



Netfinity SYMplicity Storage Manager

User's Handbook

Before using this information and product it supports, be sure to read the general information under "[Appendix B. Notices](#)" on page 139.

Second Edition (July 1999)

INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some jurisdictions do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

This publication could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. IBM may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time.

This publication was developed for products and services offered in the United States of America and the United Kingdom. It is possible that this publication may contain references to, or information about, IBM products (machines and programs), programming, or services that are not announced in your country. Such references or information must not be construed to mean that IBM intends to announce such IBM products, programming, or services in your country.

Requests for technical information about IBM products should be made to your IBM reseller or IBM marketing representative.

No part of this publication may be reproduced or distributed in any form or by any means without prior permission in writing from the International Business Machines Corporation.

©Copyright International Business Machines Corporation 1998, 1999. All rights reserved.

Note to U.S. Government Users—Documentation related to restricted rights—Use, duplication or disclosure is subject to restrictions set forth in GSA ADP Schedule Contract with IBM Corp.

Contents

Figures	ix
Tables	xiii
Chapter 1. Getting Started	1
About this User's Handbook	1
About the Installation and User's Handbook	2
About the Hardware Handbook	2
About Software Installation	2
What to Do Next	3
Chapter 2. Storage Management Software	5
Common Definitions	5
Cache Memory	5
Device Name	6
Drive Group	6
Drive Group Numbering	7
Drive Names	8
Hot Spare	9
Logical Unit (LUN)	9
Parity	10
Parity Check/Repair	10
RAID Level	10
RDAC Driver	11
RAID Controller Unit	11
RAID Controller Module	11
Reconstruction	12
Redundant Controllers	12
Segment Size	13
SNMP	13

Controller Configurations	13
Single-Host Configuration	14
Multi-Host Configuration	15
Independent Controller Configuration	17
Special Network Considerations	18
Component Failures	19
Protecting Yourself from Component Failures	19
Application Summary Charts	20
Task List	25
Chapter 3. Common Application Features.....	29
New Software Features	29
Navigation	30
Starting and Exiting an Application	30
Starting an Application	30
Exiting an Application	31
Using Online Help	31
Selecting a Module	34
RAID Controller Module Selection Screen	35
Locating a Module	36
Locating Drives	37
Viewing a Module Profile	38
Controller Profile Screen	40
Drive Profile Screen	42
LUN Profile Screen	43
Saving a Module Profile	44
Chapter 4. Configuring RAID Controllers.....	47
Configuration Overview	47
Types of Logical Unit (LUN) Creation	47
Types of Logical Unit (LUN)/Drive Group Modification	47
Module Information Display	48

Logical Unit Formatting	49
Formatting Restrictions	51
Configuration Application Main Screen	52
Creating a Logical Unit (LUN)	54
Creating a Logical Unit from Unassigned Drives	54
Creating a Logical Unit from Remaining Capacity	57
Using the Options Screens	59
Modifying Drive Groups and LUNs	63
Defragment Option	64
Defragmenting the Drive Group	66
Changing the RAID Level	67
Adding Drives	69
Modifying Segment Size	71
Deleting LUNs	72
Creating a Hot Spare	73
Deleting a Hot Spare	75
Resetting the Configuration	76
Chapter 5. Checking Status	79
Status Overview	79
Background Polling	79
Message Log File	79
Message Types	80
Status Log and Health Check	80
Event Notification	80
Viewing the Log	80
Changing Log File Settings	84
Clearing the Message File	85
Saved Log File Location	86
Using the Performance Monitor	86
Using Performance Monitor Data	90

Chapter 6. Recovery	91
Recovery Overview	91
Recovery Procedures	91
General Recovery Steps	92
Recovering from Drive Failures	92
Recovering from Controller/Connection Failures	97
Recovering from Fan/Power Supply Failures	99
Recovering from Battery Failures	100
Running Recovery Guru/Health Check	100
Manual Recovery Options	103
Manual Recovery of Drives	104
Manual Recovery of Logical Units	106
Manual Recovery of Controllers	108
Monitoring/Changing Reconstruction	109
Chapter 7. Maintenance and Tuning	113
Maintenance and Tuning Overview	113
Upgrading the Controller Firmware	113
Controller Firmware Files	114
General Procedure	114
Online and Offline Method	115
Selecting Controllers	115
Download Procedure	116
Confirming the Download	120
Running a Parity Check	122
Setting Automatic Parity Time	122
Running a Manual Parity Check	123
Changing the LUN Reconstruction Rate	125
Balancing LUNs	127
Automatic LUN Balancing	127
Manual LUN Balancing	129

Changing the Controller Mode	130
Changing Cache Parameters	133
Appendix A. Command Line Interface.....	135
Command Line Utilities	135
Appendix B. Notices	139
Trademarks	140
Index.....	141

Figures

Figure 1. Drive Groups and Logical Units (LUNs)	7
Figure 2. Single-Host to RAID Controller Configuration	15
Figure 3. Multi-Host to RAID Controller Configuration	16
Figure 4. Independent Controller Configuration	18
Figure 5. Configuration Application Summary	21
Figure 6. Status Application Summary	22
Figure 7. Recovery Application Summary	23
Figure 8. Maintenance and Tuning Application Summary	24
Figure 9. Application Icons	31
Figure 10. Main Online Help Screen	32
Figure 11. RAID Controller Module Selection Screen	35
Figure 12. List/Locate Drives Screen	38
Figure 13. Module Profile Screen	39
Figure 14. Controller Profile Screen	40
Figure 15. Drive Profile Screen	42
Figure 16. LUN Profile Screen	43
Figure 17. Save Module Profile Screen	45
Figure 18. Configuration Application Module Information Display	48
Figure 19. RAID 5 LUN Creation Example	50
Figure 20. Configuration Application Main Screen	52
Figure 21. Create LUN Main Screen (Unassigned Drives)	56
Figure 22. Create LUN Main Screen (Existing Drive Group)	58
Figure 23. LUN Capacity Screen	59

Figure 24. Drive Selection Screen	60
Figure 25. Cache Parameters Screen	61
Figure 26. Segment Size Screen	62
Figure 27. Set LUN Assignment Screen	63
Figure 28. Defragment Logical Units #1	64
Figure 29. Defragment Logical Units #2	65
Figure 30. Defragment Logical Units #3	65
Figure 31. Defragment Logical Units #4	66
Figure 32. Modify RAID Level Screen	68
Figure 33. Modify Group Capacity Screen	70
Figure 34. Modify Segment Size Screen	71
Figure 35. Delete LUNs Screen	73
Figure 36. Create Hot Spare Screen	74
Figure 37. Select Hot Spare Screen	75
Figure 38. Delete Hot Spares Screen	76
Figure 39. Selecting Message Type	81
Figure 40. Message Log Screen	82
Figure 41. Message Details Screen	83
Figure 42. Log Settings Screen	84
Figure 43. Message Limit Warning Message	85
Figure 44. Message Log Save Message	86
Figure 45. Performance Monitor Settings Screen	87
Figure 46. Performance Monitor Start-Up Screen	88
Figure 47. Performance Monitor Data Screen	88
Figure 48. Recovery Main Screen	101

Figure 49. Recovery Screen/Exceptions Found	102
Figure 50. Recovery Procedure Screen	103
Figure 51. Manual Recovery of Drives Screen	105
Figure 52. Manual Recovery of Logical Units Screen	107
Figure 53. Manual Recovery of Controllers Screen	109
Figure 54. Viewing/Changing Reconstruction Screen	111
Figure 55. Firmware Upgrade Screen	118
Figure 56. Setting the Time for Automatic Parity Check/Repair	123
Figure 57. Manual Parity Check Repair Screen	124
Figure 58. Changing the Reconstruction Rate Screen	126
Figure 59. Automatic LUN Balancing Screen	128
Figure 60. Manual LUN Balancing Screen	130
Figure 61. Change Controller Mode Screen	132
Figure 62. Cache Parameters Screen	134

Tables

Table 1. Storage Management Task List	25
Table 2. Using a Mouse and Keyboard	30
Table 3. Online Help Screen Elements	32
Table 4. Select Module Screen Elements	36
Table 5. Module Profile Screen Contents	39
Table 6. Controller Profile Screen Contents (Fibre Channel)	40
Table 7. Drive Profile Screen Contents	43
Table 8. LUN Profile Screen Contents	44
Table 9. Configuration Application Main Screen Elements	53
Table 10. Configuration Display Columns	54
Table 11. Adding Capacity Formulas	69
Table 12. Message Types	80
Table 13. Message Log Explanation	82
Table 14. Performance Monitor Screen Contents	89
Table 15. Drive Failures and Recovery (Summary)	94
Table 16. Recovering from Failed Controllers (Summary)	98
Table 17. Firmware Confirmation Screen Elements	119
Table 18. Controller Firmware Upgrade, Failed Statuses	120
Table 19. Cache Parameter Interdependencies	134
Table 20. Command Line Utilities and Program Description	135

Chapter 1. Getting Started

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

About this User's Handbook

This *User's Handbook* provides the information necessary to set up, monitor, tune, and repair your RAID Controllers using the storage management software. It is divided into 7 chapters and an appendix.

- "[Chapter 1. Getting Started](#)" provides an overview of the chapters contained in this book as well as descriptions of other companion documentation.
- "[Chapter 2. Storage Management Software](#)" describes the storage management 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 Controllers.
- "[Chapter 3. Common Application Features](#)" describes the software features common to all applications.
- "[Chapter 4. Configuring RAID Controllers](#)" describes how to configure the drives in the RAID Controllers into logical units (LUNs) and hot spares.
- "[Chapter 5. Checking Status](#)" describes how to check the status and performance of the RAID Controllers.
- "[Chapter 6. Recovery](#)" describes how to recover after device failures. These failures may involve drives, controllers, fans, and power supplies. In some cases, you also need the information found in the hardware manuals accompanying the RAID Controllers for specifics on how to do the physical tasks required in the recovery process.
- "[Chapter 7. Maintenance and Tuning](#)" describes the routine maintenance and tuning RAID Controllers.
- "[Appendix A. Command Line Interface](#)" describes the command line interface to the RAID Controllers.
- "[Appendix B. Notices](#)" lists warranty provisions and legal notices.

About the Installation and User's Handbook

In addition to this book, you also received an *Installation and User's Handbook* for the particular operating system you are using to access the RAID Controllers.

The *Installation and User's Handbook* contains:

- The installation procedure for installing the storage management software on that operating system.
- Special restrictions on the use of the storage management software on that operating system.
- Troubleshooting information for the storage management software.
- Other notes on how the storage management software runs in that specific operating system environment.

You must refer to the *Installation and User's Handbook* for installation instructions before doing anything else on the RAID Controllers as described in this book.

About the Hardware Handbook

Along with your RAID Controllers, you received one or more hardware books describing how to install the RAID hardware and replace hardware components. These hardware handbooks also contain information on troubleshooting the modules and determining which components have failed.

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

Always follow the software recovery procedures given in this *User's Handbook*, if any, when replacing a hardware component. Use the hardware books, if necessary, at the appropriate time to find how the physical replacement is performed.

About Software Installation

The software installation procedure is given in the *Installation and User's Handbook* that you received along with this book. This *Handbook* is specific to the operating system you are using. Refer to the *Installation and User's Handbook* for complete information on installation, then return to this *User's Handbook* for information on how to run the installed software.

What to Do Next

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

1. Install the RAID Controller hardware, using the information given in the hardware handbooks.
2. Install the storage management software, using the information in the *Installation and User's Handbook*.
3. If necessary, read "[Chapter 2. Storage Management Software](#)" to become acquainted with the storage management software. Be sure to read the task summary list in that chapter for a description of the routine tasks you need to perform to manage the RAID Controllers.
4. If necessary, read "[Chapter 3. Common Application Features](#)" to familiarize yourself with the common program features.
5. Go to "[Chapter 4. Configuring RAID Controllers](#)" to configure the RAID Controllers into logical units and hot spares, according to your system needs.

Chapter 2. Storage Management Software

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

Common Definitions

This section describes the concepts used by the storage management software. You will need to understand these concepts to use your RAID Controllers 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, you can specify various caching parameters for the LUNs. If you need to change any caching parameters after LUN creation, use the Maintenance and Tuning Application (see ["Changing Cache Parameters" on page 133](#)).

The following cache options are available through the storage management software. You can enable these options 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 itself.
- **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 completely, not fully charged, or if there are no batteries present. If you select this option without a UPS for protection, you could lose data if power fails.

You can change other cache options using the rdacutil utility. See "[Appendix A. Command Line Interface](#)"

Device Name

The storage management software uses device names to identify the controllers and logical units in the RAID Controller. 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 scheme, the Windows NT operating system uses a driveX scheme, and the Networked version uses an eight character scheme defined during installation.

See your *Installation and User's Handbook* 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 Controller. There are three types of drive groups:

- Unassigned drive group — the drives in the RAID Controller 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 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 with a number. These drive groups are displayed by number in all applications. See the next section for a discussion of how drive groups are numbered.

Figure 1 shows a 10-drive tray organized into drive groups and LUNs.

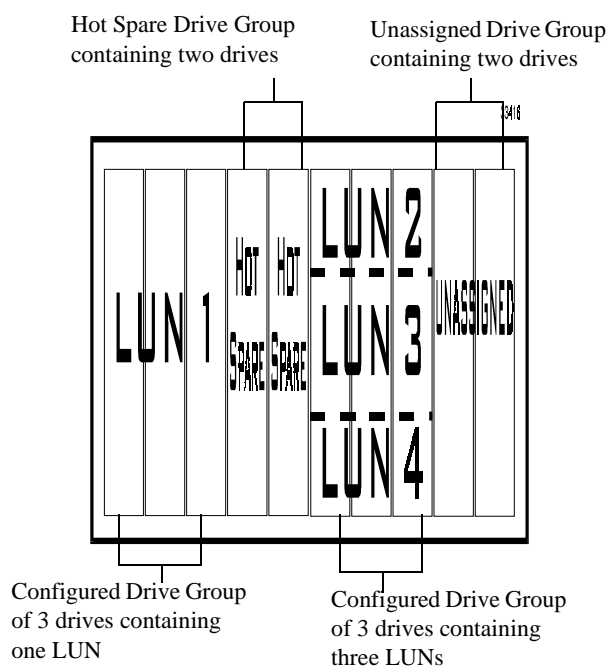


Figure 1. Drive Groups and Logical Units (LUNs)

Drive Group Numbering

Each drive group in a RAID Controller has a number from 1 to 32 (depending on how many logical units you can create on your system and on how you have configured them). These numbers may change as you create and delete logical units. Drive group numbering is based on the specific logical unit numbers associated with each drive group. Drive group numbering starts with the lowest numbered logical unit.

IMPORTANT Keep in mind that the drive group numbering can change when you are creating and deleting logical units (LUNs).

For example, the drive group containing LUN 0 will always be drive group 1. When you delete LUNs and then add new LUNs, the drive group numbers may change to reflect the new logical unit numbers associated with it.

For example, suppose you had the following drive groups:

Drive Group #	LUN #
1	0, 1
2	2
3	3, 4, 5

Now, you delete LUN 1. In this case, renumbering would not occur. The drive groups would be as follows:

Drive Group #	LUN #
1	0
2	2
3	3, 4, 5

Next, you create a new drive group. The new drive group will use the first available LUN, which in this case is 1. The drive groups would be renumbered as follows:

Drive Group #	LUN #
1	0
2	1
3	2
4	3, 4, 5

As you can see, LUN 1 is now part of drive group 2. The old drive group 2 has been renumbered to 3, and the old drive group 3 has been renumbered to 4.

Drive Names

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.

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

Hot Spare

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 Controller. If a drive fails, the hot spare takes over for the failed drive until you replace it. Once you replace the failed drive, the hot spare returns to a Standby status.

A hot spare drive is not dedicated to a specific drive group/LUN; instead, it can replace any failed drive in the RAID module that has the same or smaller storage capacity. This is true, even when the RAID modules have independent controllers.

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

See ["Creating a Hot Spare" on page 73](#) for information on how to create a hot spare. See ["Protecting Yourself from Component Failures" on page 19](#) 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 to retrieve and store your data. A LUN is a set of physical drives (that is, a drive group), that have been configured into either RAID Level 0, 1, 3, or 5.

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 new disk drive, and you must follow the same procedures on the LUN (creating file systems, volumes, and so on) that you would use to access a standard disk drive.

One or more LUNs make up a configured drive group. See [Figure 1 on page 7](#) for an example of the relationship between drive groups and LUNs.

Each LUN has several properties:

- A LUN number — this is a number (from 0 to 32, depending on your system) assigned to the LUN by the controller when you create the LUN.
- A RAID level — this determines how data is written to the drives and determines whether or not the LUN 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, all RAID 1 LUNs 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 it.
- 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 6 drive RAID 0 LUN has twice the capacity of a 6 drive RAID 1 LUN, even if both LUNs use the same size drive.
- A set of parameters — These parameters include segment size, cache options, reconstruction rate, and so on. Some of these parameters can be changed for individual LUNs, while others apply to a whole drive group (one or more LUNs occupying the same drives).

Parity

Parity is additional information stored along with the data that allows the controller to reconstruct lost data. Only RAID 3 and 5 have parity. RAID 1 uses mirroring, not parity, although you can run parity check and repair on a RAID 1 logical unit; in this case, parity check compares the data on the mirrored drives.

Parity Check/Repair

Parity check is the process of scanning the blocks in a logical unit and checking the parity for each block. RAID 1 does not use parity, but parity check can still be performed on a RAID 1 unit; in this case, parity check compares the data on the mirrored drives.

Parity repair is the process of correcting any parity inconsistencies found during parity check. However, this process 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 a Parity Check" on page 122](#) 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, so if a single drive fails, all data in the LUN 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, so 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 4-drive RAID 1 logical unit with 18 MB drives, the resulting data capacity is 36 MB).
- RAID 3 — In this version of the storage management software, RAID 3 is identical to RAID 5. However, RAID 3 is better suited to larger record performance. RAID 5 is better suited to small records, such as those used in transaction processing.
- RAID 5 — In RAID 5, data is striped across the drives in the logical unit in segments, with parity information being striped across the drives as well. Because of this parity, if a single drive fails, data can still be recovered from the remaining good drives. Two drive failures cause all data to be lost. A RAID 5 logical unit has data capacity of all the drives in the logical unit less one (for example, a 5-drive RAID 5 logical unit with 18 GB drives has a data capacity of 72 GB).

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 Controller Units with redundant controllers running under Windows NT. If a component (cable, controller, host adapter, etc.) 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 *Installation and User's Handbook* for your operating system for specific information.

RAID Controller Unit

This is the physical unit, which contains one or two controllers and associated power supplies, fans and battery. This *Handbook* sometimes refers to the RAID Controller Unit as “RAID Controller.”

RAID Controller Module

When the RAID Controller Unit is physically connected with one or more drive modules, it becomes a Raid Controller Module. You select a RAID Controller Module to perform the various RAID tasks (such as configuring, obtaining status, recovering, and so on). This *Handbook* sometimes refers to the RAID Controller Module as “RAID Controller.”

Reconstruction

Reconstruction is the process of using data and parity on all operational drives within a drive group to regenerate the data on a replacement 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 operable).

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

When you have two controllers, these controllers are 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 Controller Unit continues to operate. You can then replace the failed controller, often without shutting down the RAID Controller Unit, 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 Controller 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 or not your RAID Controller Unit has the redundant controller feature depends on the following:

- Whether your operating system supports RDAC. Currently, RDAC runs on Windows NT Version 4.
- How the RAID controllers are connected. See "[Controller Configurations](#)" on page 13 for information on the possible configurations and their effect on redundant controller support.

Redundant controllers can be configured in two ways:

- Active/passive — in this configuration, one controller handles all the I/O from the host, and the other controller is held as a spare, taking over if the first controller fails.

- Active/active — in this configuration, both controllers handle I/O from the host. Each controller is assigned specific LUNs to service. If one controller fails, the other takes over the failed controller's assigned LUNs. This is the recommended configuration, because using two controllers for I/O is usually faster than using a single controller.

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 will write 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 while installing this software. It allows this software to send remote notification of RAID events to a designated network management station (NMS) using SNMP traps.

See the *Installation and User's Handbook* for your operating system for details on enabling this notification option.

Controller Configurations

The storage management software supports three main configurations from the host systems to the RAID Controller. Be sure to consult the *Installation and User's Handbook* 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/described in the operating-system specific *Installation and User's Handbook*, this *User Handbook*, or the on-line help if other configurations are used.

Most RAID Controller Units 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 Controller Unit. 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 User's Handbook*.

Single-Host Configuration

In a single host configuration, one host machine with two host adapters is connected by separate connections to each controller in the RAID Controller. The two connections are required for maximum RDAC failover support for redundant controllers (see [Figure 2 on page 15](#)).

NOTE This is the recommended configuration with the storage management software installed on the host for fullest functionality and complete RDAC failover support with dual controllers.

IMPORTANT Although this configuration also supports RAID Controller Units 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 in the operating system hanging. For the greatest level of I/O protection, provide each controller in a RAID Controller Unit with its own connection to a separate host adapter in the host system.




Figure 2. Single-Host to RAID Controller 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 Controller Unit (see [Figure 3 on page 16](#)).

CAUTION Not every operating system supports this configuration. Be sure to consult the restrictions in the *Installation and User's Handbook* specific to your operating system for more information. Also, the host machines and operating systems must be able to handle the multi-host configuration. Refer to the appropriate hardware documentation.

