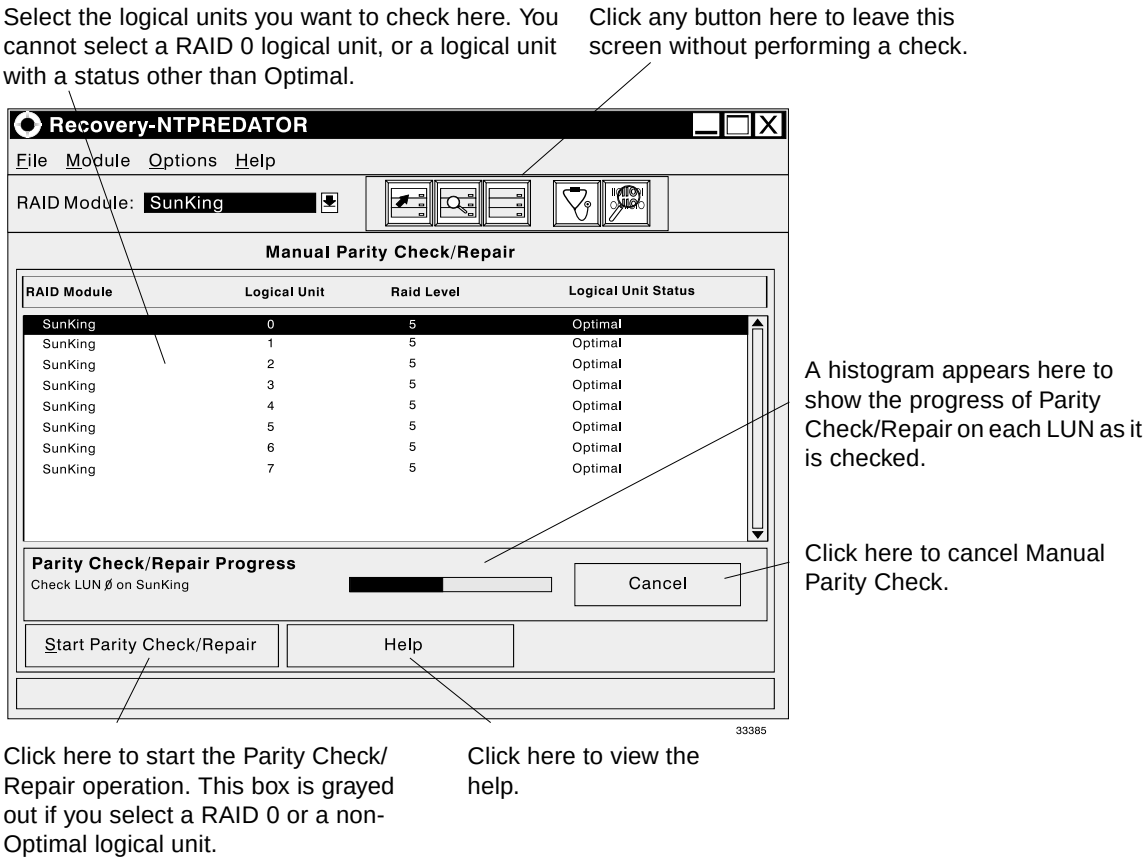

FIGURE 7-3 Manual Parity Check Repair Screen

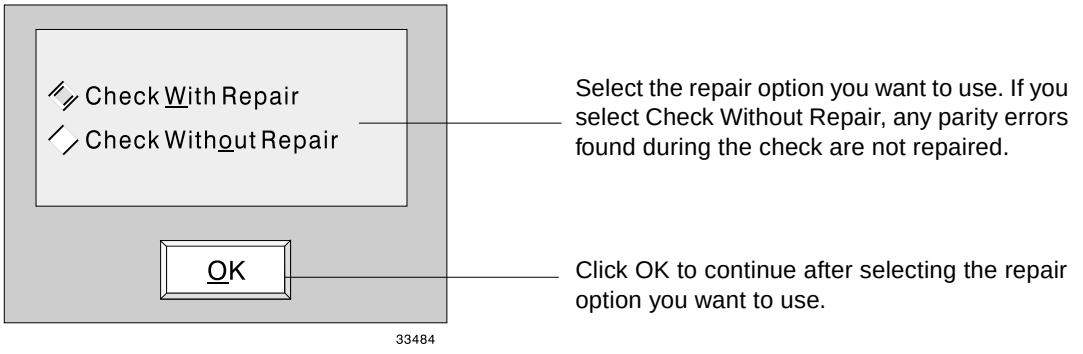

FIGURE 7-4 Manual Parity Check/Repair Option Screen

Note the following:

- You cannot run Parity Check/Repair on a RAID 0 logical unit or a non-optimal RAID 1, 3, or 5 logical unit.
- Parity repair fixes parity, not data. If the parity inconsistencies were caused by corrupted data, the data is still corrupted, even though the parity is correct.
- While parity check/repair is in progress, you cannot perform other Recovery tasks.

Changing the Reconstruction Rate

The Reconstruction Rate is the rate at which the controller rebuilds data on a replaced drive. The faster the reconstruction rate, the more time the controller spends rebuilding data and the less time it spends servicing I/O requests from the host. The slower the rate, the faster system I/O occurs, but the longer reconstruction takes to finish. See "Reconstruction" on page 17 for more information.

The Reconstruction Rate also affects the speed of background formatting (see "Logical Unit (LUN) Formatting" on page 64).

Use this option to change the LUN Reconstruction Rate of logical units that are not currently being reconstructed. To change the rate of logical units that are currently being reconstructed, see "Monitoring and Changing Reconstruction" on page 138.

To change the Reconstruction Rate, perform the following procedure.

1. **Start the Maintenance and Tuning application.**
2. **Select the RAID Module you want to change reconstructions rates on. Do not select All RAID Modules or the option will be grayed out.**
3. **Click on the LUN Reconstruction button or select Options → LUN Reconstruction from the drop-down menus.**

   A screen like FIGURE 7-5 is displayed.
4. **For each logical unit, move the tab in the slider bar to the rate you want.**
   - Moving the tab to the left increases system performance during reconstruction, but slows reconstruction.
   - Moving the tab to the right decreases the amount of time reconstruction takes, but slows system performance.
5. **After you have made all the changes you want, click Save to implement the changes.**