With the storage management software installed on each host machine, both hosts have complete visibility of both controllers, all data paths, and all configured drive groups/ logical units (LUNs) in a RAID Controller, 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 the two hosts do not send conflicting commands to the controllers in the RAID Controller.

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 16-32 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 (many UNIX systems can), the storage management software will honor them. This means that each host could have reservations to specified drive groups/logical units (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 LUNs, 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. That feature requires third-party software.




Figure 3. Multi-Host to RAID Controller Configuration

Independent Controller Configuration

In the Independent controller configuration, two host machines are connected to a dual-controller RAID Controller. One host machine is connected to one RAID Controller controller, and a second host machine is connected to the other controller (see [Figure 4 on page 18](#).) Each host machine and its storage management software see the controller and the drive groups/LUNs that it owns as independent of the other (alternate) controller. That is, each host machine acts as if it is connected to a single-controller RAID Controller.

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. It only reports real-time status for the host-controller data path on which it is installed, but displays and reports an “Inaccessible” status for drive groups/LUNs owned by the alternate controller.

IMPORTANT Because it is possible to open a Configuration application from both hosts, be sure that you have only one configuration session open at a time (from only one host) when creating and deleting LUNs 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 16-32 LUNs).
- This configuration is indicated by a special setting in Select Module (the “Indep. Cntrls?” column says Yes).
- The controllers in the RAID Controller Unit 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 path-related failures *only* for the controller that is connected to the host machine running the storage management software. For example, if host 1 has a controller failure, host 1 reports the failure, but host 2 will *not* report a controller failure using its Health Check. Also, these applications detect drive-related failures *only* for configured drive groups/LUNs that are owned by the controller (connected to the host machine running the storage management software) or for any unassigned or hot spare drive.

- Hot spares are always available to LUNs owned by either controller.

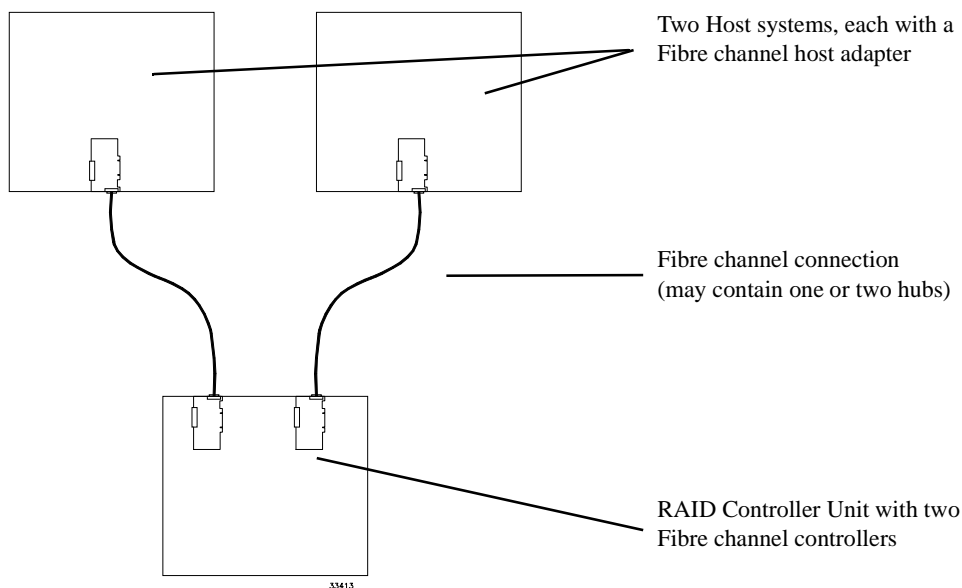


Figure 4. Independent Controller Configuration

Special Network Considerations

The Networked version of the storage management software always sees both controllers in a dual-controller RAID Controller Unit regardless of which configuration you have. However, the Networked version will be able to tell if the currently selected RAID Controller has an independent controller configuration *if* independent controllers was selected using the Select Module option.

If you are using the Networked version of this software, the following restrictions apply to any of the host-RAID Controller Unit configurations:

- Your RAID Controller Unit does not have RDAC failover protection unless there is failover protection installed on the host connected to the unit through the Fibre Channel connection.
- This software does *not* provide I/O connection failure detection or recovery. However, any problems with a loop connection to the controllers or a problem with the controllers themselves are shown as a data path failure. Recovery Guru/Health Check provides assistance for these problems.
- This software has no way to recognize any exclusive access operations that may be performed by other software installed on the host machine (not even another storage management package). This requires you to use caution before starting certain operations that need exclusive access because without it file systems are not detected

and multiple operations could be launched without logical units being protected. Furthermore, because it is possible to open a Configuration application from any networked station, be sure that you have only one configuration session open at a time when creating and deleting LUNs or the operation will fail for all but one station.

CAUTION Drive groups/LUNs and their data could be lost if more than one “destructive” operation is launched. No other operations should be attempted on the same drive group/LUN if one of these operations is still being completed. Operations requiring exclusive access to the LUNs include Delete for LUNs and File → Reset Configuration (Configuration); formatting a LUN with Options → Manual Recovery → Logical Units (Recovery); and Firmware Upgrade → Offline method (Maintenance/Tuning).

Component Failures

The RAID Controller Units serviced by the storage management software are designed to be fault tolerant. All units provide drive redundancy, so that if a drive fails, the data contained on the drive is still available. The units also feature redundant controllers, fans, and power supplies as well, so that the RAID Controller Unit can keep functioning even if a single component fails. The controllers 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 data loss in the event of a power failure. See the hardware books accompanying your RAID Controller Unit for complete information on the redundant features of your hardware.

The storage management software provides notification of most component failures. When you do replace the 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 books 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 book.

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, and a RAID 0 logical unit will lose 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. 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.
- When creating LUNs, try and 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 LUNs. Note that this may not be possible if you create LUNs with more drives than you have channels.

Application Summary Charts

The storage manager software consists of four separate applications, Configuration, Recovery, Status, and Maintenance and Tuning. These applications can be run independently of each other or concurrently.

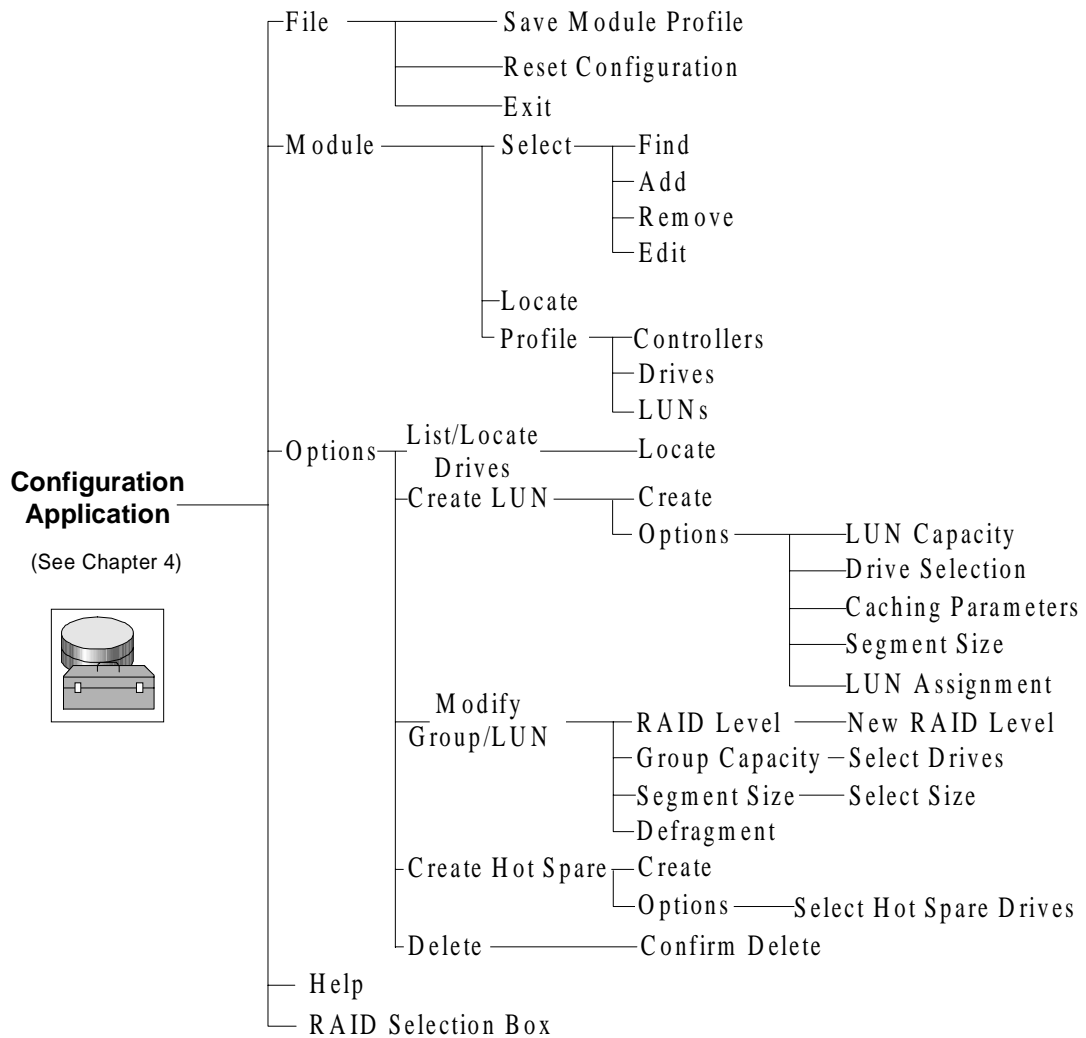


Figure 5. Configuration Application Summary

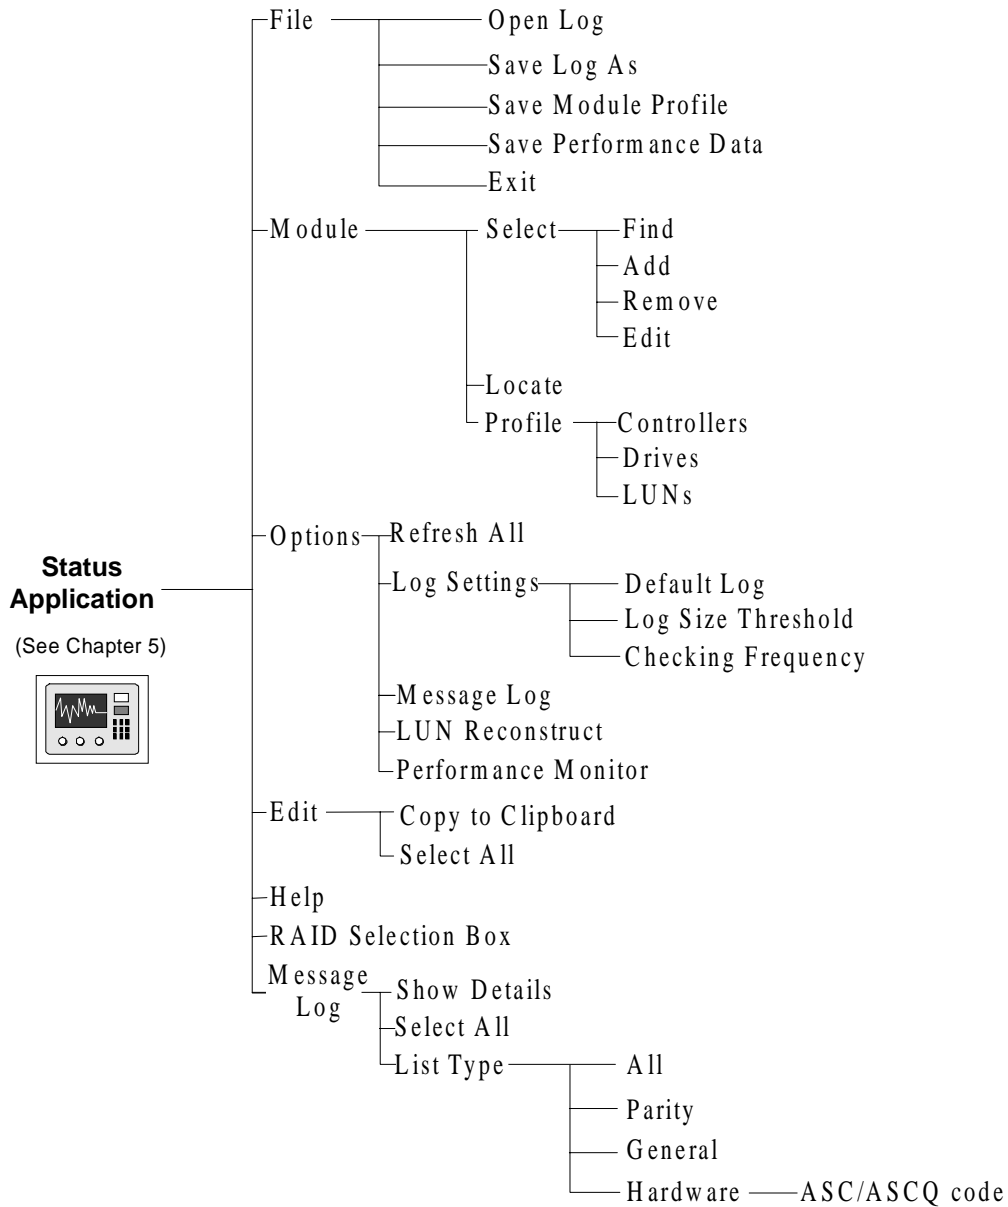


Figure 6. Status Application Summary

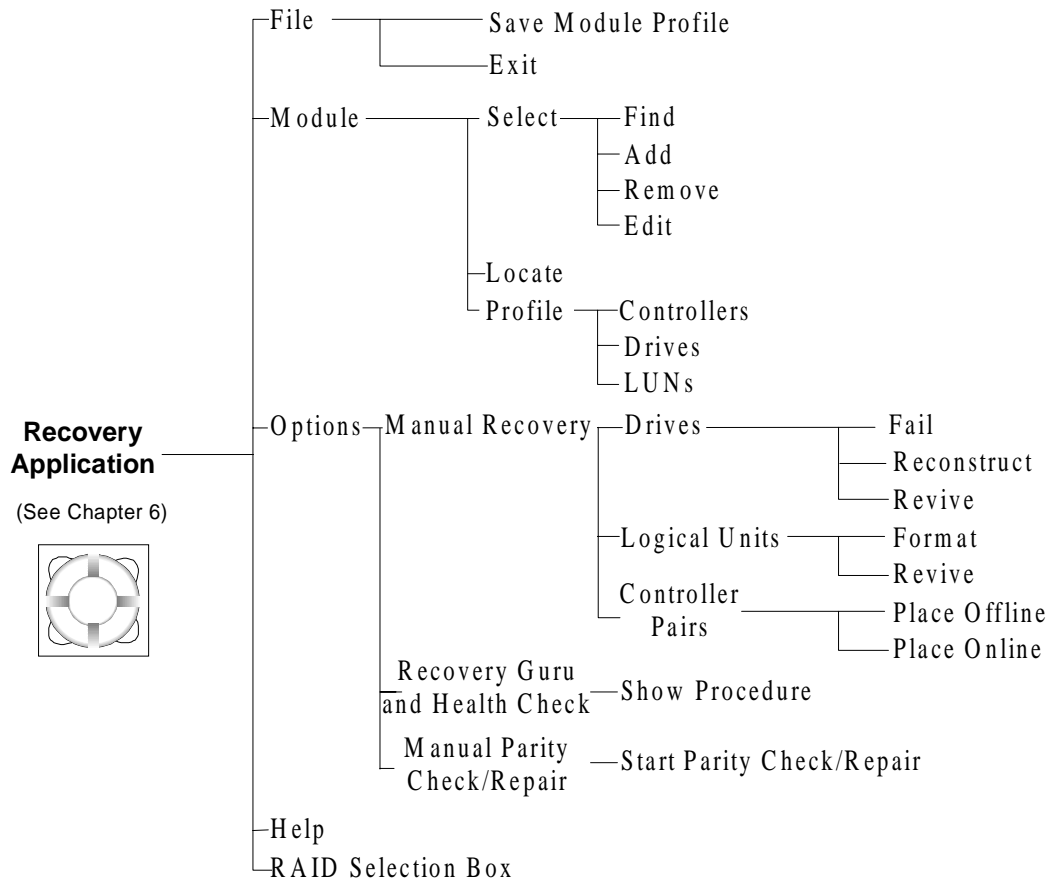


Figure 7. Recovery Application Summary

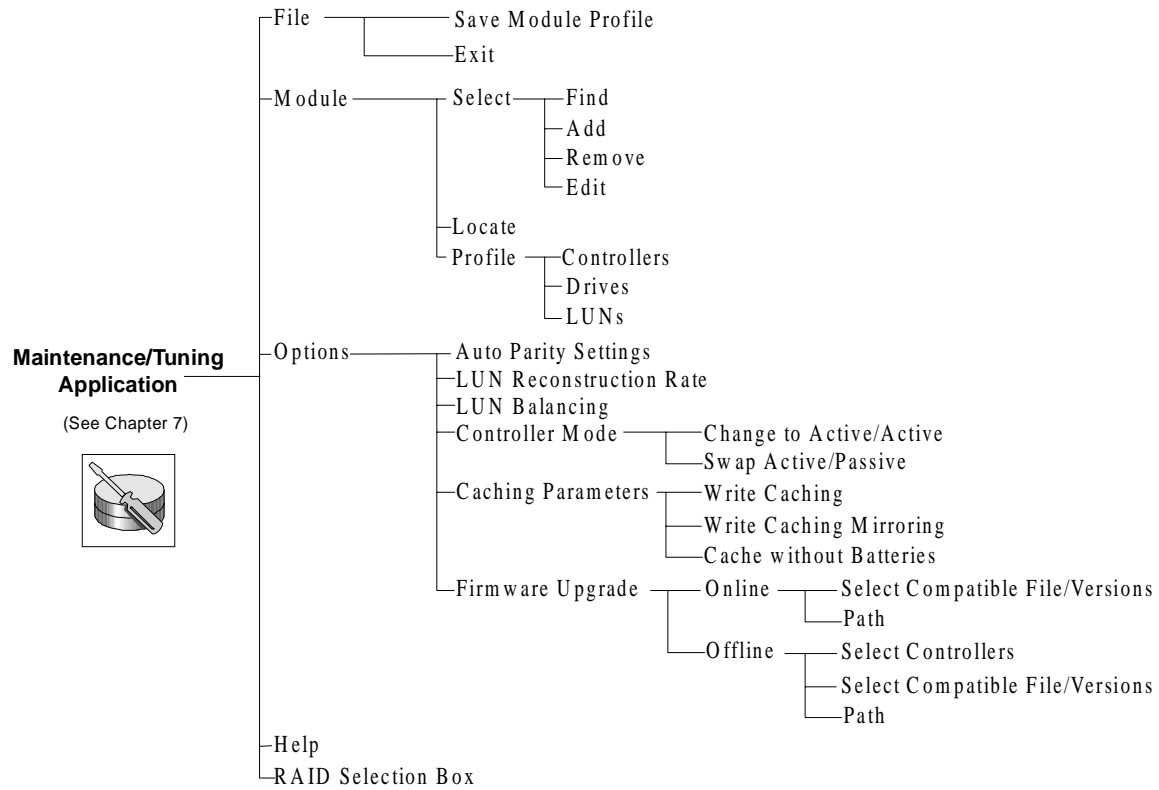


Figure 8. Maintenance and Tuning Application Summary

Task List

The table below lists the tasks you must perform to manage and support the RAID Controllers, and indicates when they need to be performed.

Table 1: Storage Management Task List

Task	When to Perform	Reference
Informational Tasks (Locating RAID Controllers, Checking Configurations)		
Locate a RAID Controller Module	When you want to physically locate a RAID Controller Module. Perform this task when you first install the module and mark the module for future reference.	"Locating a Module" on page 36
Locate a Logical Unit or Specific Drives	When you want to physically locate a drive or a logical unit.	"Locating Drives" on page 37
Assign Names to Modules and Controllers	To make the RAID Controller Module components easily identifiable. Only the networked version allows you to name controllers.	"Selecting a Module" on page 34
View RAID controller Configuration Information	When you need information on your RAID Controllers, such as current firmware version or drive information.	"Viewing a Module Profile" on page 38
Check the Status Log	When you want to see if any errors have occurred on the RAID Controller Module.	"Viewing the Log" on page 80
Run the Performance Monitor	To gather information on the performance of the RAID controller unit, with the goal of discovering any changes that could be made to speed performance.	"Using the Performance Monitor" on page 86
Configuration Tasks (Creating and Modifying Logical Units)		
Create Logical Units (LUNs)	After you install a RAID Controller Module.	"Creating a Logical Unit (LUN)" on page 54
Create a Hot Spare	When you configure a RAID Controller Module. You should always have at least one hot spare on your system to protect against failures.	"Creating a Hot Spare" on page 73
Add Drives to an Existing Drive Group	When you want to add capacity in a Drive Group	"Modifying Drive Groups and LUNs" on page 63

Table 1: Storage Management Task List (continued)

Task	When to Perform	Reference
Change RAID Level of an Existing Drive Group/Logical Unit	If you want a different configuration in your RAID Controller after you have created a LUN.	"Modifying Drive Groups and LUNs" on page 63
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 LUNs" on page 63
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 LUNs" on page 72
Delete a Hot Spare	If you need the hot spare drive for a new LUN.	"Deleting a Hot Spare" on page 75
Maintenance Tasks (Loading new firmware, changing parameters)		
Download Controller Firmware	<ul style="list-style-type: none"> • When you receive new firmware. • When adding RAID controllers with new firmware into an existing site (you will need to upgrade the existing controllers). 	"Upgrading the Controller Firmware" on page 113
Download NVSRAM	<ul style="list-style-type: none"> • When you receive new firmware. • When adding RAID controllers with new firmware into an existing site (you will need to upgrade the existing controllers). 	"Upgrading the Controller Firmware" on page 113
Reset the Configuration	If all attempts to use the RAID controller unit have failed. This is a last resort option that will delete all data and configuration information from your RAID Controller.	"Resetting the Configuration" on page 76
Change Controller Modes	If the controllers in a RAID Controller are currently active/passive and you want to change them to active/active.	"Changing the Controller Mode" on page 130
Balance Logical Units across controllers	<ul style="list-style-type: none"> • If you have changed controller mode from active/passive to active/active and didn't balance the LUNs at that time. • If you want to reassign logical units to controllers. 	"Balancing LUNs" on page 127
Change Cache Parameters	To turn cache on and off after you have created a logical unit.	"Changing Cache Parameters" on page 133

Table 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.	"Changing the LUN Reconstruction Rate" on page 125
Set the Time for Automatic Parity Check/Repair	If you want to change the default time automatic parity check is performed. You might want to do this if the current time is interfering with system operations.	"Setting Automatic Parity Time" on page 122
View the Results of Automatic Parity Check/Repair	If you want to see if automatic parity check/repair detected any parity errors.	"Viewing the Log" on page 80
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. The Health Check will locate the problem and give information on how to correct it.	"Running Recovery Guru/Health Check" on page 100
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"
Check Data Parity	If Recovery Guru tells you to run parity check after a component failure, or you have indications that I/O errors have occurred.	"Running a Parity Check" on page 122

Chapter 3. Common Application Features

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

New Software Features

This release of the storage management software contains the following new features:

- **Faster LUN creation.** The controller now performs background formatting to create logical unit numbers (LUNs) faster.
- **SMART/PFA support.** Through the use of Self-Monitoring Analysis and Reporting Technology (SMART), the storage management software now reports predictive failure analysis (PFA) flags on drives.
- **Firmware autosynch.** When you replace a controller with firmware level 3.00.X or higher, the firmware on the new controller automatically synchronizes with the firmware on the remaining controller.
- **Parity Check/Repair Enhancement.** You now can choose to not repair errors during automatic parity check/repair.

For additional information about the new features for this release, see your *SYMplicity Storage Manager Installation and User's Guide*.

Navigation

This software requires that you use a mouse for fullest functionality; however, you can also use your keyboard to access the task options. [Table 2](#) describes how to use the mouse and keyboard.

Table 2: Using a Mouse and Keyboard

When using a mouse...	<ul style="list-style-type: none">• 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 mouse over the appropriate button and read the description near the bottom of the screen.• To highlight items, do one of the following:<ul style="list-style-type: none">• Single-click to highlight a single item.• Press <Shift> + click to highlight a series of items. For example, single-click to highlight the top item in a list, then press <Shift> + click on the last item in the list to highlight all the items in that list.• Press <Control> + click to highlight items not in a series. For example, single-click to highlight one item in a list, then press <Control> + click on another item to highlight it as well. Do this for every item you want to highlight.
When using a keyboard...	<ul style="list-style-type: none">• 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 key to highlight Save Module Profile and press <Enter>.

Starting and Exiting an Application

The following sections explain how to start and exit an application.

Starting an Application

Select the application ([Figure 9](#)) you want to start.

- The first screen you see will prompt you to select the RAID Controller you want to access (see "[Selecting a Module](#)" on page 34).

- If the message log file is currently over its maximum size limit, you will be prompted to rewrite the file (see ["Clearing the Message File" on page 85](#)).




Figure 9. Application Icons

Exiting an Application

To exit any application, select File → Exit from the top menu. 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. Fortunately, 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 Handbook*. 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 10. Table 3 on page 32](#) details the features common to Online Help.

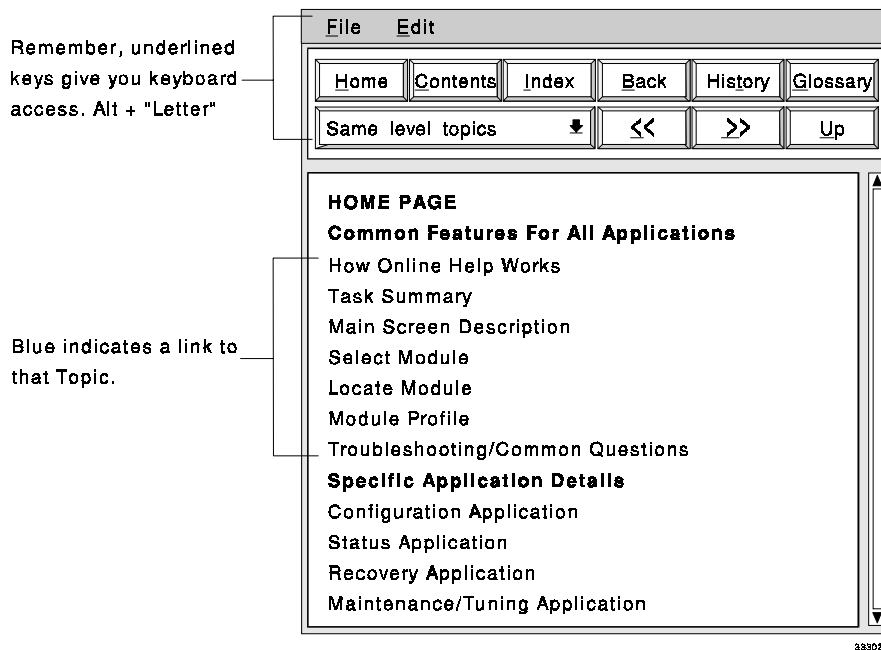


Figure 10. Main Online Help Screen

Table 3: Online Help Screen Elements

Selection	Description
File	Allows you to do the following: <ul style="list-style-type: none"> • Print the currently displayed topic to a file or printer. • Set up your printer (landscape/portrait, margins, etc.). • Exit Online Help.
Edit	Copies text to a 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.

Table 3: Online Help Screen Elements (continued)

Selection	Description
Contents	<p>Displays all the help topics organized by hierarchy and appearance on the Home Page.</p> <p>Press a letter to quickly advance through the list of topics. For example, pressing <M> will take 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.</p>
Index	<p>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.</p> <p>Press a letter to quickly advance through this alphabetical list. For example, pressing <M> will take 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>.</p>
Back	<p>Takes you back (one topic at a time) through the topics you have viewed since selecting Help.</p>
History	<p>Creates a list of all topics you view in the order you have selected them.</p> <p>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.</p>
Glossary	<p>Displays an alphabetical list of defined terms.</p> <p>Press a letter to quickly advance through this alphabetical list. For example, pressing <M> will take 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. You can use the <Home> and <End> keys on your keyboard to quickly move to the beginning and end of the glossary.</p>
Same level topics	<p>Displays topics of the same level using the << and >> keys to move forward or backward.</p> <p>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.</p>
<< and >>	<ul style="list-style-type: none">• 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	<p>Moves you to the next higher level of topics.</p>

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, such as configuration or recovery, on it.

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

RAID Controller Modules are identified either by a device name, assigned by the operating system, or by a user-specified name. You can name the module through the Module Selection screen (see the next section).

There are several ways to select a RAID Controller Module:

- If the correct parameter is turned on in your rmparms file (it is turned on by default), you are given the chance to select a RAID Controller Module when you first start an application (see [Figure 11 on page 35](#)). Either highlight a module and select OK or double-click on the module you want.
- Once in an application, you can:
 - Select the RAID Controller Module you want from the RAID Controller 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 tree (in the Configuration application only).

In all cases except when selecting a module from the Module drop-down list box, the RAID Controller Module selection screen is displayed. See "[RAID Controller Module Selection Screen](#)" on page 35 for details on that screen.

RAID Controller Module Selection Screen

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




Figure 11. RAID Controller Module Selection Screen

Notes:

- 1 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 17](#)), and displays any comments you may have entered with the Edit option.
- 2 [Table 4 on page 36](#) gives more information on the buttons you can select.

Table 4: Select Module Screen Elements

Screen Element	Description
Find	<p>Allows you to quickly locate a RAID Controller. It will be most useful when you have many modules.</p> <p>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.</p>
Add	<p>Applicable for the Networked version only. Allows you to add new modules to your system so that this software can access and monitor it.</p> <p>See the <i>Installation and User's Handbook</i> specific to your operating system for details on adding new modules to your system.</p>
Remove	<p>Allows you to remove RAID Controllers from your system.</p> <p>You can only remove one module at a time. This option will be grayed out if All RAID Controllers is selected.</p> <p>IMPORTANT</p> <ul style="list-style-type: none">• 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 Controllers from your system, but do not use this Remove option, the storage management software will continue trying 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	<p>Allows you to assign a name to the RAID Controllers and to enter comments about it. Use the comments area to provide detailed information about the RAID Controller (such as location information, independent controllers, etc.) to help you identify it.</p> <p>On Networked versions, you can also assign names to the RAID Controller controllers.</p> <p>You can only edit information for one module at a time. This option will be grayed out if All RAID Controllers is selected.</p>

Locating a Module

The storage management software allows you to physically locate a RAID Controller Module. This is useful in trying to physically locate the 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 Drives" on page 37](#)).

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 activity light's normal blinking during I/O activity.

To locate a RAID Controller Module, do the following.

1. Start the application you want to perform.
2. Select the Module you want to locate. Do not select All Modules, because this will gray 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 activity lights on the selected module begin to blink.
5. After you have physically located the module, click Stop and return to the current application. This is a good time to label the module for future reference.

Locating Drives

The Configuration application allows you to physically locate individual drive groups/logical units and drives. Drives are identified by channel number and SCSI ID. See your hardware documentation for more information on drive numbering.

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

1. Start the Configuration application.
2. Select the RAID Controller Module containing the logical units or drives you want to locate.
3. Click on the drive group or drives you want to locate in the directory tree.
4. Click on the List/Locate Drives button or select Options → List/Locate Drives from the drop-down menus.
5. A screen like [Figure 12](#) appears. 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 activity lights on the selected drives begin to blink.
7. After you have physically located the drives, click Stop and return to the Configuration application.

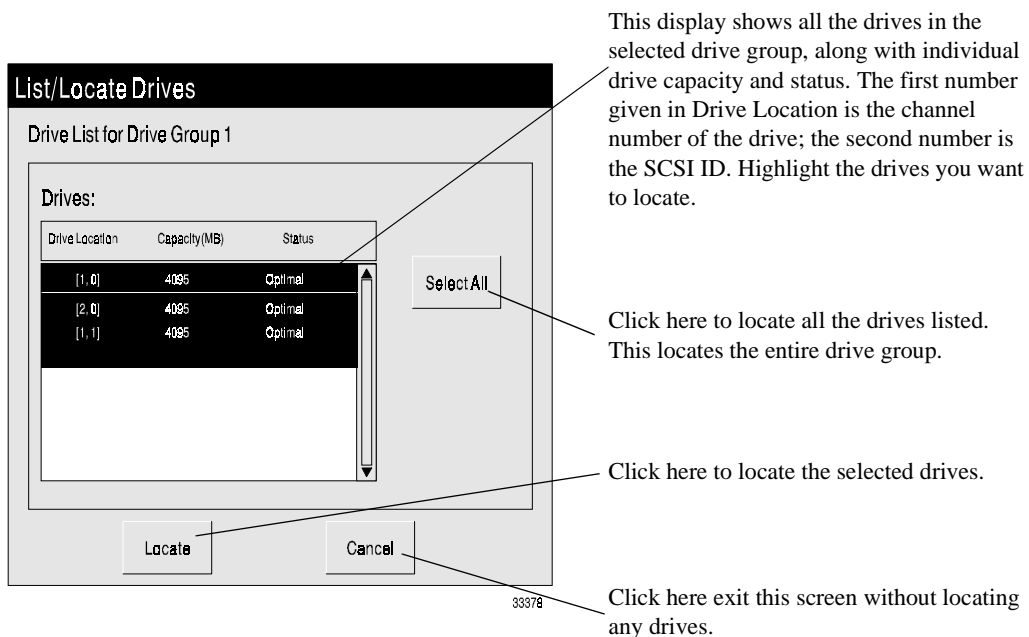


Figure 12. List/Locate Drives Screen

Viewing a Module Profile

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

To display a module profile, do the following:

1. Start any application (Configuration is the recommended application).
2. Select the module you want information on. Do not select All Modules (the Module Profile option is grayed out if All Modules is selected).
3. Click on the Module Profile button or select Module → Profile from the drop-down menus. The Module Profile screen is displayed (see [Figure 13](#)).
4. After you have viewed all the profile information you want, click OK to return to the application.

- After exiting Module Profile, you can save the profile information to a file. See ["Saving a Module Profile" on page 44](#).

See [Table 5](#) for an explanation of the items in this display.

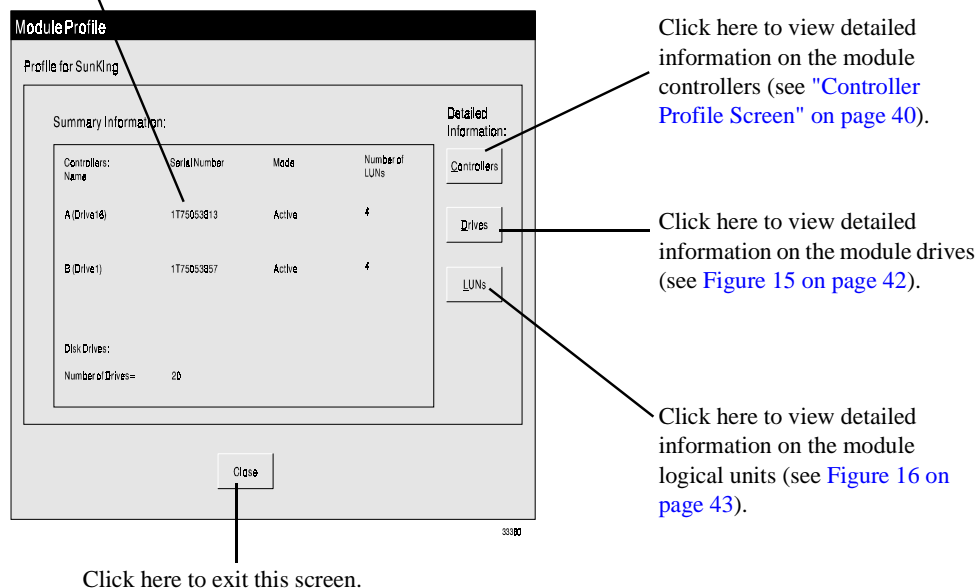


Figure 13. Module Profile Screen

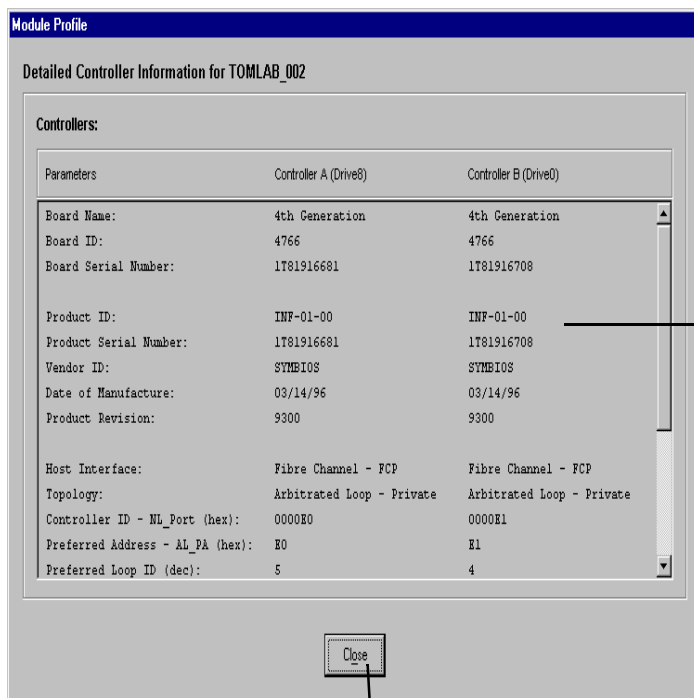
Table 5: Module Profile Screen Contents

Screen Element	Explanation
Name	Identifies the controllers in the selected RAID Controller 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 Controller Module has an independent controller configuration. IMPORTANT If you do not see Mode information or other information in this screen is incomplete, there may be a connection problem. Select Recovery Guru and correct any problems indicated.
Number of LUNs	Indicates how many logical units are owned by the particular controller.
Disk Drives	Indicates how many drives make up the selected RAID Controller Module.

Controller Profile Screen

This section describes the Fiber Channel Controller Profile screen.

Fibre Channel Controller Profile



See [Table 6](#) for an explanation of the items in this display.

Click here to exit this screen.

Figure 14. Controller Profile Screen

Table 6: Controller Profile Screen Contents (Fibre Channel)

Screen Element	Explanation
Board Name	Controller type designation.
Board ID	Controller model number (4766).
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.

Table 6: Controller Profile Screen Contents (Fibre Channel) (continued)

Screen Element	Explanation
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. <ul style="list-style-type: none">• For Arbitrated Loop, this is a one-byte Arbitrated Loop Physical Address.• For Fabric and Point-to-Point, this is a three-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 an 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 (optical).
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).

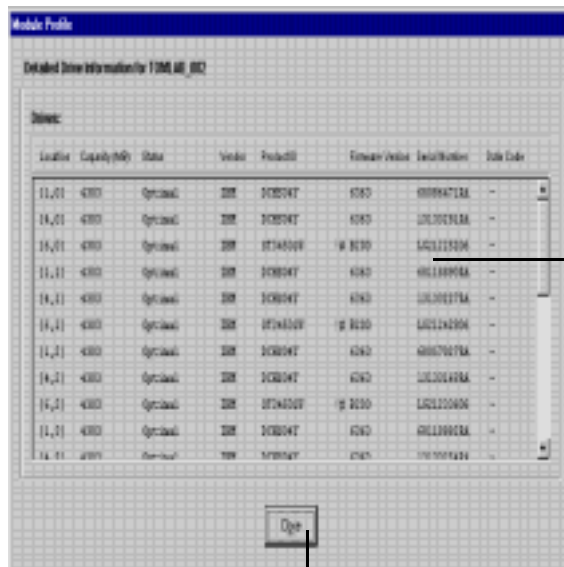
Table 6: Controller Profile Screen Contents (Fibre Channel) (continued)

Screen Element	Explanation
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” in a column means that the software can not return the indicated value. This is not necessarily an error; it is usually a limitation on the fibre channel firmware.

Drive Profile Screen

Figure 15 shows the Drive Profile screen.



See [Table 7 on page 43](#) for an explanation of the items in this display.

Click here to exit this screen.

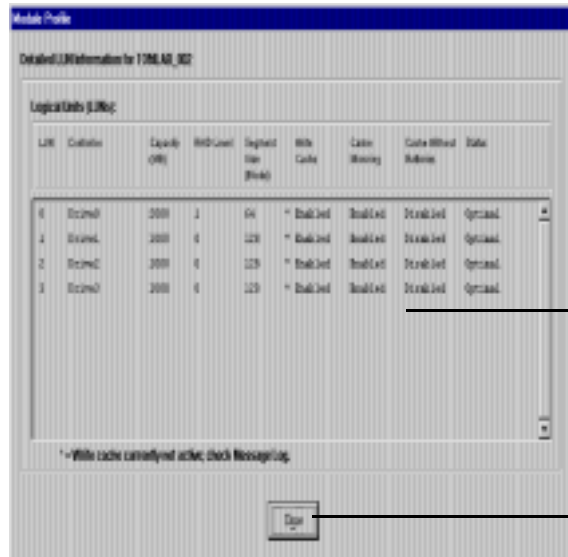
Figure 15. Drive Profile Screen

Table 7: Drive Profile Screen Contents

Screen Element	Explanation
Location	Designation indicating the unique location of the drive in the selected RAID Controller Module. This information is displayed as [x,y] on screen where the channel number is always listed first. For example, [2,1] corresponds to the drive at SCSI Channel 2 and SCSI ID 1.
Capacity (MB)	Amount of storage space on the drive (in megabytes).
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.

LUN Profile Screen

Figure 16 shows the LUN Profile screen.



See [Table 8 on page 44](#) for an explanation of the items in this display.

Click here to exit this screen.

Figure 16. LUN Profile Screen

Saving a Module Profile

Table 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).
RAID Level	Indicates the way the controller reads and writes both data and parity on the drives. Possible 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 option has been enabled for a particular LUN.
Cache Mirroring	Indicates whether the cache mirroring option has been enabled for a particular LUN.
Cache Without Batteries	Indicates whether the cache without batteries option 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.