   End of Procedure

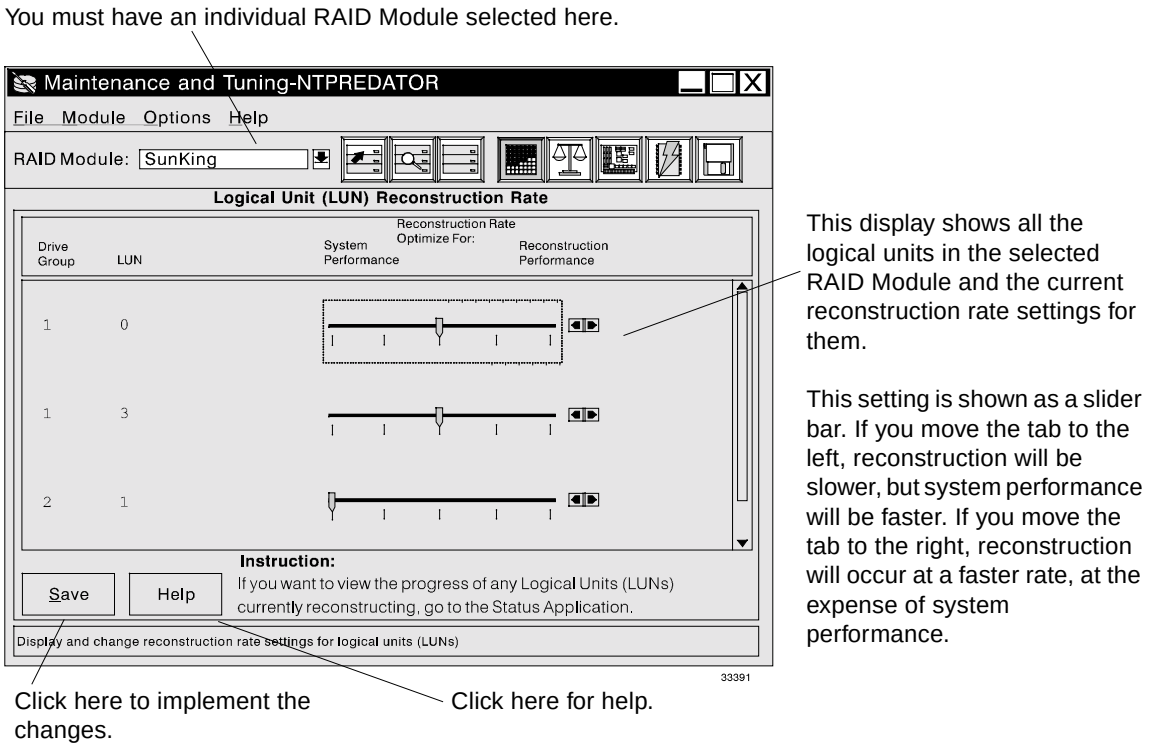

FIGURE 7-5 Changing the Reconstruction Rate Screen

Note the following:

- You cannot perform this option if you select All RAID Modules.
- The new reconstruction rate goes into effect the next time the selected logical units are reconstructed.
- You can use this option to change the reconstruction rate of LUNs currently being reconstruction; however, it is best to use the procedure given in “Monitoring and Changing Reconstruction” on page 138.
- From left to right, the points on the Slider bar indicate the following reconstruction rates (blocks / seconds delay). Note that each 1024 block = 512 KB.
  - Slow – 1024 blocks / 0.8 second delay
  - Medium slow – 1024 blocks / 0.4 second delay
  - Medium – 1024 blocks / 0.2 second delay
  - Medium fast – 1024 blocks / 0.1 second delay
  - Fast – Reconstruct with no delays

Balancing Drive Groups/LUNs

Balancing drive groups/LUNs involves assigning drive groups/LUNs to controllers in an active/active redundant controller pair (see “Redundant Controllers” on page 17 for a definition of redundant controllers). Some balancing occurs when you create drive groups/LUNs. The first drive group you create in a configuration session is assigned to controller A, the second to controller B, the third to controller A again, and so on. However, if you repeatedly create one drive group per session, you will end up with all drive groups assigned to controller A. You can manually balance the drive groups to change this default balancing, or if you need to reassign drive groups after a controller failure.

Logical units are balanced by *drive group.* That is, all of the LUNs in a drive group must be assigned to the same controller. This may result in more actual logical units assigned to one controller than to the other, even if each controller is assigned the same number of drive groups. For example, if drive groups 1 and 2 have one logical unit each and drive group 3 has four logical units, relying on automatic balancing results in one controller servicing five logical units while the other controller services only one.

Using the LUN balancing option, you can either:

- Automatically balance drive groups/LUNs (see “Automatic LUN Balancing” on page 161).
- Manually balance the drive groups/LUNs in one RAID Module (see “Manual LUN Balancing” on page 163).

Automatic LUN Balancing

Note – If a RAID Module does not have RDAC protection, it will not be balanced unless all I/O to that module is stopped. If none of the displayed RAID Modules are eligible for balancing, the Balance button is grayed out.

To automatically balance drive groups/logical units in RAID Modules, use the following procedure.

1. **Start the Maintenance and Tuning application.**
2. **Select All RAID Modules. (If you select a single RAID Module, you can only balance the logical units manually; see** "Manual LUN Balancing" on page 163**).**
3. **Click on the LUN Balancing button or select Options → LUN Balancing.**

   A screen like FIGURE 7-6 is displayed.
4. **Select the RAID Modules you want to balance and click the Balance button to automatically balance them.**

   End of Procedure

   Note the following:

   - You can only select RAID Modules with active/active controller pairs.
   - You cannot select controllers in an Independent Configuration (you must use manual LUN balancing to reassign drive groups in this configuration; see "Manual LUN Balancing" on page 163).
   - Balancing is done by *drive group*, so the number of LUNs assigned to each controller will not necessarily be the same after the LUNs are balanced.
   - Odd-numbered drive groups are assigned to one controller and even numbered drive groups are assigned to the second controller.
   - Logical units/drive groups will not be balanced on RAID Modules without RDAC protection unless all I/O to that module is stopped.
   - If none of the displayed RAID Modules are eligible for balancing, the Balance button is grayed out.

You must have All RAID Modules selected here. If you have selected a specific RAID Module, see “Manual LUN Balancing” on page 163.