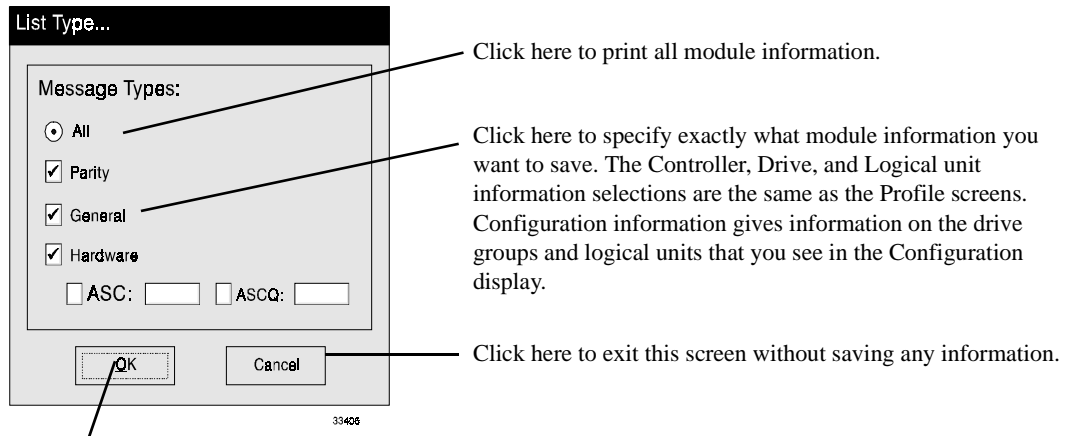
NOTE: You might see an asterisk next to the caching parameters column. This indicates that the parameter is enabled, but is currently *not* active. The controller has disabled the parameter for some reason (such as low batteries). If you see this condition, use Recovery Guru to determine the correct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

You can save module profile information for future reference. The module information is stored in a text file. It can not be used to restore module settings, but it 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 steps:

1. Start any storage management application.
2. Select the RAID Controller you want to save profile information for. Do not select All RAID Controllers (this option is grayed out if All RAID Controllers is selected).

3. Select File → Save Module Profile from the drop-down menus. A screen like [Figure 17](#) is displayed.



Click here to save the information selected. You then must assign a name to the file to save it.

Figure 17. 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 and location to the file. Enter the name and click OK.

Chapter 4. Configuring RAID Controllers

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

Configuration Overview

IMPORTANT Before you attempt to configure a RAID Controller, make sure you understand the concepts of RAID Controller, Drive Group, Logical Unit, and RAID Level. See ["Common Definitions" on page 5](#) for information on these important concepts, which will be needed to perform the tasks in this section.

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

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

Types of Logical Unit (LUN) Creation

There are two types of LUN creation:

- Creating a new LUN from the unassigned drives in the RAID Controller. 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 54](#).
- Creating a new LUN from remaining capacity in a drive group. In this configuration mode, you cannot assign a new RAID level or new drives, as all the LUNs in a drive group must use the same RAID level and drives. See ["Creating a Logical Unit from Remaining Capacity" on page 57](#).

Types of Logical Unit (LUN)/Drive Group Modification

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

- You change the RAID level of the LUNs in the drive group. This change applies to all LUNs in the drive group. See ["Changing the RAID Level" on page 67](#).
- You can modify LUN parameters (such as segment size, caching, reconstruction rate and so on). See:
 - ["Modifying Segment Size" on page 71](#)
 - ["Changing Cache Parameters" on page 133](#)
 - ["Changing the LUN Reconstruction Rate" on page 125](#)

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

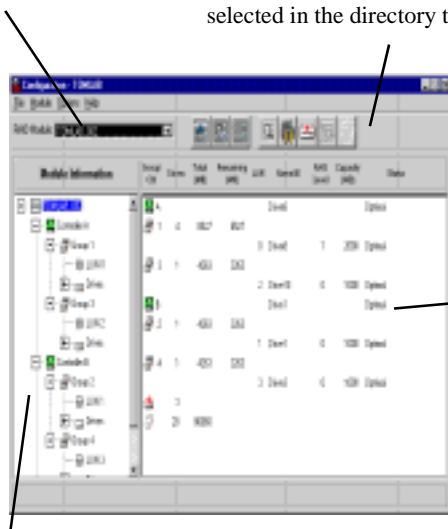
- You can increase Drive Group capacity by adding new drives to the drive group. This does not increase the size of the LUNs in the drive group. See ["Adding Drives" on page 69](#).
- You can change the controller assignment of the LUN/drive group. See ["Balancing LUNs" on page 127](#).

Module Information Display

The Configuration application displays module information as shown in [Figure 18](#)

This is the currently selected RAID Controller.

These buttons are active or grayed out depending on what RAID Controller structure is currently selected in the directory tree.



This gives information on the drive groups and logical units in the RAID Controller. The items displayed here depend on what is selected in the directory tree display to the left.

The display shown here is the master display, shown when the RAID Controller itself is selected.

See [Table 10 on page 54](#) for a description of the elements in the display.

This shows the RAID Controller as a directory tree structure. At the top is the RAID Controller, then the controllers, then the drive groups owned by the controller, then the logical units and drives in that drive group. At the bottom of the display (not shown here) are the hot spares and unassigned drives. Clicking on the different structures changes the display on the right of the screen.

Figure 18. Configuration Application Module Information Display

Notes:

- 3 Click on the “+” or “-” signs to expand or collapse the tree.
- 4 Clicking on an element in the tree display changes the information displayed in the right half 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.
- 5 You cannot click on items in the right-hand display.

Logical Unit Formatting

Whenever you create a logical unit (LUN) of more than 20 MB, the RAID controller formats (writes zeros to) only the first 10 MB and the last 10 MB of the LUN; then, the controller reports the format operation as completed and the LUN status becomes Optimal. At this point, you can use the logical unit, but some restrictions apply. Before you start writing data to the LUN, review the information provided in ["Formatting Restrictions" on page 51](#). After you create the logical unit, the RAID controller will complete the format operation in the background, depending upon the assigned RAID level, as follows:

- RAID 0 logical unit: The RAID controller will perform no further format actions. Data is written to the unformatted portion of the logical unit normally.
- RAID 1 logical unit: The RAID controller will format the remaining portion of the logical unit in the background. Stripe by stripe, the controller will rewrite 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 will write the data normally, so that the contents of each mirrored pair is the same.
- RAID 3 or 5 logical unit: The RAID controller will format the remaining portion of the logical unit in the background. One at a time, the controller will rewrite the contents of each stripe so that the parity is correct in each stripe (see [Figure 19, "RAID 5 LUN Creation Example"](#)). If the controller needs to write user data in the unformatted portion of the LUN, it will write the data normally, so that the contents of each stripe is correct.

Only the first and last 10 MB in the logical unit are formatted when you create the LUN. The LUN is then available for use. The controller completes the format operation in the background if the LUN is RAID 1, 3, or 5.

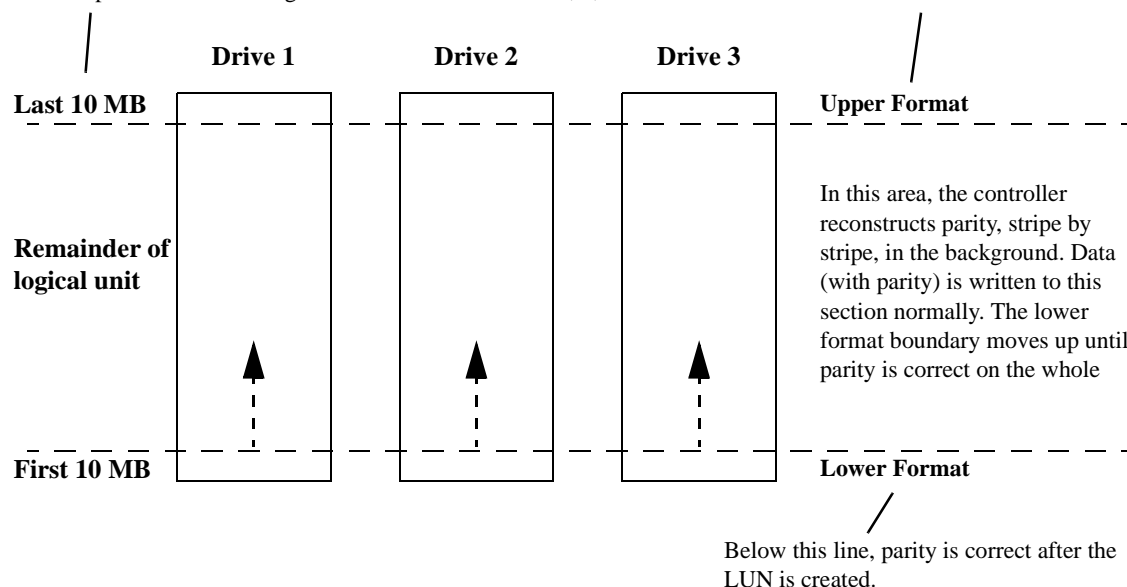


Figure 19. RAID 5 LUN Creation Example

While the controller is formatting in the background, there is a *boundary* in the LUN. Below the boundary, the parity or the mirroring in the logical unit is correct. Above 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 LUN Reconstruction Rate" on page 125](#)). Depending upon the size of the LUN and the amount of I/O activity taking place, the background format operation can take up to 8 hours. In addition, the system might run slower during a background format operation because a reconstruction process is taking place. Again, this will depend upon the amount of I/O activity taking place and the Reconstruction Rate parameter settings.

A LUN being formatted in the background has an Optimal status. There is no indication in the storage management software that the background format operation is in progress, unless you attempt to perform a restricted operation as listed in ["Formatting Restrictions" on page 51](#). Aside from the restricted operations listed, you can use the LUN normally while the background format operation is in progress.

If you need to determine whether a LUN is still formatting in the background, try performing a modification on the affected drive group, such as changing the segment size (see ["Modifying Segment Size" on page 71](#)). If a background format operation is still in

progress, a message will appear stating that the LUN is still initializing. If the background format operation has completed, the normal Modify Screen will appear. You can then cancel out of that screen without making any modifications.

Formatting Restrictions

You can not perform the following operations on a logical unit while it is formatting in the background:

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

The background format operation will stop if the logical unit being formatted becomes degraded during the format operation. (In a RAID 1 LUN, the formatting stops only if each mirrored pair has a failed drive. In this case, the logical unit is still degraded.) When you replace the failed drive, the entire logical unit will be reconstructed before the status of the logical unit becomes Optimal.

The background format operation also will stop if the logical unit fails during the format operation. If this occurs, the logical unit must be reformatted as part of a recovery operation.

Configuration Application Main Screen

Figure 20 shows the Configuration application Main Screen.

This is the currently selected RAID Module.

These buttons are active or grayed out depending on which RAID Module structure is currently selected in the Module Information area to the left.

Group/Ctrl	Drives	Total (MB)	Remaining (MB)	LUN	Name/ID	RAID Level	Capacity (MB)	Status
A	1	8111	8011	0	Drive16	5	100	Optimal
	3	8642	8442	2	Drive16	1	200	Optimal
	5	8111	8036	4	Drive20	3	75	Optimal
	7	17284	16634	6	Drive22	5	500	Optimal
	2			7	Drive23	5	150	Optimal
B	2	8111	7881		Drive1			Optimal
	4	8111	7886	1	Drive1	3	150	Optimal

The Detailed Information area gives information on the drive groups and logical units in the RAID Module. The items displayed here depend on what is selected in the Module Information area to the left.

The information shown in this example is displayed when the RAID Module itself is selected.

See Table 10 on page 54 for a description of the elements in this area.

The Module Information area shows the RAID Module as a directory tree structure. At the top is the RAID Module, then the controllers, then the drive groups owned by the controller; then the logical units and drives in that drive group. At the bottom of the display (not shown here) are the hot spares and unassigned drives. Clicking on the different structures changes the display in the Detailed Information area on the right.

Figure 20. Configuration Application Main Screen

Notes:

- 1 Click on the plus (+) or minus (-) signs to expand or collapse the Module Information area directory tree.
- 2 Clicking on an element in the Module Information area changes the information that appears 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 appears on the right. If you click on the unassigned drive group, information on the drives in the unassigned group (drive ID and capacity) appears on the right.
- 3 You cannot click on items in the Detailed Information area.

Table 9: Configuration Application Main Screen Elements

Column Heading	Explanation
Group/Ctl	The RAID 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	<p>The name or ID of the controller, logical unit, or drive, depending on the display.</p> <ul style="list-style-type: none">• 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 Controller Module Selection Screen" on page 35).• Drive IDs are the channel number and SCSI ID of the drives. See "Drive Names" on page 8 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 a Logical Unit (LUN)

Table 10: Configuration Display Columns

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	<p>The name or ID of the controller, logical unit, or drive, depending on the display.</p> <ul style="list-style-type: none">• 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 Controller Module Selection Screen" on page 35).• Drive IDs are the channel number and SCSI ID of the drives. See "Drive Names" on page 8 for more information.
RAID Level	The RAID level of the logical unit.
Capacity (MB)	The capacity of the logical unit.
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.

You can create a new LUN in two ways:

- By creating a new LUN from unassigned drives (see ["Creating a Logical Unit from Unassigned Drives"](#)).
- By creating a new LUN from capacity remaining in a drive group (see ["Creating a Logical Unit from Remaining Capacity" on page 57](#)).

Creating a Logical Unit from Unassigned Drives

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

Use the following steps to create a LUN from unassigned drives.

1. Start the Configuration application.
2. Select the RAID Controller on which you want to create the new LUNs.
3. Select the RAID Controller or the Unassigned Drive group and click on the Create LUN button. A screen like [Figure 21 on page 56](#) 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 3 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 change the value by selecting Options. 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 here, 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 Controller. Each LUN created gets an equal amount of the displayed capacity. For example, if you specify two LUNs, each LUN gets half the displayed capacity. You can change the capacity value by selecting Options.
7. The new drive group will automatically be assigned to a controller if the RAID Controller has redundant controllers. The first drive group you create in a configuration session will be 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 will be assigned to the same controller. You can change controller assignment by selecting Options. To balance LUNs after you create them, see ["Balancing LUNs" on page 127](#).
8. 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 59](#) for a description of those screens.

IMPORTANT If your RAID Controller contains drives of different sizes, you should use the Options → Drive Selection option 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 in the LUN, and capacity will be wasted.

Select one of these RAID buttons to assign a RAID level to the new LUN. RAID level helps determine the possible capacity and Number Of Drives settings.

Click here to start LUN creation with the current parameters.

Information describing the currently selected RAID level is displayed here. Available Group Capacity gives the capacity of the group after considering the RAID level selected.

Specify the number of drives in the new LUN here. The number of drives you may assign depends on the RAID level selected and on the number of unassigned drives available. Changing this value may change the selected RAID level.

Select the number of new LUNs to create here. This value may be limited by the number of LUNs available on your system.

This part of the screen gives help information

Click here to exit without creating a LUN.

Click here to view or change LUN options (capacity, drives, caching, segment size, and controller assignment). See [Figure 23](#) through [Figure 27](#) for examples of those options.

Figure 21. Create LUN Main Screen (Unassigned Drives)

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

IMPORTANT If you change any parameters in the Create LUN screen after making changes in the Options screens, any Options changes are negated.

10. The new LUNs are created.

NOTE If you are creating the *first* LUN on the module, wait for the create/format 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.

11. After the LUNs are created, you need to do whatever is necessary on your operating system to make the LUNs available for use. This might involve restarting the system, creating file systems and volumes, and so on. Check your operating system documentation and the *Installation and User's Handbook* for your operating system for more information.

Creating a Logical Unit from Remaining Capacity

To create a new 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.

Use the following steps to create a LUN from the remaining capacity in a drive group.

1. Start the Configuration application.
2. Select the RAID Controller on which you want to create the new LUNs.
3. Select the Drive Group in the Module directory tree and click on the Create LUN button. A screen like [Figure 22 on page 58](#) 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 Controller. By default, each LUN created gets 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 59](#)).
6. After you have set all the LUN parameters you want, start LUN creation by clicking Create in the Create LUN screen.

IMPORTANT If you change any parameters in the Create LUN screen after making changes in the Options screens, any Options changes are negated.

7. The new LUNs are created.
8. After the LUNs are created, you need to do whatever is necessary on your operating system to make the LUNs available for use. This might involve restarting the system, creating file systems and volumes, and so on. Check your operating system documentation and the *Installation and User's Handbook* for your operating system for more information.

This area displays the RAID level, number of drives, and available group capacity for the drive group. You cannot change any of these parameters.

Set the number of LUNs you want to create here. This may be limited by the total number of LUNs available on your system.

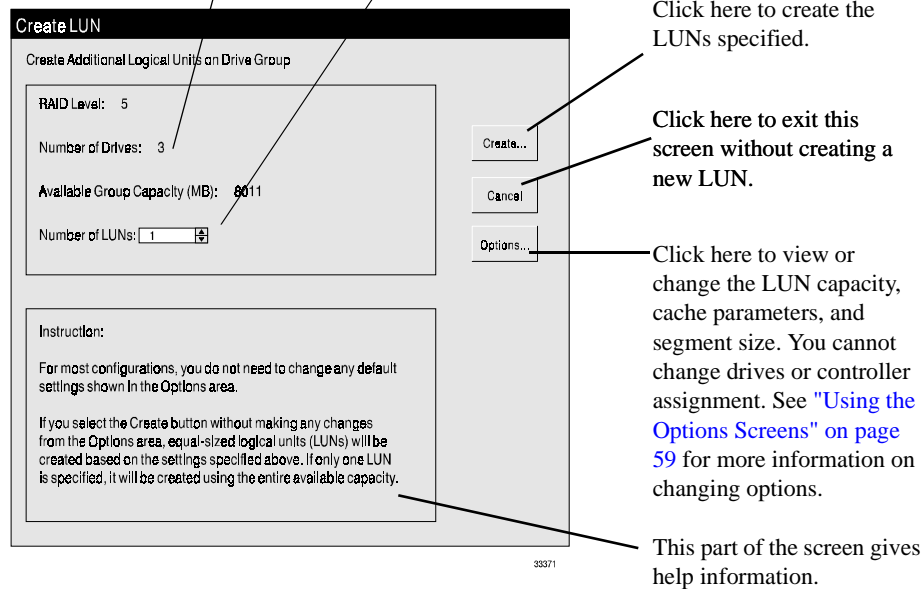


Figure 22. Create LUN Main Screen (Existing Drive Group)

Notes:

- 1 You cannot change the RAID level or drives in the LUN.
- 2 The new drive group will automatically be assigned the controller that currently owns the 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 23 on page 59](#)).
- The drives in the LUN (see [Figure 24 on page 60](#)).
- The caching parameters (see [Figure 25 on page 61](#)).
- The segment size (see [Figure 26 on page 62](#)).
- The controller that “owns” the new LUN. This option is allowed only if your RAID Controller has active/active redundant controllers (see [Figure 27 on page 63](#)).

NOTE When you move from screen to screen, you automatically save the values displayed in the current screen.

LUN Capacity Option

Use the LUN Capacity 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 2 logical units, each is assigned half the total capacity.

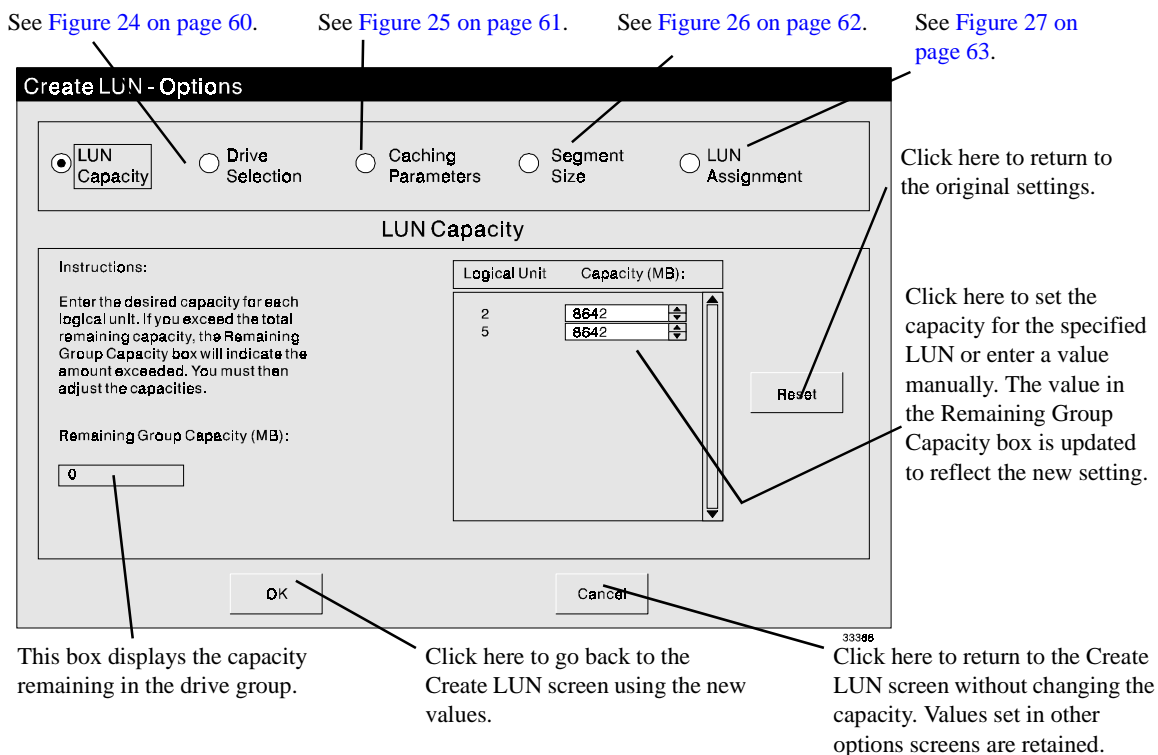


Figure 23. LUN Capacity Screen

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

Drive Selection Option

Use the Drive Selection option to specify the drives assigned to the new LUN/drive group. In particular, you should make sure that all the drives in the new LUN/drive group 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.

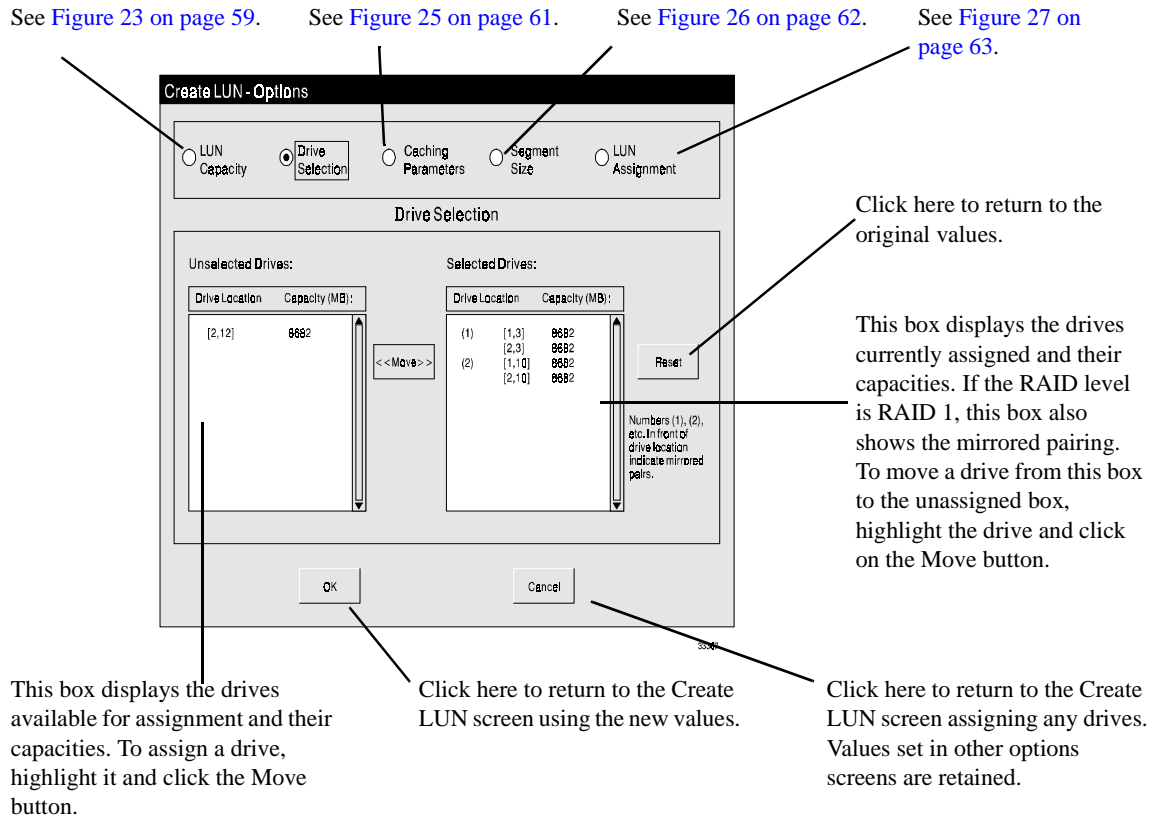


Figure 24. Drive Selection Screen

NOTE You cannot exit this screen if you have specified more drives than you entered in the Create LUN Main screen ([Figure 21 on page 56](#)).

Cache Parameters Option

Use the Cache Parameters option to change the cache parameters assigned to the new LUNs. Parameters are assigned on a LUN basis. See ["Cache Memory" on page 5](#) and ["Changing Cache Parameters" on page 133](#) for more information on the cache parameters.




Figure 25. Cache Parameters Screen

Notes:

- 1 See [Table 19 on page 134](#) for a description of cache parameter interdependencies.
- 2 Write-cache mirroring is effective only for RAID modules with redundant controller pairs (active/active or active/passive) that have the same size cache.
- 3 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 cache is also enabled
 - If you deselect write caching, write-cache mirroring and cache without batteries are also deselected (if currently enabled)

Segment Size Option

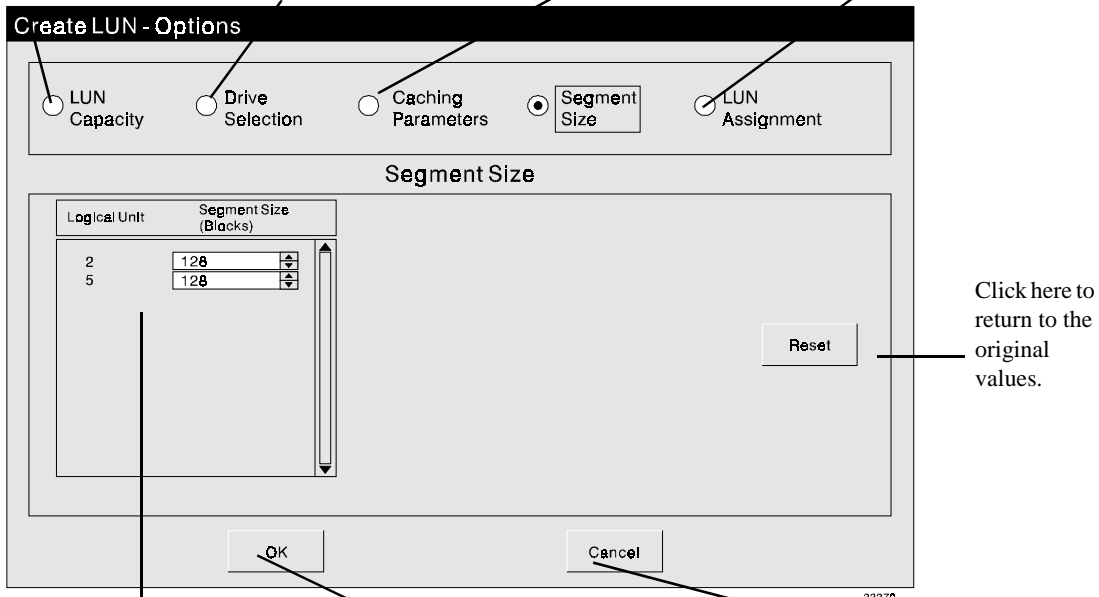
This option changes the segment size used by the new LUNs.

See Figure 23 on page 59.

See Figure 24 on page 60.

See Figure 25 on page 61.

See Figure 27 on page 63.



Click here to return to the original values.

The current segment size for each LUN is displayed here. Use the spinner to increase or decrease this size.

Click here to return to the Create LUN screen using the new values.

Click here to return to the Create LUN screen without changing segment sizes. Values entered in other options screens are retained.

Figure 26. Segment Size Screen

Notes:

- 1 You can set the segment size for individual LUNs; LUNs in a drive group do not need to have the same segment size value.
- 2 You can change the segment size of a LUN after you create it; see "[Modifying Segment Size](#)" on page 71.

LUN Assignment Option

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

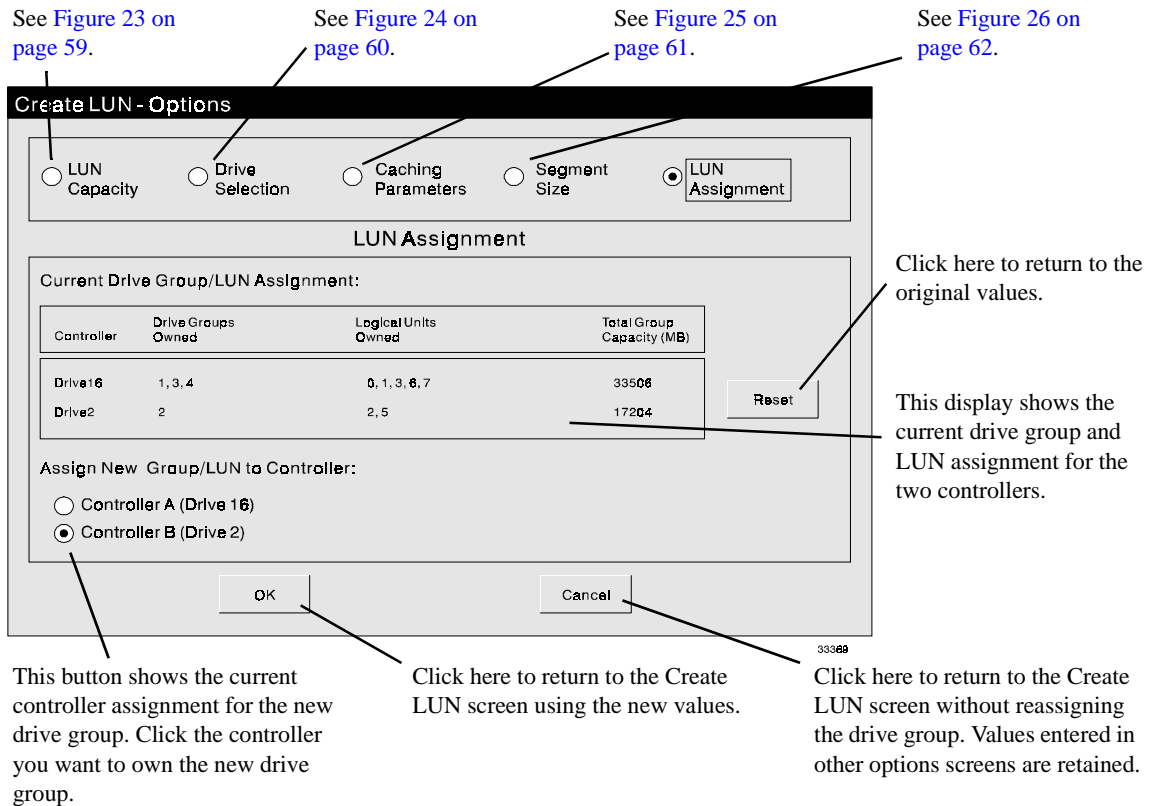


Figure 27. Set LUN Assignment Screen

Modifying Drive Groups and LUNs

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

- Change the RAID level — you can dynamically 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 ["Changing the RAID Level" on page 67](#).
- 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 increase depends on the number of drives added and the RAID level (see [Table 11 on page 69](#)). See ["Adding Drives" on page 69](#).
- 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 71](#).

- Defragment the logical units on the drive group — this may increase total available capacity. See ["Defragment Option" on page 64](#) for a description of defragmenting and when it applies.
- 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 133](#). You can change other cache parameters using the rdacutil utility. See the rdacutil entry on [page 136](#).
- Change the reconstruction rate — you can change the reconstruction rate assigned to individual logical units. See ["Changing the LUN Reconstruction Rate" on page 125](#).

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.

IMPORTANT You cannot change the RAID level, capacity, or segment size or defragment the drive group until the background formatting has completed. Also, be sure to close all other storage management applications before starting a LUN modification operation. See ["Logical Unit Formatting" on page 49](#) for more information.

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 28](#) 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.

Drive 1	Drive 2	Drive 3
LUN 1 -- 200 MB	LUN 1 -- 200 MB	LUN 1 -- 200 MB
LUN 2 -- 200 MB	LUN 2 -- 200 MB	LUN 2 -- 200 MB
LUN 3 -- 200 MB	LUN 3 -- 200 MB	LUN 3 -- 200 MB
Empty -- 300 MB	Empty -- 300 MB	Empty -- 300 MB

Figure 28. Defragment Logical Units #1

Now assume LUN 2 (the middle LUN) is deleted. This is shown in [Figure 29](#).

Drive 1	Drive 2	Drive 3
LUN 1 -- 200 MB	LUN 1 -- 200 MB	LUN 1 -- 200 MB
Empty -- 200 MB	Empty -- 200 MB	Empty -- 200 MB
LUN 3 -- 200 MB	LUN 3 -- 200 MB	LUN 3 -- 200 MB
Empty -- 300 MB	Empty -- 300 MB	Empty -- 300 MB

Figure 29. Defragment Logical Units #2

Even though there is actually 1500 MB of empty space in this drive group, for purposes of actual logical unit creation, 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 will show the remaining capacity as 900 MB, ignoring the additional 600 MB.

Assume you create a new LUN 2 of 200 MB. The new logical unit will be written in the 300 MB remaining capacity, as shown in [Figure 30 on page 65](#).

Drive 1	Drive 2	Drive 3
LUN 1 -- 200 MB	LUN 1 -- 200 MB	LUN 1 -- 200 MB
Empty -- 200 MB	Empty -- 200 MB	Empty -- 200 MB
LUN 3 -- 200 MB	LUN 3 -- 200 MB	LUN 3 -- 200 MB
LUN 2 -- 200 MB	LUN 2 -- 200 MB	LUN 2 -- 200 MB
Empty -- 100 MB	Empty -- 100 MB	Empty -- 100 MB

Figure 30. Defragment Logical Units #3

Total remaining capacity is 900 MB, but the configuration utility will show 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 30](#), the resulting drive group will look like [Figure 31](#).

Drive 1	Drive 2	Drive 3
LUN 1 -- 200 MB	LUN 1 -- 200 MB	LUN 1 -- 200 MB
LUN 3 -- 200 MB	LUN 3 -- 200 MB	LUN 3 -- 200 MB
LUN 2 -- 200 MB	LUN 2 -- 200 MB	LUN 2 -- 200 MB
Empty -- 300 MB	Empty -- 300 MB	Empty -- 300 MB

Figure 31. Defragment Logical Units #4

Note that the logical units have been “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 will tell you if there is any space that needs to be reclaimed on the drive group. See "[Defragmenting the Drive Group](#)" for procedural information.

Defragmenting the Drive Group

After you delete a logical unit from the drive group, you can free up space by defragmenting the drive group. See "[Defragment Option](#)" on [page 64](#) for more information.

Use the following procedure to defragment the drive group.

IMPORTANT You cannot defragment the drive group while background formatting is in progress (see "[Logical Unit Formatting](#)" on [page 49](#)). If the selected drive group contains a LUN that is still formatting in the background, a message will appear stating 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, you can select the option from the drop-down menus.)
- 6 A message appears 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 are finished with 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.

Changing the RAID Level

You can change the RAID level of all logical units in the drive group. Keep in mind that you might 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 69](#) and then change its level to RAID 1).

Use the following procedure to change the RAID level of all the logical units in a drive group.

IMPORTANT You cannot change the RAID level while background formatting is in progress (see ["Logical Unit Formatting" on page 49](#)). If the selected drive group contains a LUN that is still formatting in the background, a message will appear stating 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 Controller containing the logical units you want to modify.
4. Select the Drive Group you want to modify in the Module directory tree.
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 32](#) is displayed.
6. Select the RAID level you want for the drive group.
7. After you select the new RAID level, click OK to make the change.
8. As the controller changes the RAID level, the configuration display will show a status of “Modifying” for the drive group and for the individual logical units. The status will give 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.

This shows the current RAID level, number of drives, and logical units in the drive group.

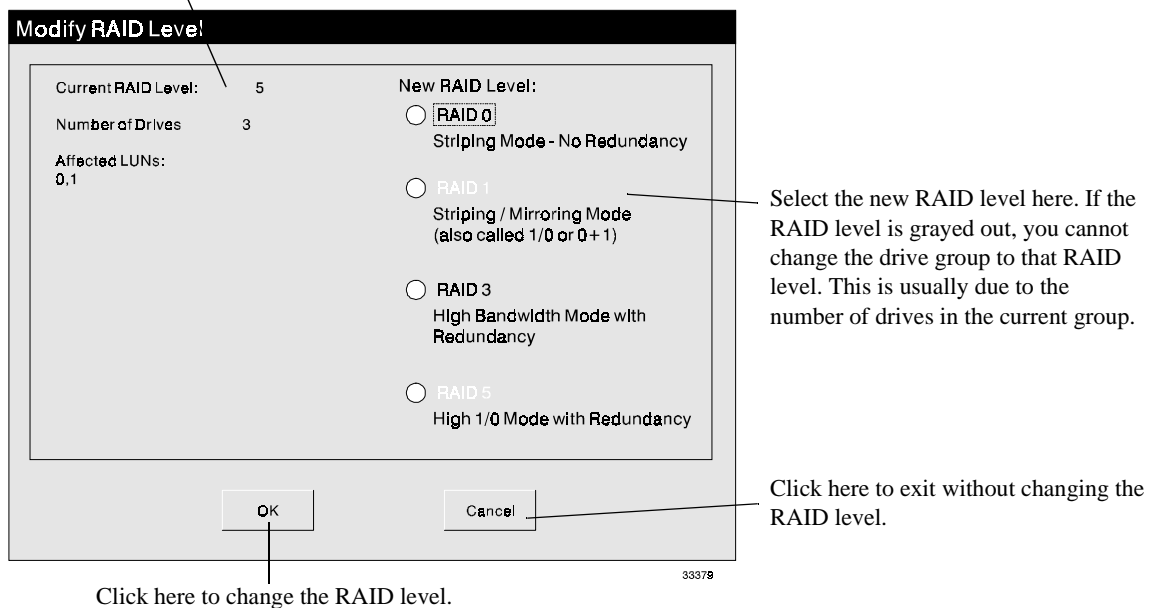


Figure 32. Modify RAID Level Screen

Adding Drives

Use this procedure to increase the capacity of a drive group by adding drives to the group. [Table 11](#) gives the formulae for calculating how much space is added to the drive group.

Table 11: Adding Capacity Formulas

RAID Level	The Drive Group Capacity Is <i>Increased By</i> :
RAID 1	$1/2$ 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 5-drive 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 2×18 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.

Perform the following procedure.

1. Start the Configuration application.
2. Select the RAID Controller containing the logical units you want to modify.
3. Select the Drive Group you want to modify in the Module directory tree.
4. Click on the Modify Group/LUNs button, then select Group Capacity, or select the option from the drop-down menus. A screen like [Figure 33](#) is displayed.
5. Select the drives you want to add to the drive group. You should pick drives of the same size as the current drives in the group.
6. After you have selected the drives, click OK to make the change.

- As the controller adds the drives, the configuration display will show a status of “Modifying” for the drive group and for the individual logical units. The status will give 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.

This shows the current RAID level, number of drives, and logical units in the drive group.



This shows the unassigned drives available to be added to the drive group. You can select more than one drive (for a RAID 1 logical unit, you must select an even number of drives).

Click here to exit without adding drives to the drive group.

Click here to add the selected drives to the drive group.

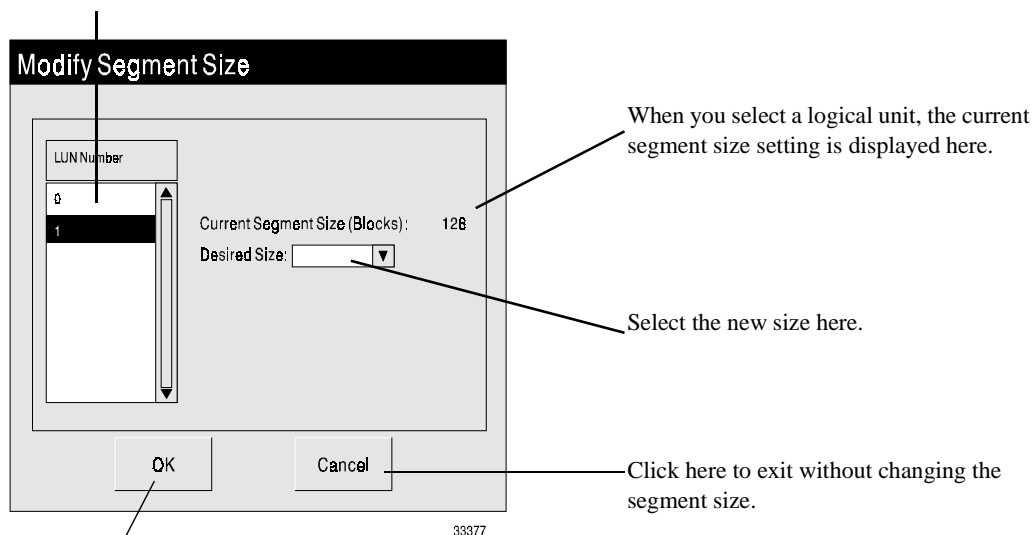
Figure 33. Modify Group Capacity Screen

Modifying Segment Size

Use this procedure to modify the segment size of individual logical units. You do not have to change all the logical units in a drive group. Perform the following procedure.

1. Start the Configuration application.
2. Select the RAID Controller containing the logical units you want to modify.
3. Select the Drive Group or LUN you want to modify in the Module directory tree.
4. 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 34](#) is displayed.
5. One at a time, select the logical units you want to modify, then pick the new segment size.
6. After you have changed the segment size, click OK to make the change.
7. As the controller changes the segment size, the configuration display will show a status of “Modifying” for the drive group and for the individual logical units. The status will give 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.

Select the logical units you want to modify here.



Click here to change the segment size as specified.

Figure 34. Modify Segment Size Screen

Deleting LUNs

When you delete a LUN, any data in the LUN is lost. If you delete all 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 Controller and exit configuration without creating at least 1 LUN. Doing so could prevent you from accessing the RAID Controller.

Before you delete a LUN, check your *Installation and User's Handbook* to see if there are any special restrictions or special requirements on LUN deletion. For example, you may be required to first delete partitions, deassign 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 Network Configuration](#)" on page 73.