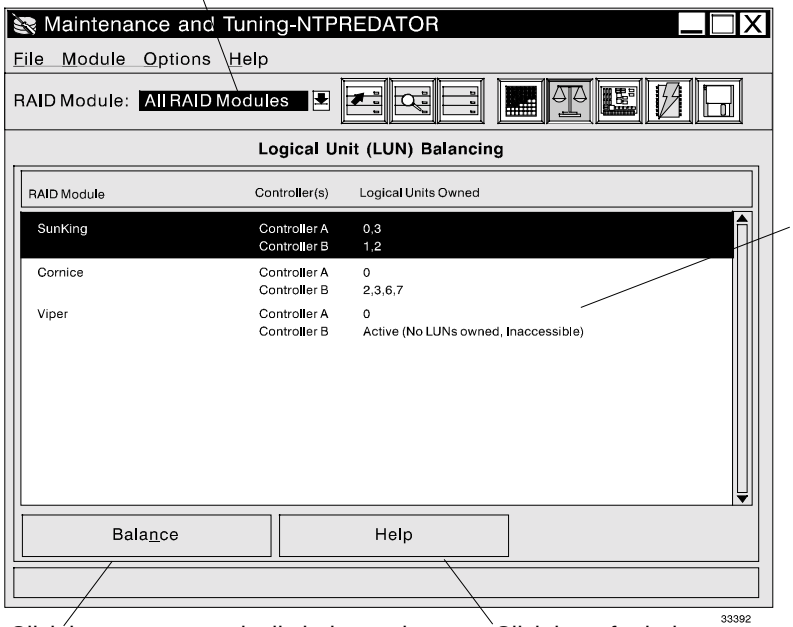

This display shows all the RAID Modules on your system and the current LUN assignments on them. Select the RAID Modules you want to balance and click on the Balance button.

Remember that balancing is by drive group, so the number of LUNs assigned to each controller will not necessarily be the same after the LUNs are balanced.

Click here to automatically balance the LUNs on the selected RAID Modules.

Click here for help.

FIGURE 7-6 Automatic LUN Balancing Screen

Manual LUN Balancing

To manually balance the drive groups/LUNs in a single RAID Module, perform the following procedure.

Caution – If the RAID Module does not have RDAC failover protection, you must stop all I/O to that module before balancing the drive groups.

1. **Start the Maintenance and Tuning application.**
2. **Select the individual RAID Module containing the drive groups/LUNs you want to reassign.**

   (If you select All RAID Modules, you can only automatically balance all of the drive groups in all RAID Modules; see"Automatic LUN Balancing" on page 161).
3. **Click on the LUN Balancing button or select Options → LUN Balancing.**

   A screen like FIGURE 7-7 is displayed.
4. **Make the re-assignments you want by highlighting the drive groups and then clicking on the Move button.**
5. **After you have made all the changes you want to make, click Save to make the changes.**

   End of Procedure

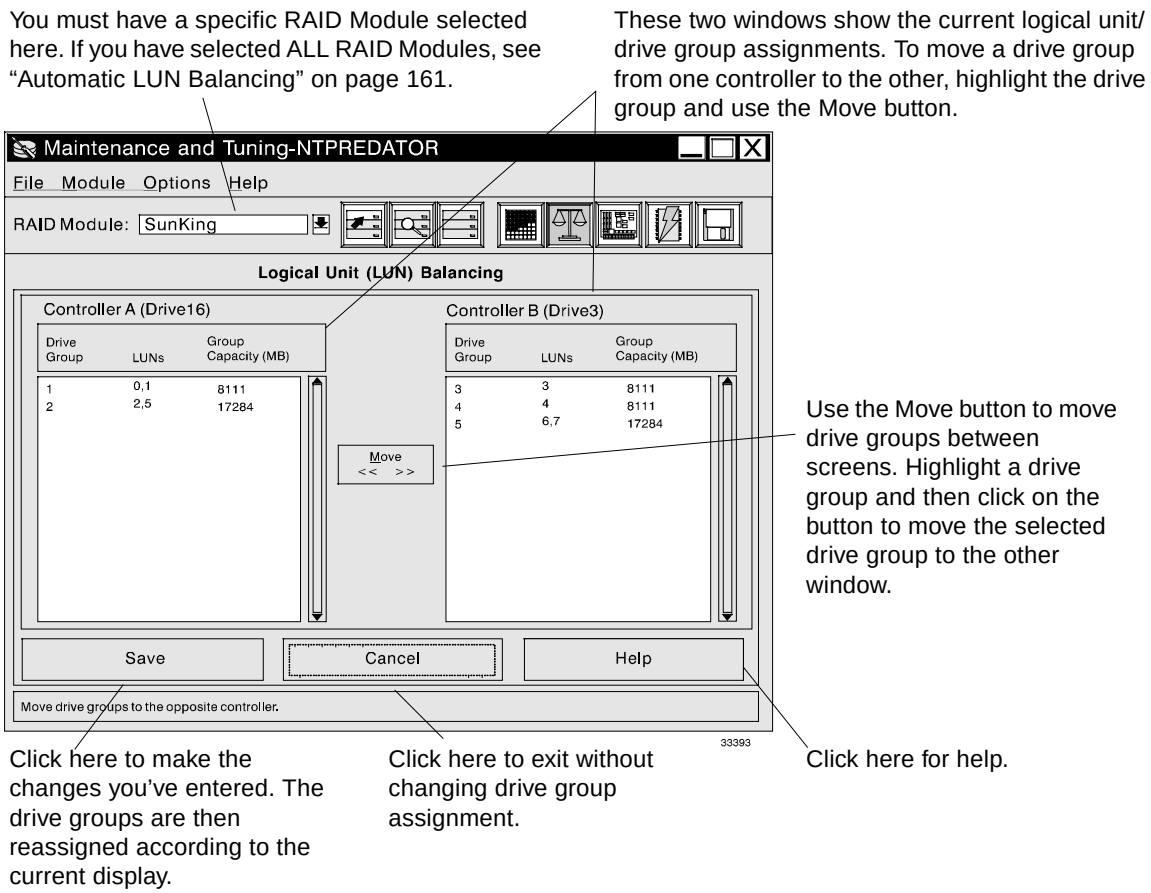

FIGURE 7-7 Manual LUN Balancing Screen

Note the following:

- You can only select RAID Modules with active/active controller pairs.
- If the RAID Module does not have RDAC protection, you must stop all I/O to the module before balancing the LUNs.
- In an Independent Controller configuration, you can only reassign LUNs *from* the inaccessible controller *to* the accessible controller (the controller connected to the host system you are using). You cannot assign LUNs from the accessible controller to the inaccessible controller.

Changing Controller Mode

Tip – You can use this option only if the redundant controllers in the RAID Module are active/passive.

Use this option to:

- Swap the active/passive controllers in a RAID Module so that the active controller becomes passive and the passive controller becomes active.
- Change the active/passive status to active/active.

Note that you cannot use this option to change controllers from active/active to active/passive. This can only be done with the `rdacutil` utility. For information on the `rdacutil` utility, see Appendix A.

To change the controller mode of your controllers, perform the following procedure.

Caution – If the selected RAID Module does not have RDAC protection, you must stop all I/O to that module before changing controller modes.

1. **Start the Maintenance and Tuning application.**
2. **Select the RAID Module with the controllers you want to change or select All RAID Modules to change controller mode in more than one RAID Module.**

Note – You cannot select this option unless the selected RAID Module contains an active/passive controller pair, or, if you selected All RAID Modules, at least one controller pair is active/passive.

3. **Click on the Controller Mode button or select Option → Controller Mode from the drop-down menu.**

   A screen like FIGURE 7-8 is displayed.

4. **Perform *one* of the following steps:**
   - Swap the controllers, making the active one passive and the passive one active.
   - Change the controller mode to active/active. A confirmation message appears, giving you the chance to automatically balance the drive groups/LUNs across the two controllers during the mode change.