To delete a LUN, perform the following steps:

1. Perform whatever steps are necessary on your operating system to prepare for LUN deletion. See your operating system documentation and the *Installation and User's Handbook* for information.
2. Start the Configuration application.
3. Select the RAID Controller containing the LUN you want to delete.
4. Select the drive group containing the LUN you want to delete in the Module directory tree.
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 35 on page 73](#) appears.
6. Select the LUNs you want to delete, or click Select All to delete all the LUNs in them drive group.
7. When the LUNs you want to delete are highlighted, click Delete.
8. After the LUNs are deleted, you need to 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 *Installation and User's Handbook* for your operating system for more information.

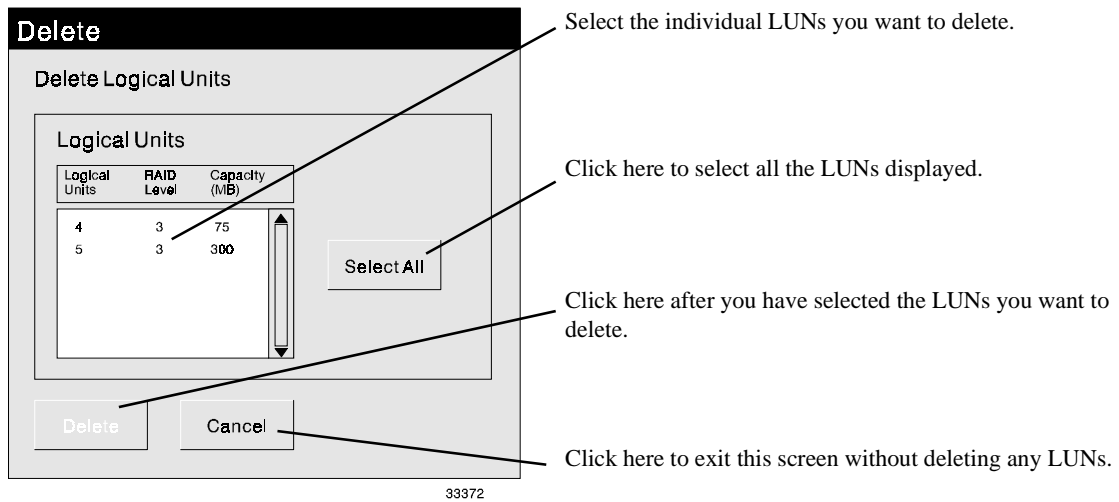


Figure 35. Delete LUNs Screen

Deleting LUNs in the Network Configuration

If you are using the Networked version or a host-RAID Controller configuration with more than one host (see "[Controller Configurations](#)" on page 13), it is possible to open a Configuration Application and begin Delete 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 a Hot Spare

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

IMPORTANT A hot spare cannot replace a drive with a larger capacity. If you have drives of two different capacities in your RAID Controllers, you should select a drive of the *larger* capacity as your hot spare drives.

To create a hot spare drive, do the following.

1. Start the Configuration application.
2. Select the RAID Controller 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 36](#) 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 Controller (this can be further limited by the number of unassigned drives remaining in the RAID Controller).

Set the number of hot spares you want to create here. You may be limited by the number of channels on your RAID Controller, by the number of hot spares already created, and by the number of unassigned drives available.

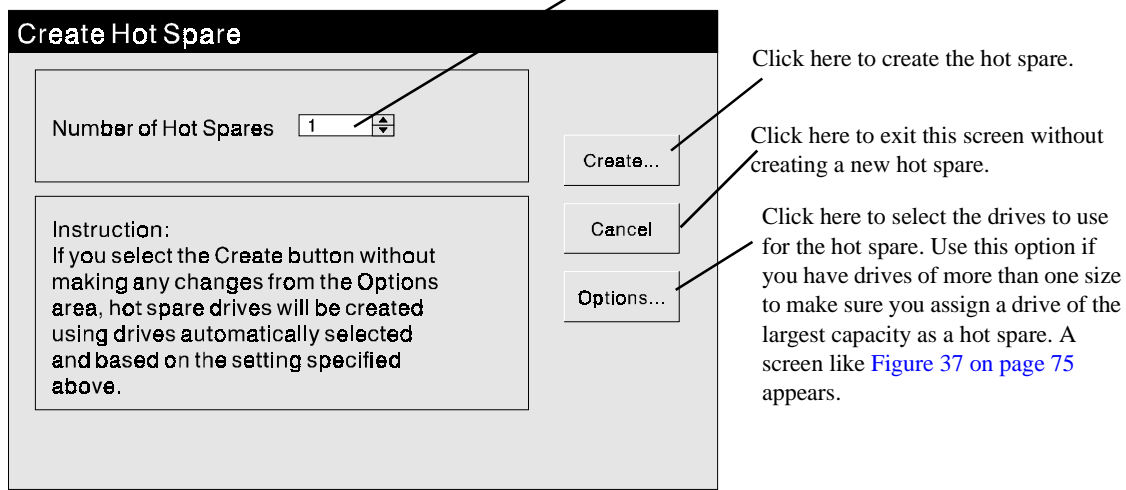


Figure 36. Create Hot Spare Screen 33374

6. If you want, select Options to physically select the drive you want to assign as a hot spare. In most cases, you should just take the default choice and not use this option. 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.

This box shows the unassigned drives available to be assigned as hot spares. Highlight a drive and click the Move button to assign them as hot spares.

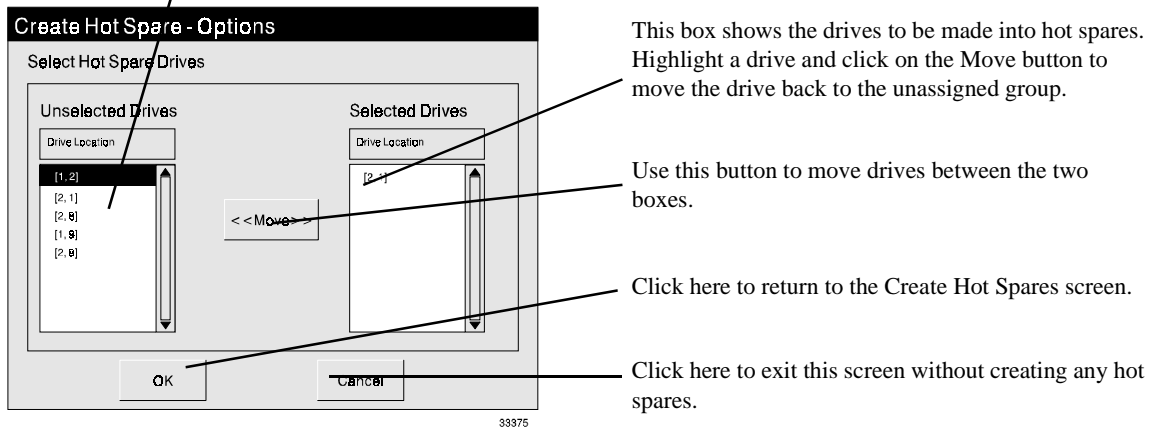


Figure 37. Select Hot Spare Screen

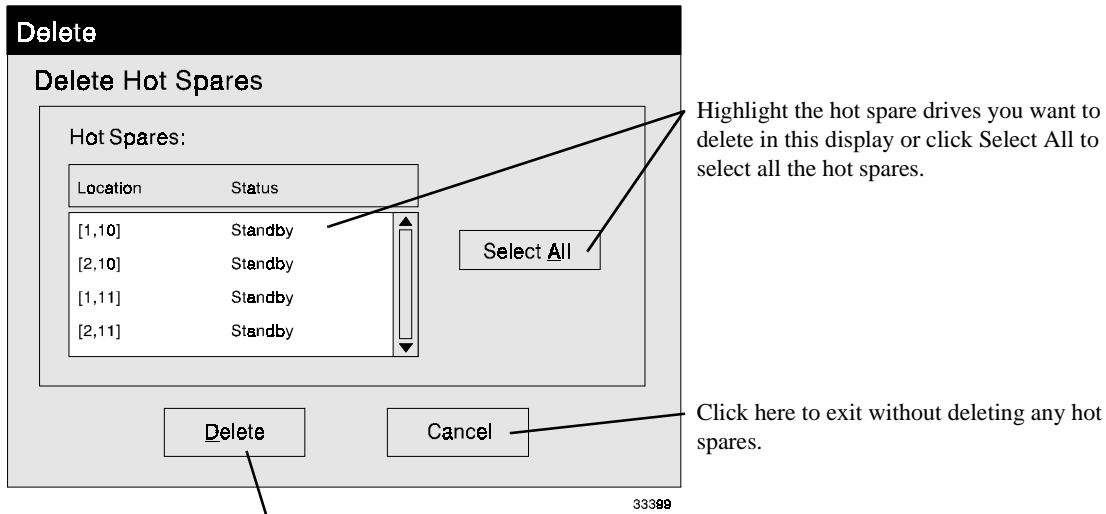
Notes:

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

Deleting a Hot Spare

When you delete a hot spare, the drive returns to the unassigned drive group. To delete a hot spare, perform the following steps.

1. Start the Configuration application.
2. Select the RAID Controller 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 38](#) 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.



Highlight the hot spare drives you want to delete in this display or click Select All to select all the hot spares.

Click here to exit without deleting any hot spares.

After you have selected the hot spares you want to delete, click here to delete them and return the drives to the unassigned drive group.

Figure 38. 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 Controller (*all* drive groups and their logical units) is reset back to a default configuration based on NVSRAM settings specified in the controller.

- Because deleting LUNs causes data loss, back up data on *all* the drive group/LUNs in the RAID Controller. This operation also deletes any file systems mounted on the LUNs.
- You must first stop I/Os to the affected RAID Controller and you should ensure no other users are on the system.
- You cannot reset the configuration on a RAID Controller with an independent controller configuration. For a description of this configuration, see ["Independent Controller Configuration"](#) on page 17.
- You cannot reset the configuration if the controllers are currently active/passive. You must change them to active/active before resetting. See ["Changing the Controller Mode"](#) on page 130 for more information.

- This operation could also fail if the storage management software cannot gain exclusive access to the drive groups/LUNs, such as if file systems are mounted.
- You must restart your system after resetting the configuration for the operating system to recognize the new configuration.

Use the following procedure to reset the configuration on the RAID Controller.

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 will destroy all data on all the logical units in your RAID Controller! [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 have 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, including adding drives and possibly restarting the operating system. Refer to the *Installation and User’s Handbook* specific to your operating system (restrictions and troubleshooting sections) and to the appropriate system documentation for specific details.

Chapter 5. Checking Status

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

Status Overview

The storage management software monitors the status of the RAID Controllers 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 Controllers periodically to find any status changes or other “events.” When a status change is found, the software writes detailed information on the event into the log file.

You can set the polling interval for this background check (see ["Changing Log File Settings" on page 84](#)). The default time is 5 minutes. You can set any time from 1 to 59 minutes. However, setting too small a time may slow system performance because of the time required to poll, and setting too large a time may delay notification of serious failures.

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 *Installation and User's Handbook*), but can be changed to any file (see ["Changing Log File Settings" on page 84](#)).

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 84](#)). 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 85](#) for a description of this procedure). Until you open an application and reset the log file, however, messages will 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 Controller 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 12](#):

Table 12: Message Types

Type	Description
Parity	Messages related to parity check/repair (either manual or automatic). The message details give the location of any bad blocks found, the affected logical unit, and the number of blocks repaired.
Hardware	Messages related to component failures. The message details 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 message 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 message details will try and give a further description of the event. In many cases, they will require no action.

Status Log and Health Check

Note that in some cases, you may discover a problem when running Recovery Guru/Health Check that is not 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 will find 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 *Installation and User's Handbook* for your operating system for details.

Viewing the Log

To view the status log, perform the following steps.

1. Open the Status Application.

2. Select the RAID Controller you want to view status information for (or select All RAID Controllers).
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 40 on page 82](#) 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 ([Figure 39](#)).
6. Continue with [Step 7 on page 83](#).

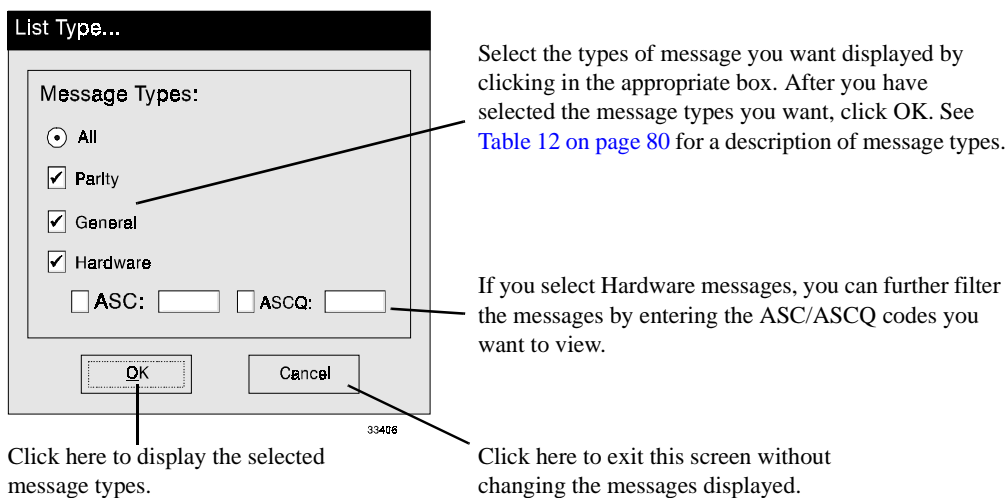


Figure 39. Selecting Message Type

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

The screenshot shows the 'Status-NTPREDATOR' application window. The 'RAID Module' is set to 'SunKing'. The 'Message Log - Summary Information' table contains the following data:

Date	Time	RAID Module	Type	Code	Controller
04/24/1998	18:27:58	SunKing	Hardware	A000	Drive2
04/24/1998	18:27:58	SunKing	Hardware	A000	Drive1
04/24/1998	18:27:58	SunKing	Hardware	A000	Drive0
04/24/1998	18:27:58	SunKing	Hardware	3FC7	Drive0
04/24/1998	14:30:58	SunKing	General		
04/24/1998	14:30:07	SunKing	General		
04/24/1998	14:24:58	SunKing	Hardware	0404	Drive1
04/24/1998	14:02:57	SunKing	General		
04/24/1998	14:02:56	SunKing	General		

Callouts from the image:

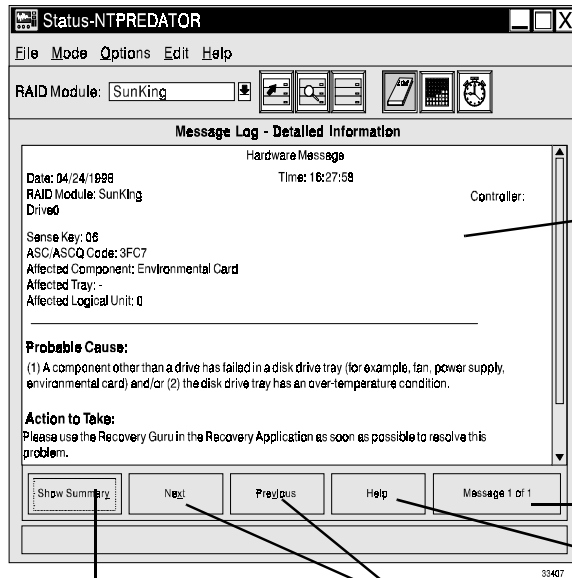
- Click here to display detailed information on the selected messages (see Figure 41 on page 83).
- Click here to select all the messages, then click Show Details.
- Click here to specify the types of messages you want to display (see Figure 39 on page 81).
- Click here for Help.
- This display gives summary information on the messages in the currently selected message log. Table 13 explains the entries shown. Highlight the messages you want to view, then click Show Details.
- This bar shows the selected message log, the number of messages in the log, and the number of messages selected.

Figure 40. Message Log Screen

Table 13: Message Log Explanation

Column Heading	Explanation
Date & Time	Indicates when the detected event was logged. <ul style="list-style-type: none"> 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 Controller event; either Hardware, Parity, or General (see "Message Types" on page 80).
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.

- Highlight a displayed message or messages and click Show Details to get more information on the message. A screen like [Figure 41](#) is displayed.



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 40](#) on page 82).

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

Click here for help.

Figure 41. Message Details Screen

- View the details and take the appropriate action.

- You can also perform the following actions:

- Save the message log to another file — select File Save → Log As from the drop-down menus. See "[Message Log File](#)" on page 79.
- 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 Details screen (see [Figure 41](#)).
- 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 Summary screen (see [Figure 40](#) on page 82).

- When you are finished viewing the messages, close the Status application or select another option.

Changing Log File Settings

You can change the following log settings (see "Message Log File" on page 79):

- 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 85). Note that setting a size does not limit the size of the file. The file will to grow 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 Controllers. Setting too small a value may affect system performance. Setting too large a value may delay notification of serious errors.

Use the following procedure to change log file settings.

1. Start the Status Application.
2. Select *any* RAID Controller (this change affects all RAID Controllers).
3. Select Options → Log Settings from the drop-down menus. A screen like Figure is displayed.
4. Enter the values you want, then click Save (or view the current values and click cancel to exit without changing values).

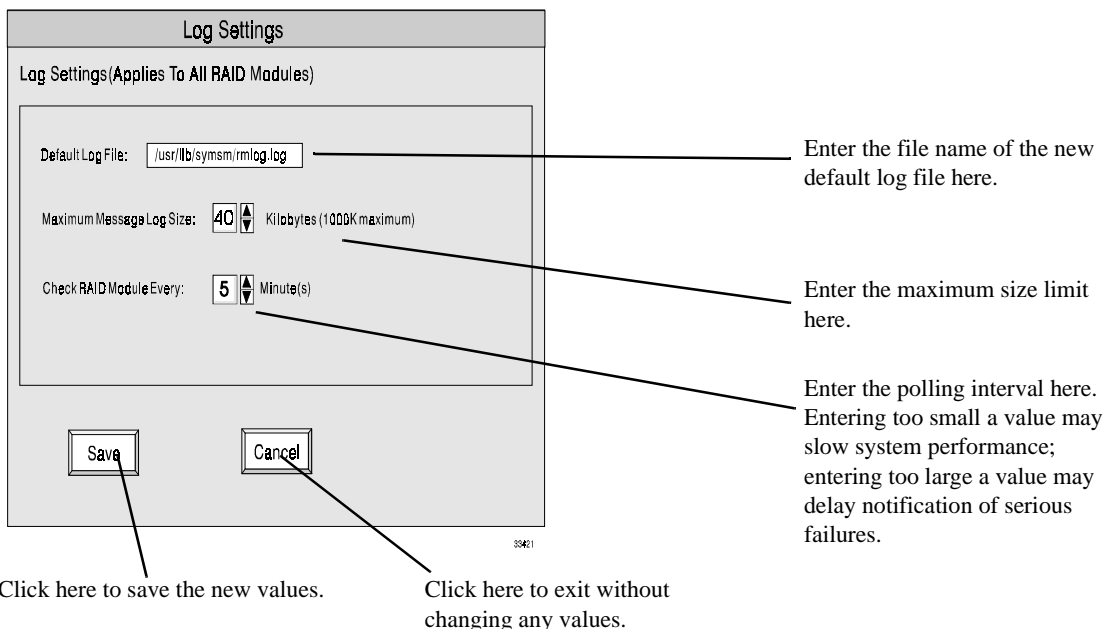
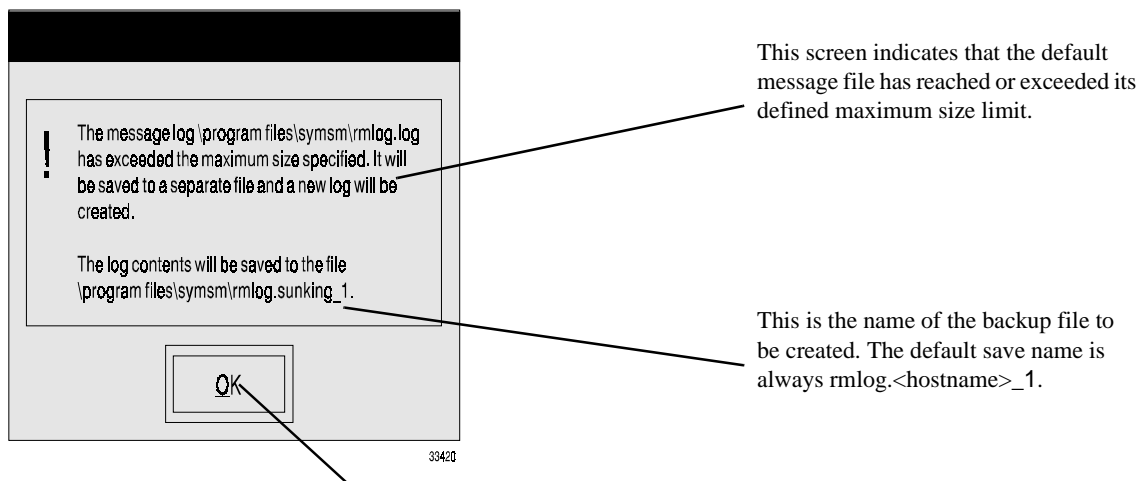


Figure 42. Log Settings Screen

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 84), 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 43](#) is displayed.
2. Click OK.
 - If the file `rmlog.<hostname>_1` 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 44 on page 86](#) 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.



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 44 on page 86](#) is displayed when you click OK.