End of Procedure

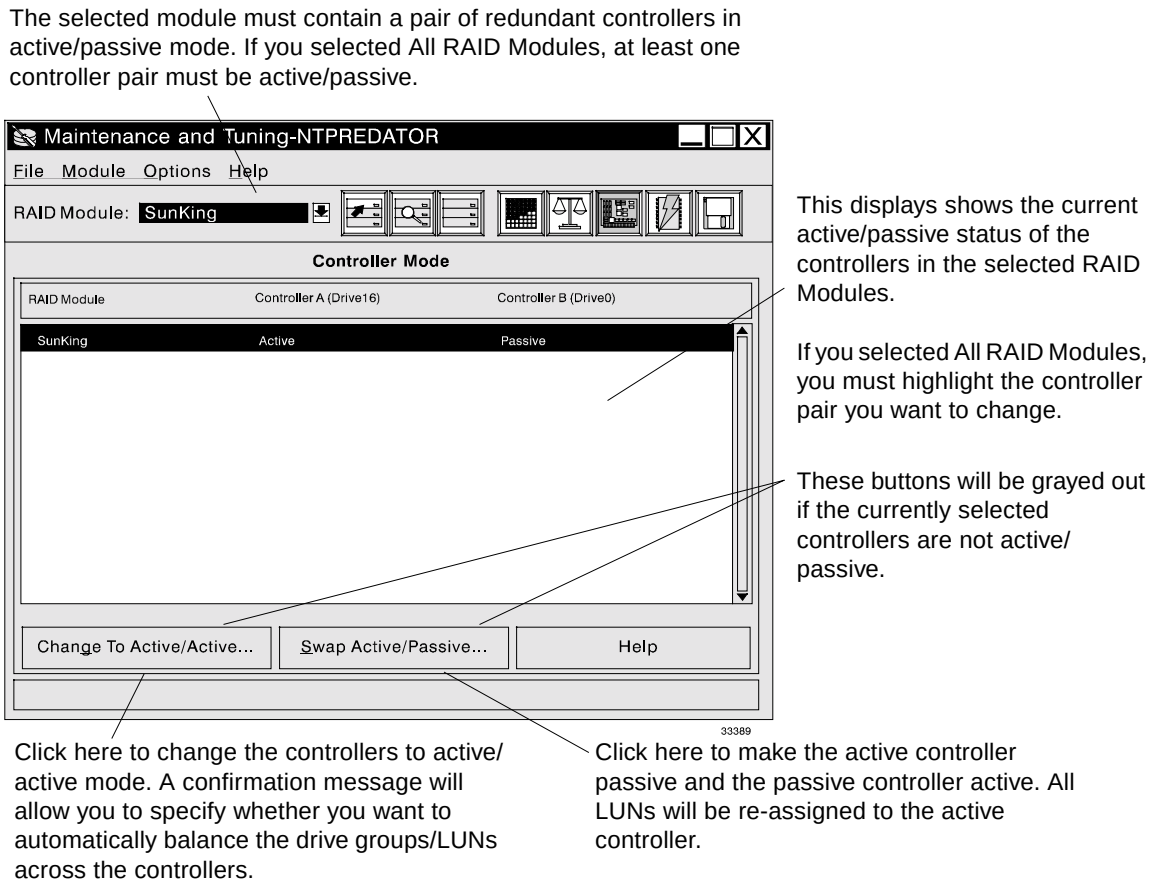

FIGURE 7-8 Change Controller Mode Screen

Note the following:

- If the RAID Module does not have RDAC protection, you must stop all I/O to the module before changing controller mode.
- You cannot select this option unless the selected RAID Module contains an active/passive controller pair, or, if you selected All RAID Modules, at least one controller pair is active/passive.
- You cannot change active/active controllers to active/passive using this option. See Appendix A for more information about the `rdacutil` utility.
- When you swap active/passive controllers, the logical units are automatically switched to the other controller.

Changing Cache Parameters

Cache is a memory area on the controller that stores read/write data from the host. Using cache memory can speed up I/O operations. See "Cache Memory" on page 9 for more information on cache memory.

You can set three cache parameters: Write Caching, Write Cache Mirroring, and Caching Without Batteries. You can change other cache parameters using the `raidutil` utility. For information about the `rdacutil` utility, see Appendix A.

You can only change the Write Cache Mirroring parameter if:

- Both controllers in the RAID Module (if there are two) have the same cache memory size (at least 8 MB).
- Write cache mirroring can be assigned to a single controller, but this has no effect.

You can check the cache capacity of the controllers through the Module Profile controller information screen (see "The Controller Profile Screen" on page 50).

To change the cache parameters, perform the following procedure.