Figure 43. Message Limit Warning Message

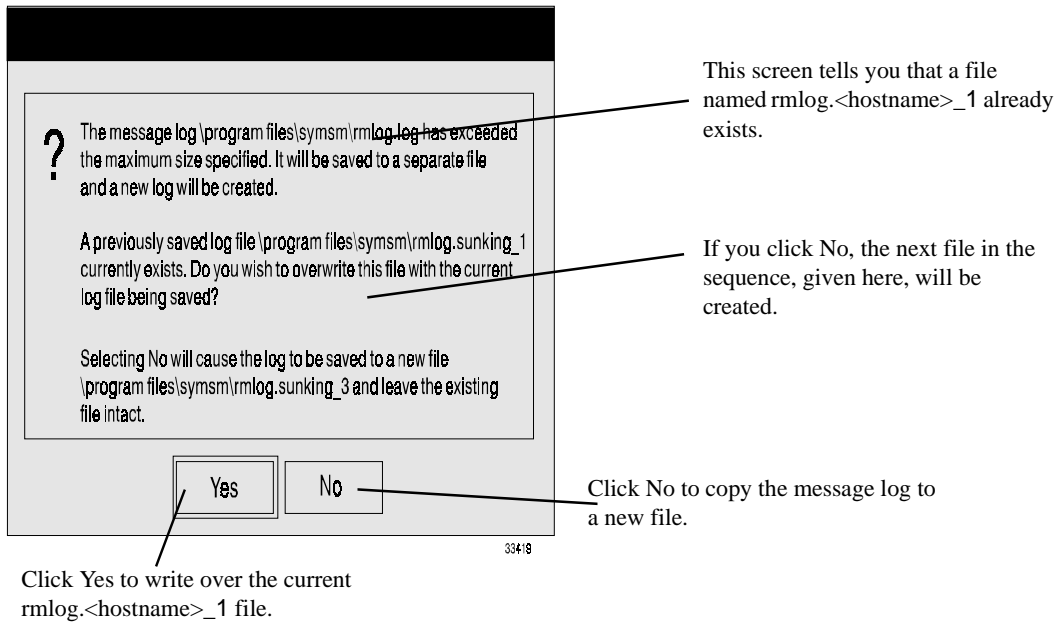


Figure 44. Message Log Save Message

Saved Log File Location

The saved log file is put in a default directory, named in the message. You can change the directory by changing the `System_LogFileDefaultBackupPath` parameter in the `rmparams` file. For more information about the `rmparams` file, see page 137 in this manual or see your *SYmplicity Storage Manager Installation and User's Handbook*.

Using the Performance Monitor

IMPORTANT Before you start the Performance Monitor, make sure you have I/O going to the RAID Controllers you want to monitor.

To run the Performance Monitor, perform the following steps:

1. Start the Status application.

2. Select the RAID Controller you want to monitor.
(Do not select All RAID Controllers.)
3. Click on the Performance Monitor button or select Options → Performance Monitor from the drop-down menus. A screen like [Figure 45](#) is displayed.

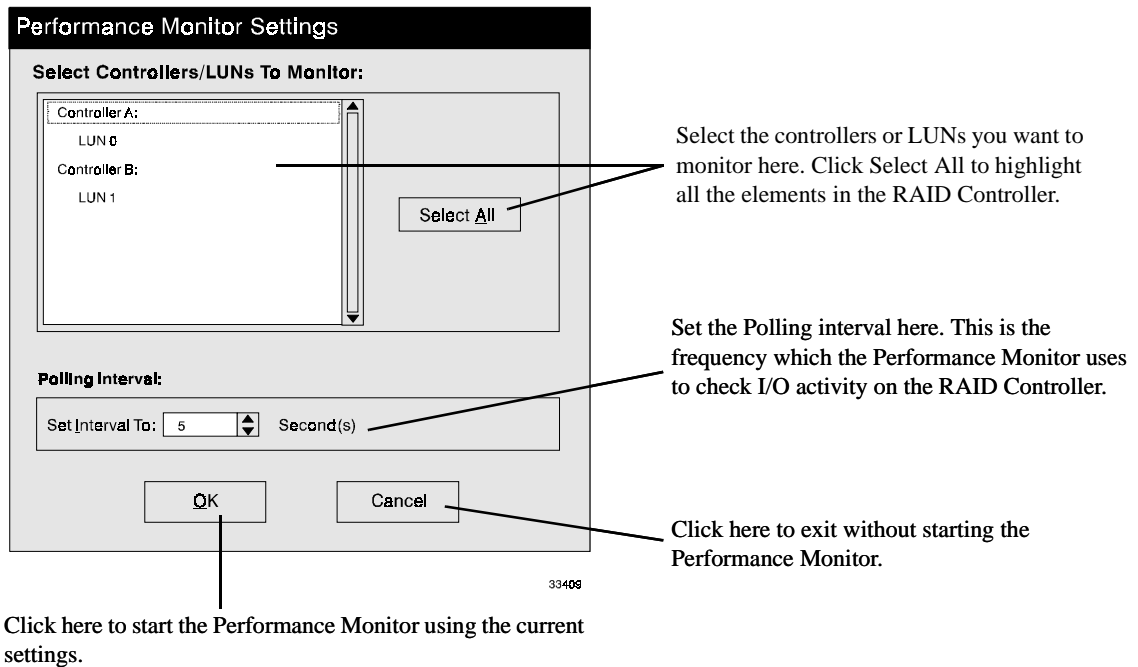


Figure 45. Performance Monitor Settings Screen

4. Select the Controllers/LUNs you want to monitor and click OK. A screen like [Figure 46](#) is displayed.
5. Click Start to start the Performance Monitor. A screen like [Figure 47 on page 88](#) is displayed.

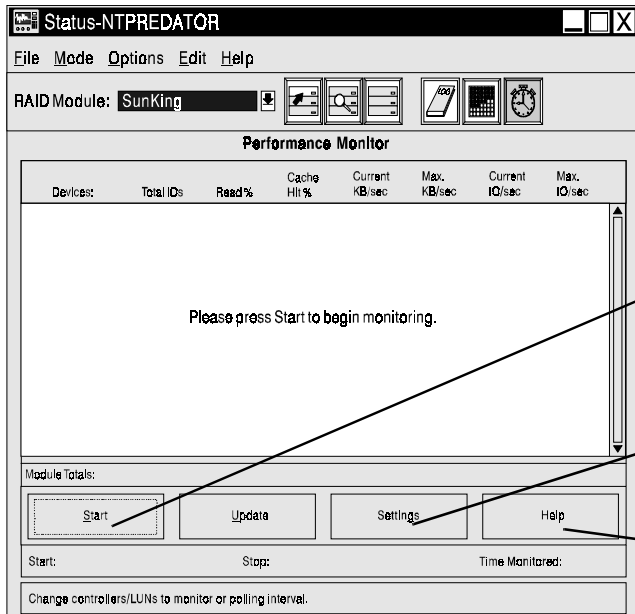
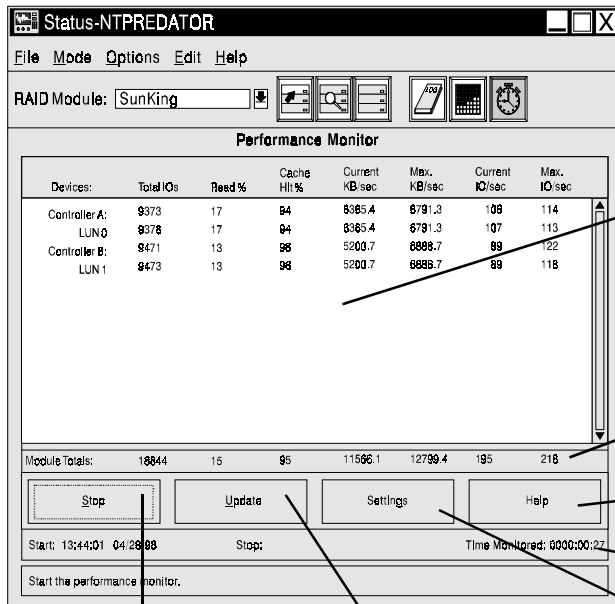


Figure 46. Performance Monitor Start-Up Screen

Click here to start the Performance Monitor.

Click here to change the Performance Monitor settings (see [Figure 45 on page 87](#)).

Click here for help.



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 45 on page 87](#)). See [Table 14](#) 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 stop the Performance Monitor.

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

Click here to change the polling setting (see [Figure 45 on page 87](#)).

Figure 47. Performance Monitor Data Screen

Table 14: Performance Monitor Screen Contents

Column	Explanation
Devices	Controller name listed as Controller A or B; or LUNs listed as LUN 0, LUN 1, etc. The LUNs are listed in numerical order following the controller that owns them.
Total I/Os	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 will be displayed as 100K until the number reaches 9999K, at which time it will be displayed as “M” units. For amounts less than 100M, the value will be 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.

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.

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 Controller. If a particular array controller has considerably more I/O activity, consider moving the LUN 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 will increase 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 will result in lower throughput rates (but will provide 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 rates with relatively small I/O sizes, 4096 bytes, is approximately 17MB/second. So, if 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 will result in higher I/O rates. Higher write I/O rates will be 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 component, such as a drive, controller, fan, or power supply. Because most RAID Controllers have redundancy built into them, a single failure will not stop the RAID Controller 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 manuals accompanying your RAID Controller for the redundant features on that module. In most cases, the module has redundant controllers, fans, and power supplies.

IMPORTANT You may need to refer to the hardware manuals accompanying your RAID Controllers for information on how to physically replace failed components. The information in this *User Handbook* provides information on what you must do with the software to enable component replacement.

If the failed component has redundancy, in most cases you can replace it without shutting down the RAID Controller or stopping I/O to it. For example, if your RAID Controller has redundant fans or power supplies, you can replace one of these components if they fail without shutting down the RAID Controller.

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. This notification could be via:
 - Software error messages
 - An error message in the Message Log
 - Fault lights or audible alarms on your RAID Controller
 - SNMP notification
 - Console messages from your operating system
 - Results of a Recovery Guru/Health Check
 - A non-optimal status appearing for any RAID Controller 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 100).
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 49](#) on page 102).
4. You take the steps indicated, which often involve replacing the failed component. In some cases, you may need to refer to your hardware documentation for information on how to physically replace failed components. You may also need to perform the functions given in "[Manual Recovery Options](#)" on page 103, but do not do so unless instructed by Recovery Guru.
5. You perform whatever functions are necessary on your operating system to recover from the failure (see your Operating System documentation and the *Installation and User's Handbook*).

Recovering from Drive Failures

CAUTION	This section gives the general principles involved in replacing failed drives. However, you should <i>always</i> 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 is 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 to 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 procedures that you take if a standard drive fails. See the documentation accompanying your operating system 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 15](#) summarizes the recovery procedures required. Note that in the case of single drive failures on RAID 1, 3, or 5 logical units, you will not need to shut down the RAID Controller to recover from the failure. [Table 15](#) gives general information only; *always* run Recovery Guru/Health Check before attempting to replace drives.

Table 15: 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 will continue 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 will then reconstruct the data on the replaced drive, and the logical unit will return to optimal status.
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. <ul style="list-style-type: none"> • 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 will then reconstruct the data on the replaced drives, and the logical unit will return to optimal status.
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.

NOTE 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. They do this by automatically replacing the failed drive. For example, if your RAID Controller has 1 hot spare drive, *three* drives must fail before data is lost in a RAID 5 logical unit. See "[Hot Spare](#)" on page 9 and "[Protecting Yourself from Component Failures](#)" on page 19 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, you should try and 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 and revive the logical unit using manual recovery (see ["Manual Recovery Options" on page 103](#)). 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 of Drives" on page 104](#) for a description of the Revive Drive procedure.

Revive a drive only if:

- You are instructed to by Recovery Guru.
- You accidentally failed an *optimal* drive in an *optimal* RAID 0 LUN (with Manual Recovery → Fail Drive) and have not performed any writes to the LUN in the meantime. In this case, the LUN 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 LUN (with Manual Recovery → Fail Drive) and have not performed any writes to the LUN in the meantime. In this case, the LUN 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 LUN and you were writing to it at the time (the LUN status will change to Dead, and the operating system will send 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 LUN and you were writing to it at the time (the LUN status will change to Dead, and the operating system will send 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* LUN and the LUN status remains optimal (that is, you were not writing to the LUN at the time you removed it). The drive will automatically be returned to an optimal status after you reinsert it.
- You accidentally remove an *optimal* drive from an *optimal* RAID 1, 3, or 5 LUN and the LUN status changes to degraded (that is, you were writing to the LUN at the time you removed it). The drive will automatically be reconstructed after you reinsert it.

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 12](#) for more information on reconstruction.

You can monitor the progress of reconstruction, and change the reconstruction rate, by using the Status Application. See ["Monitoring/Changing Reconstruction" on page 109](#) for more information.

Using Configured Drives as Spares

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 use 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 must then revive the drive using the procedure in ["Reviving Drives" on page 95](#) *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) PFA (Predictive Failure Analysis) reporting. 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 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 a 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 recovery procedure.

In general:

- If a hot spare is currently being reconstructed on the logical unit affected by the problem drive, wait for reconstruction to complete 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 back up.
- If the drive is in a degraded RAID 1, use the List/Locate Drives option to determine if the problem drive is the mirrored pair of the drive that already failed. See "[Locating Drives](#)" on page 37).
 - 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 back up.
- If the drive is in a degraded RAID 3 or 5 LUN, 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 back up.

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

- The RAID Controller Unit has redundant controllers (see "[Redundant Controllers](#)" on page 12).

- The RAID Controller Unit is configured in the single host configuration using two host busses (see "[Controller Configurations](#)" on page 13).
- The RAID Controller has RDAC protection (see "[RDAC Driver](#)" on page 11).

Always consult your hardware documentation 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 16](#) summarizes the steps you need to take to recover from a failed controller, depending on whether or not your RAID Controller has fail-over protection.

Table 16: Recovering from Failed Controllers (Summary)

Fail-Over Protection?	Recovery Procedure
Yes	<p>If a controller fails, the other controller in the pair will take over for it. LUNs assigned to the failed controller will be automatically assigned to the remaining good controller. Data will continue 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.</p> <p>Recovery: You will probably be able to replace the failed controller without shutting down the RAID Controller Unit. Follow the instructions on the Recovery Guru procedure screen.</p>
No	<p>If the RAID Controller Unit 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 destroyed, but is inaccessible.</p> <p>Recovery: You will need to shut down the RAID Controller Unit to replace the controller. Follow the instructions on the procedure screen.</p>

Controller Replacement and Firmware (Autosynch)

When you replace a failed controller, the firmware on the replaced controller is automatically synchronized with the firmware on the remaining 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.

NOTE This feature applies only to controllers with firmware levels of 3.0 or greater, and only if certain NVSRAM bits are set. See your *Installation and User's Handbook* for information on these bit settings on your system.

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 new controller online, the firmware on the new 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 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 automatic synchronization is occurring.

Note that if either controller has a firmware level earlier 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 remaining good controller.

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

Controller Replacement and Cache Parameters

After you replace a failed controller, it might take as long 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 133](#)).

Recovering from Fan/Power Supply Failures

CAUTION This section gives the general principles involved in replacing failed RAID Controller 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.

The RAID Controller Unit contains redundant fans and power supplies. A single failure, therefore, will not shut down the RAID Controller Unit. You can then replace the failed component while the RAID Controller Unit is running. The replaced component will automatically be returned to service in the RAID Controller Unit.

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

Recovering from Battery Failures

RAID controllers have a battery unit to preserve cache memory in the event of a power failure. This battery has a limited life-span. 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 on 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, enter the following from the command line:

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

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

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

Running Recovery Guru/Health Check

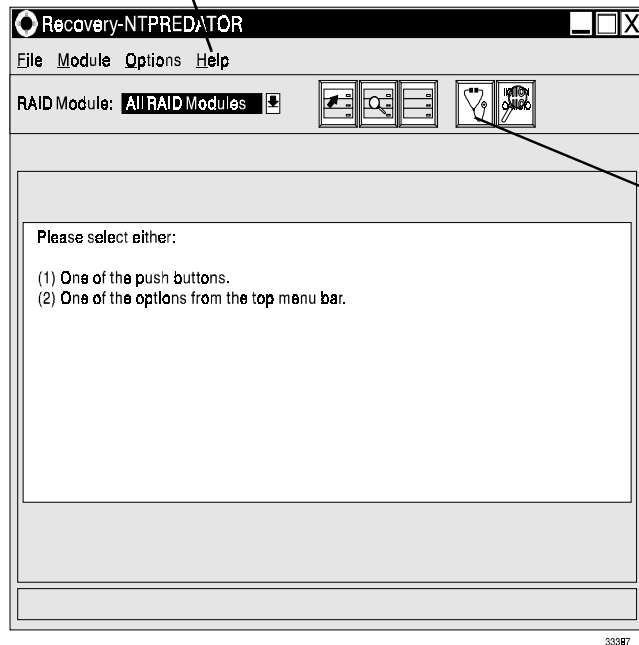
Run Recovery Guru/Health Check when you have some indication that there is a failure in some array component or if you want to check the status of your modules. This Health Check locates the failure and gives detailed information on the steps you need to take to correct it.

Use the following procedure to run Recovery Guru/Health Check.

1. Start the Recovery application.

2. Select the RAID Controllers you want to check, or select All RAID Controllers to check all the modules on your system. A screen like [Figure 48](#) appears.

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

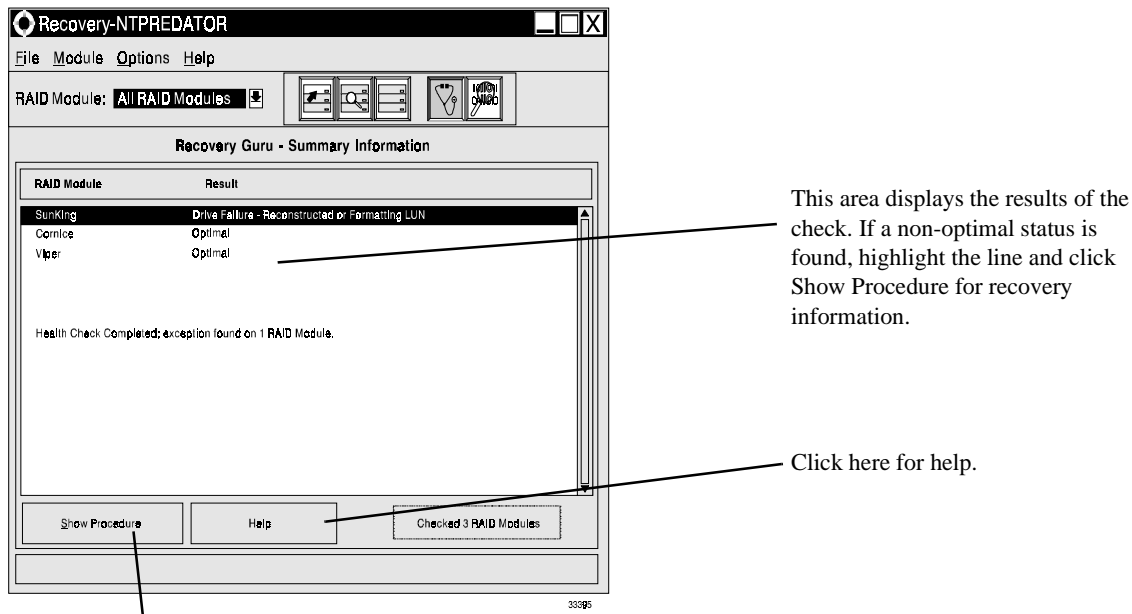


Click here to start the Health Check on the selected modules. You can also start the Health Check by selecting Options → Recovery Guru.

Figure 48. 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 49](#) is displayed.



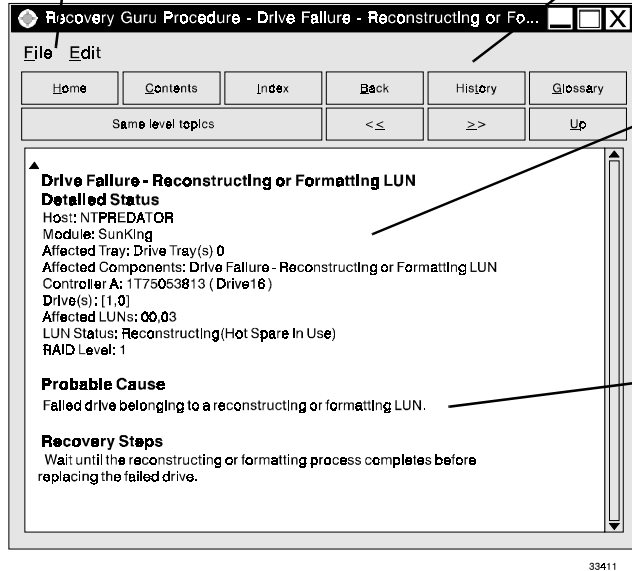
Click here to display information on the highlighted failure. This information will include recovery information. See [Figure 50 on page 103](#) for a sample screen.

Figure 49. Recovery Screen/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 50](#)).
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 103.

Select File here to print the information displayed in the procedure screen, or to set up the printer before printing. Select Edit to copy text to the clipboard.

Although this looks like a standard Help screen, these buttons are grayed out.



The first part of this area gives detailed status information on the failure. In this case, the message indicates that drive [1,0] has failed. A hot spare drive has taken over for the failed drive and is currently reconstructing. LUNs 0 and 3 are affected by this failure.