1. **Start the Maintenance and Tuning application.**
2. **Select the RAID Module containing the logical units you want to set cache parameters for.**
3. **Click on the Cache Parameters button or select Option → Cache Parameters from the drop-down menus.**

   A screen like FIGURE 7-9 is displayed.

Note – You might see an asterisk next to the caching parameters column. This indicates that the parameter is enabled, but is *not* currently active. The controller has disabled the parameter for some reason (such as low batteries). If you see this condition, use Message Log (Status application) to determine the correct action to take.

---

4. **Click in the boxes to select or deselect individual cache parameters for the logical units.**

   The parameters are interdependent; clicking on some automatically selects or deselects others. See TABLE 7-3.

5. **After you have made all the changes you want, click Save to apply the changes.**

   End of Procedure

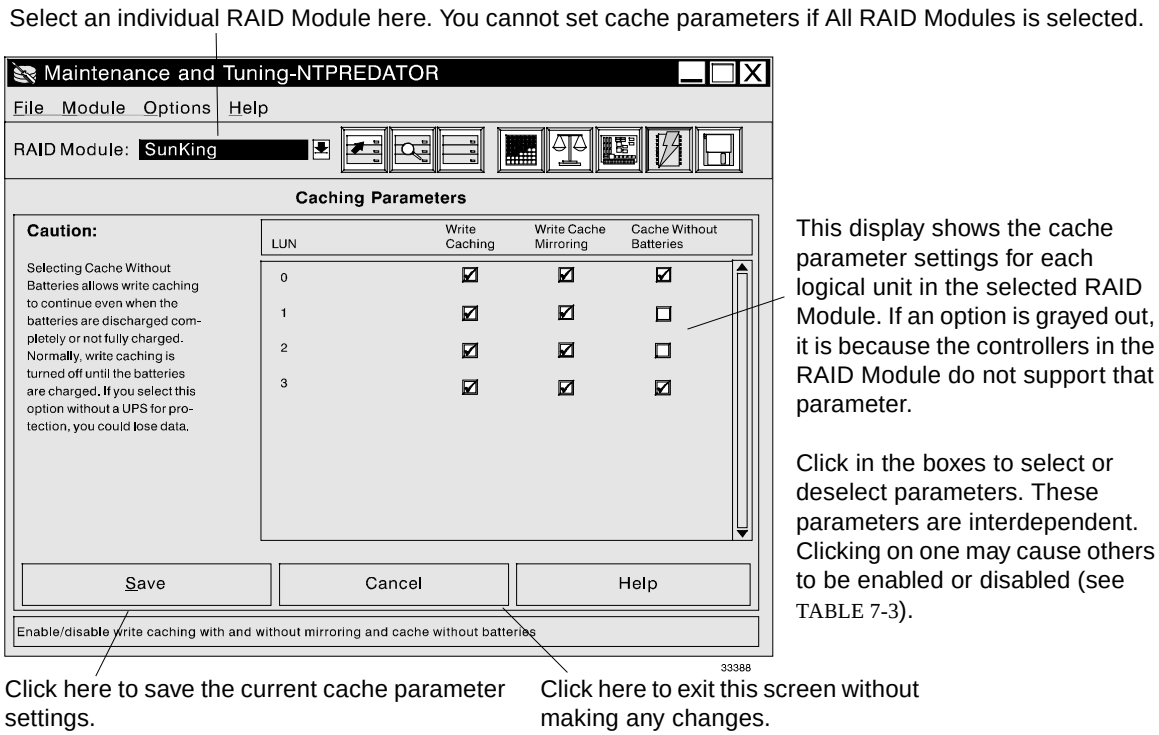

FIGURE 7-9 Cache Parameters Screen

TABLE 7-3 Cache Parameter Interdependencies

If you select:	**There parameters are enabled:**
Write Caching	Write Cache Mirroring[1]
Write Cache Mirroring[1]	Write Caching
Cache Without Batteries	Write Caching
If you deselect:	**these parameters are disabled:**
Write Caching	Write Cache Mirroring Cache Without Batteries

[1] Write Cache Mirroring is only effective for modules with redundant controller pairs (active/active or active/passive) that have the same size cache. Use Module Profile → Controllers to determine if both controllers in the pair have the same cache size *before* enabling this parameter. Note that the parameter can be assigned to any controller, even single controllers.

APPENDIX **A**

Command-Line Interface

This appendix describes the command-line interface to the storage management software.

Command-Line Utilities

TABLE A-1 summarizes the command-line utilities, scripts, and background processes available in the storage management software. All of these items have associated Help files available as `man` pages (UNIX) or as Help text files (Windows NT). The utilities are fully explained in these Help files.

Most of these utilities duplicate operations in the GUI, and therefore need to be performed only as a last resort.

TABLE A-1 Command-Line Utilities and Program Descriptions

Option	Description
Informational	
`rm6`	Gives an overview of the software's graphical user interface (GUI), command-line programs, background process programs and driver modules, and customizable elements.
`rdac`	Describes the software's support for RDAC (Redundant Disk Array Controller), including details on any applicable drivers and daemons.
`rmevent`	The RAID Event File Format. This is the file format used by the applications to dispatch an event to the `rmscript` notification script. It also is the format for Message Log's log file (the default is `rmlog.log`).
`raidcode.txt`	A text file containing information about the various RAID events and error codes.
Command-Line Utilities	
`drivutil`	The drive/LUN utility. This program manages drives/LUNs. It allows you to obtain drive/LUN information, revive a LUN, fail/revive a drive, and obtain LUN reconstruction progress.
`fwutil`	The controller firmware download utility. This program downloads *appware*, *bootware*, or an NVSRAM file to a specified controller.
`healthck`	The health check utility. This program performs a health check on the indicated RAID module and displays a report to standard output.
`lad`	The list array devices utility. This program identifies the RAID controllers and logical units that are connected to the system.
`logutil`	The log format utility. This program formats the error log file and displays a formatted version to the standard output.

TABLE A-1 Command-Line Utilities and Program Descriptions *(Continued)*

Option	Description
`nvutil`	The NVSRAM display / modification utility. This program views and changes RAID controller non-volatile RAM settings, allowing for some customization of controller behavior. It verifies and fixes any NVSRAM settings that are not compatible with the storage management software.
`parityck`	The parity check / repair utility. This program checks and, if necessary, repairs the parity information stored on the array.
`raidutil`	The RAID configuration utility. This program is the command-line counterpart to the graphical Configuration application. It allows you to create and delete RAID logical units and hot spares from a command line or script. It also allows certain battery management functions to be performed on one controller at a time.
`rdacutil`	The redundant disk array controller management utility. This program permits certain redundant controller operations such as LUN load balancing and controller failover and restoration to be performed from a command line or script.
`storutil`	The host store utility. This program performs certain operations on a region of the controller called host store. You can use this utility to set an independent controller configuration, change RAID module names, and clear information in the host store region.
Background Process Programs and Driver Modules	
`arraymon`	The array monitor background process. The array monitor watches for the occurrence of exception conditions in the array and provides administrator notification when they occur.
`rdaemon` (UNIX only)	The redundant I/O path error resolution daemon. The `rdaemon` receives and reacts to redundant controller exception events and participates in the application-transparent recovery of those events through error analysis and, if necessary, controller failover.
`rdriver` (Solaris only)	The redundant I/O path routing driver. The `rdriver` module works in cooperation with `rdaemon` in handling the transparent recovery of I/O path failures. It routes I/Os down the proper path and communicates with the `rdaemon` about errors and their resolution.

TABLE A-1 Command-Line Utilities and Program Descriptions *(Continued)*

Option	Description
Customizable Elements	
`rmparams`	The storage management software's parameter file. This ASCII file has a number of parameter settings, such as the array monitor poll interval, what time to perform the daily array parity check, and so on. The storage management applications read this file at startup or at other selected times during their execution. A subset of the parameters in the `rmparams` file are changeable under the graphical user interface. For more information about the `rmparams` file, see the *Sun StorEdge RAID Manager Installation and Support Guide.*
`rmscript`	The notification script. This script is called by the array monitor and other programs whenever an important event is reported. The file has certain standard actions, including posting the event to the message log (`rmlog.log`), sending email to the superuser / administrator and, in some cases, sending an SNMP trap. Although you can edit the `rmscript` file, be sure that you do not disturb any of the standard actions. For more information about the `rmscript` file, see the *Sun StorEdge RAID Manager Installation and Support Guide.*

Index

D

E

F

G

H

I

K

L

M

N

O

P

R

S

T

U

V

W