The rest of this area gives information on the cause of the failure and on how to recover from it. Read the instructions carefully. If necessary, print them out by selecting File → Print Topic from the drop-down menus.

In this case, there is no action to take until the reconstruction is finished. Then you should run Health Check again for information on how to replace the failed drive.

Figure 50. 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 a Customer Services representative. Doing so could result in the loss or corruption of data. Do not attempt to manually recover without understanding the circumstances of the failure.

The manually recovery options provide a means to perform operations on drives, logical units, and controllers. You may need to perform these procedures if instructed by Recovery Guru.

Use the Manual Recovery options if:

- 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.
- 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 95](#)).

- 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 Controller.
- 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 of Drives](#)" on page 104.
- Manual recovery of logical units — See "[Manual Recovery of Logical Units](#)" on page 106.
- Manual recovery of controllers — See "[Manual Recovery of Controllers](#)" on page 108.

Manual Recovery of 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 a Customer Services representative. Doing so could result in the loss or corruption of data. Do not attempt to manually recover without understanding the circumstances of the failure.
----------------	--

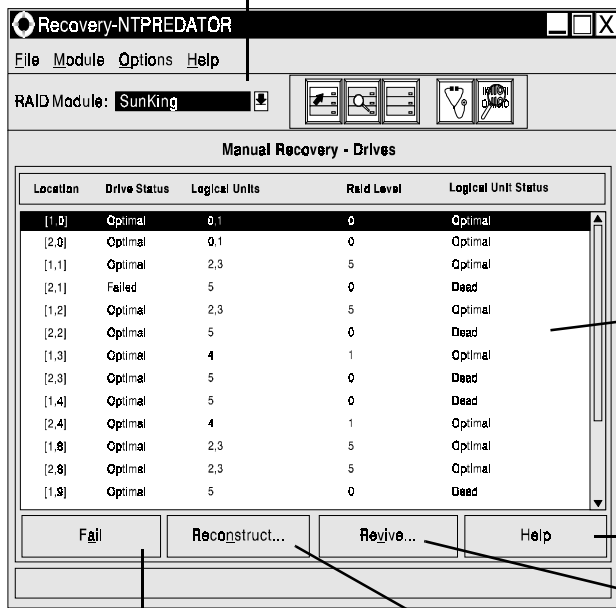
Using manual recovery on drives, you can:

- 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 might need to do this if you accidentally fail or remove a drive. See "[Reviving Drives](#)" on page 95 for a description of when to use this option.

To perform manual recovery of drives, do the following:

1. Start the Recovery application.
2. Select the RAID Controller containing the drives you want to recover.
3. Select Options → Manual Recovery → Drives from the drop-down menus. A screen like [Figure 51 on page 105](#) is displayed.
4. Select the individual drive you want to fail, reconstruct, or revive, then click on the appropriate button to perform the operation.

Select the RAID Controller containing the drives you want to recover.



Information on the drives in the module is displayed here. Highlight the drive you want to fail, reconstruct, or revive.

Click here for help.

Click here to Fail the currently highlighted drive. This button will be grayed out if the selected drive is already failed. Note that failing a drive could result in the loss of data, especially in a RAID 0 logical unit.

Click here to reconstruct data on the selected drive. You can only perform this option on a drive with a status of Replaced or Failed in a RAID 1, 3, or 5 logical unit. The logical unit status must be Degraded or Optimal.

Click here to Revive the drive. This option works only if the drive status is Failed. See ["Reviving Drives" on page 95](#) for a description of when to use this option.

Figure 51. Manual Recovery of Drives Screen

Notes:

- 1 Failing a drive in a RAID 0 logical unit will result in the loss of data on the logical unit.
- 2 Failing a drive in a Degraded RAID 3 or 5 logical unit will result in the loss of data.
- 3 Failing a drive in a degraded RAID 1 logical unit will result in the loss of data if the drive you fail is the mirror drive of a failed drive.
- 4 Reconstruction should only be performed on a drive after you have replaced it.
- 5 Reviving a drive should be done only under the circumstances described in ["Reviving Drives" on page 95](#).

Manual Recovery of 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 a Customer Services representative. Doing so could result in the loss or corruption of data. Do not attempt to manually recover without understanding the circumstances of the failure.

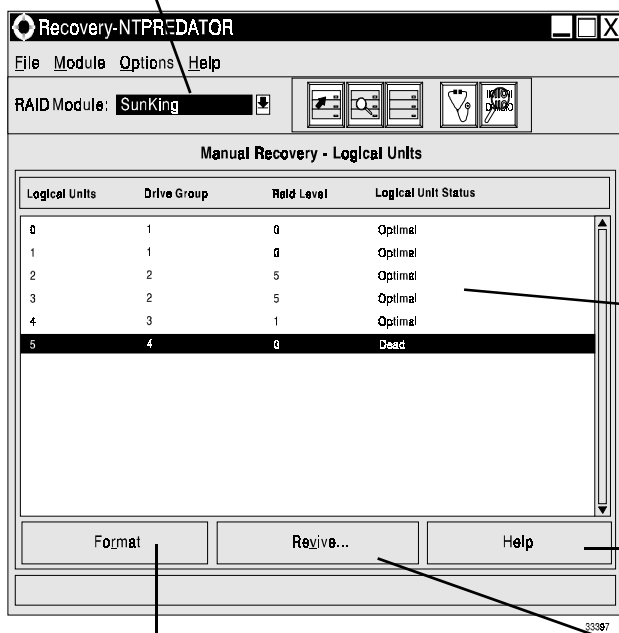
Using manual recovery on logical units, you can:

- 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 perform manual recovery of logical units, do the following:

1. Start the Recovery application.
2. Select the 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 52 on page 107](#) is displayed.
4. Select the logical units you want to format or revive, then click on the appropriate button to perform the operation.

Select the RAID Controller containing the logical units you want to recover.



This display shows the logical units on the selected RAID Controller and the current status of those units. Highlight the logical units you want to format or revive and click on the appropriate button

Click here for help

Click here to format the selected logical units. You must first replace any drives that caused the failure. Formatting deletes all data on the logical unit.

Click here to revive the selected logical units. You should only revive a logical unit if the logical unit failure was caused by a channel failure. Use the Recovery Guru to determine the cause of the failure. Reviving a logical unit under other circumstances will most often result in corrupted data or more failures.

Figure 52. Manual Recovery of Logical Units Screen

Notes:

- 1 Do not format a logical unit until you have replaced any failed drives that caused the failure.
- 2 Formatting deletes all data on the logical unit (it is probably already lost anyway, assuming the logical unit status is Dead).
- 3 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 94 for more information.

Manual Recovery of 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 a Customer Services representative. Doing so could result in the loss or corruption of data. Do not attempt to manually recover without understanding the circumstances of the failure.

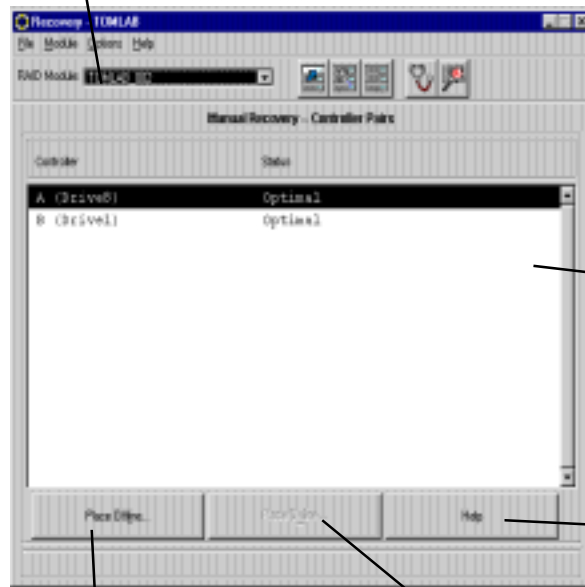
Using manual recovery on controllers, you can:

- 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 was not put back online automatically.

To perform manual recovery of controllers, do the following:

1. Start the Recovery application.
2. Select the 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 53 on page 109](#) is displayed.
4. Select the controller you want to place online or offline and then click on the appropriate button to perform the operation.

Select the RAID Controller containing the controllers you want to recover.



This display shows the controllers in the selected RAID Controller Unit and their current status. Highlight the controller you want to put online or offline and click on the appropriate button below.

Click here for help.

Click here to put the selected controller Offline. This option is only available if the selected controller is currently online.

Click here to put the selected controller Online. This option is only available if the selected controller is currently offline. After you put the controller online, you may need to change controller mode (see ["Changing the Controller Mode" on page 130](#) for more information).

Figure 53. Manual Recovery of Controllers Screen

Notes:

- 1 You can only place a controller offline if it is one of a redundant pair of controllers. The other controller takes over for the offline controller (see ["Redundant Controllers" on page 12](#)).
- 2 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 the Controller Mode" on page 130](#)).

Monitoring/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 LUN Reconstruction Rate" on page 125](#).

To view reconstruction progress and change the rate of reconstruction, perform the following steps:

1. Start the Status Application.
2. Select the RAID Controller containing the logical units currently being reconstructed. Do not select All RAID Controllers.
3. Click on the LUN Reconstruction button or select Options → LUN Reconstruction from the drop-down menus. A screen like [Figure 54 on page 111](#) is displayed.
4. The display shows the progress of reconstruction on each logical unit. If you want to change the reconstruction rate, move the slider bar for that logical unit.

IMPORTANT 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 speed 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 will be completed on the hot spare before it begins on the replaced drive. This may increase the time required to complete the recovery procedure.

This must be a module containing logical units currently being reconstructed.

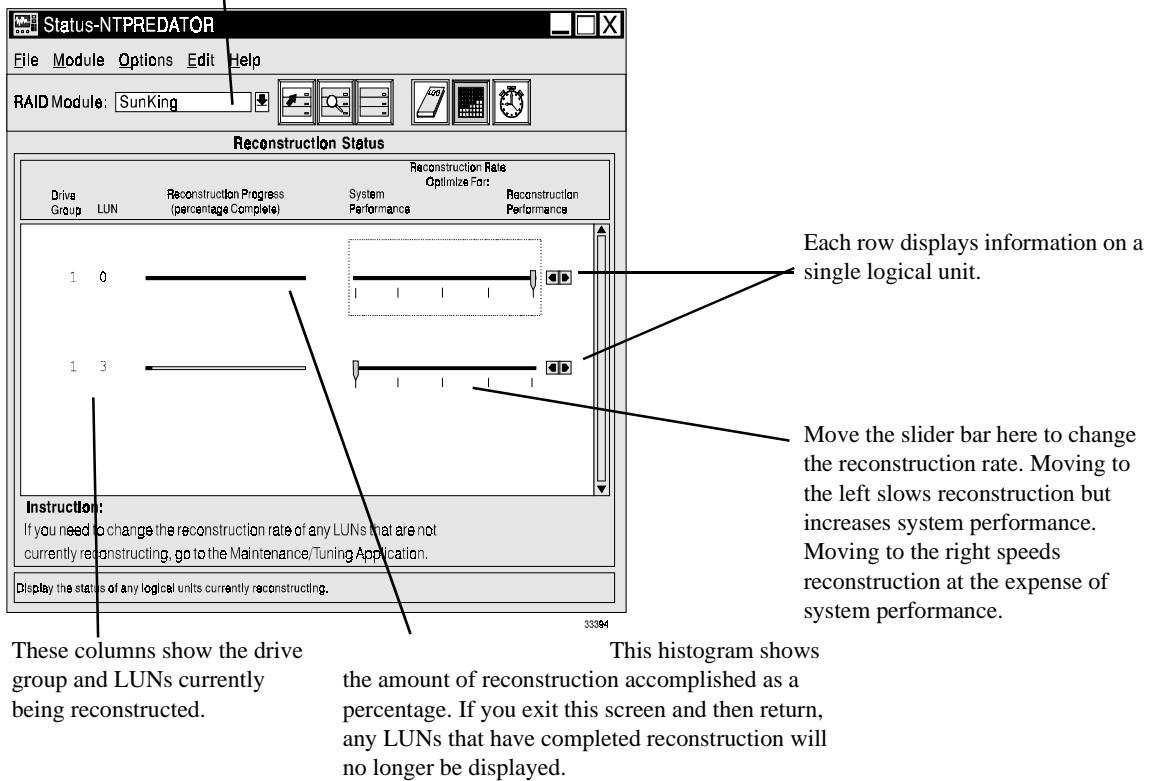


Figure 54. Viewing/Changing Reconstruction Screen

Notes:

- 1 If you exit this screen and then return to it, any logical units that have completed reconstruction will no longer be displayed.
- 2 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.
- 3 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.
 - a Slow — 1024 blocks/0.8 second delay
 - b Medium slow — 1024 blocks/0.4 second delay
 - c Medium — 1024 blocks/0.2 second delay
 - d Medium fast — 1024 blocks/0.1 second delay
 - e Fast — reconstruct with no delays.

Chapter 7. Maintenance and Tuning

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

Maintenance and Tuning Overview

The Maintenance and Tuning application performs functions needed to maintain and tune your RAID Controllers. With the Maintenance and Tuning application, you can:

- Download new controller firmware and NVSRAM files (see ["Upgrading the Controller Firmware" on page 113](#)).
- Manually run Parity Check (see ["Running a Parity Check" on page 122](#)).
- Set the time for Automatic Parity Check (see ["Running a Parity Check" on page 122](#)).
- Change the Reconstruction Rate for logical units (see ["Changing the LUN Reconstruction Rate" on page 125](#)).
- Balance drive groups/logical units across active/active redundant controller pairs (see ["Balancing LUNs" on page 127](#)).
- Swap active/passive controllers or change them to active/active (see ["Changing the Controller Mode" on page 130](#)).
- Change the write cache, write cache mirroring, and cache without batteries parameters for logical units (see ["Changing Cache Parameters" on page 133](#)).

Upgrading the Controller Firmware

Controller firmware resides on the array controller in the RAID Controller. It contains the commands necessary to operate the RAID Controllers. Normally, you do not need to load new firmware when you install a RAID Controller. You may need to perform this function as part of an upgrade. You will receive new firmware files from your Customer Services representative.

IMPORTANT Controller firmware is different from the drive firmware. Use this option only to upgrade controller firmware.

Controller Firmware Files

Downloading controller firmware involves up to five files:

- The NVSRAM file — The NVSRAM file specifies default settings for the controller. You must download this file, if present, *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 Fibre Channel firmware — this file contains firmware to operate the RAID Controller in a fibre channel environment. The fibre channel information is combined in the firmware (appware) file and is not a separate file.
- 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, Firmware, and Fibre Channel version levels through the Controller Profile screen; see "[Controller Profile Screen](#)" on page 40.

General Procedure

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

1. 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).
2. 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 *Installation and User's Handbook* for information on which directory to use).

IMPORTANT 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 filenames, the software is unable to check the firmware files for compatibility or to provide you with a list of compatible files to select for downloading.

3. Download the new NVSRAM file using the Firmware Upgrade option in the Maintenance and Tuning application.
4. Run `nvutil -vf` from the command line to ensure that NVSRAM settings are set-up correctly for the storage management software.

5. Reset the controller or turn the power to the RAID Controller off and then on again, to establish the new NVSRAM settings.
6. Download the new firmware files to your system.

Read through the following sections, then go to the complete download procedure on [page 116](#).

Online and Offline Method

You can download controller firmware using one of two methods, Online and Offline, depending on your controllers and configuration.

IMPORTANT Note that because you must turn the RAID Controller off and then on again, or reset the controller, after upgrading NVSRAM, you must use the Offline method if a new NVSRAM file must be downloaded.

- Select Online to upgrade firmware while the selected RAID Controller 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 Controllers have a status of Optimal.
- Select Offline to upgrade firmware when the selected RAID Controller 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 will not take effect until after you turn the RAID Controller off and then on again).
 - The RAID Controller has an independent controller configuration (in this case, you must download the firmware to each controller from its host).
 - The RAID Controllers have single controllers.
 - The Offline option also requires exclusive access to the logical units in the selected RAID Controllers; that is, no other operations can be running on the RAID Controller.

The software will notify 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 Controller you select:

- When you select All RAID Modules → Firmware Upgrade, or one RAID module and the Online method, you will be downloading NVSRAM files or upgrading 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 that has 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/controller firmware *unless* you are replacing a failed controller and the replacement controller has an earlier firmware version than the original pair was using.

Download Procedure

Perform the following steps to download new controller firmware.

1. Copy the NVSRAM file, firmware files, and the fwcompat.def file to the correct subdirectory in the installation directory of your system (see your *Installation and User's Handbook*).
2. Determine whether you can upgrade the firmware using the Online or Offline method. If you must use the Offline method, make sure that you stop all I/O to the RAID Controllers you are going to upgrade.
3. Start the Maintenance and Tuning application.
4. Select the RAID Module containing the controllers you want to upgrade. Select All RAID Controllers to download firmware to all controllers in all RAID Controllers.
5. Click on the Firmware Upgrade button or select Options → Firmware upgrade from the drop-down menus.
6. Read the Important Notes screen and click OK.
7. A screen appears asking you to select the online or offline procedure. Select either:

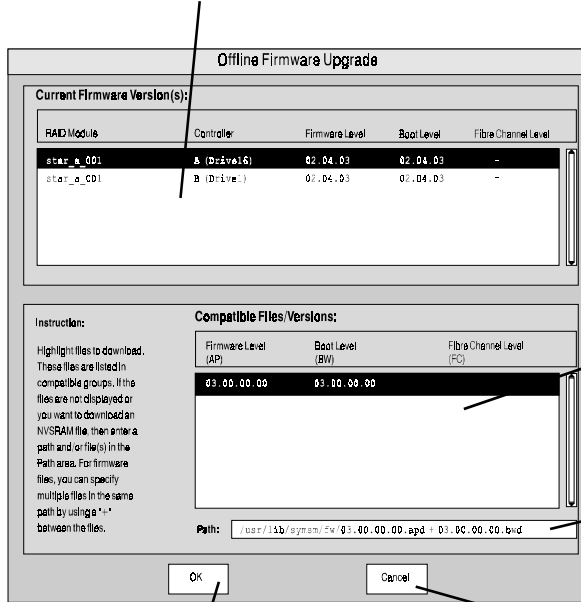
- Online to upgrade firmware while the selected RAID Controller receives I/O.
 - Offline to upgrade firmware when the selected RAID Controller is not receiving I/O.
8. After selecting Online or Offline, the screen displays “Verifying the controller state” while the software checks the selected RAID Controllers for restrictions based on the type of firmware upgrade you selected. If there are no restrictions, a screen like [Figure 55 on page 118](#) is displayed.
 9. Depending on whether you are downloading NVSRAM files or upgrading controller firmware, do one of the following:

IMPORTANT Remember, you need to download any NVSRAM files first, *before* downloading any new firmware files.

- To download an NVSRAM files, type its complete path information in the path box and select OK. Continue with [Step 11 on page 118](#).
 - To select controller firmware, highlight the version level you want to download. Continue with [Step 10](#).
10. Select OK with the correct version level highlighted.

You either receive notification that a problem occurred, or you have a final confirmation that the upgrade process is about to begin.

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 by fwcompat.def. That program checks the default subdirectory for files. 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 55. Firmware Upgrade Screen

IMPORTANT Once you click OK at the “Firmware is about to start” prompt in [Step 11](#), do *not* select any other options or exit the Maintenance/ Tuning Application until the upgrade process is completed. You can, however, monitor the upgrade progress.

11. Select OK and follow 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 17 on page 119](#) 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 final summary report should indicate which modules were not successful and give an appropriate cause. For more information, see ["Confirming the Download" on page 120](#).

Table 17: 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 at the Compatible Files/Versions screen (Figure 55 on page 118).
RAID Controller	Identifies the specific RAID Controller.
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 18 on page 120 if you see any Failed download statuses.

12. After the download is completed, select OK to return to the Maintenance/ Tuning screen.

13. Depending on whether you are downloading NVSRAM files or upgrading controller firmware, do one of the following:

- If you have successfully downloaded NVSRAM files, continue with [Step 14](#).
- If you have successfully upgraded controller firmware, you are finished with this procedure.

14. At the command line, type:

```
nvutil -vf
```

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

15. Turn the power to the RAID Controller Unit off and then on again to establish the new settings.

16. Go back to [Step 9](#) and download any new firmware files.

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 Controllers.

- If you see that the upgrade was `Successful`, you should still verify that all the logical units (LUNs) are *not* assigned to only one controller. See ["Balancing LUNs" on page 127](#) 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 18](#) for possible actions to take to correct a failed upgrade.

Table 18: 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 filesystems 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 are 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, call your Customer Services representative.
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 , "Chapter 6. Recovery."
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.

Table 18: Controller Firmware Upgrade, Failed Statuses (continued)

Reasons for Failed Status	Actions
<p>The files you selected are not compatible with the current firmware version on the selected module's controllers.</p>	<p>Most likely, the current directory does not contain <i>all</i> the necessary firmware files. Copy the firmware files and the <i>fwcompat.def</i> 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.</p>
<p>The software was unable to access the controllers during the upgrade process.</p>	<p>Use Recovery Guru/Health Check in the Recovery Application to determine if the module has a failure. See Chapter , "Chapter 6. Recovery."</p> <ul style="list-style-type: none"> • 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.
<p>You tried to load a pre-2.04 firmware version, which is not supported by this software or the redundant controller configuration.</p>	<p>Do <i>not</i> try to load this firmware version again.</p>
<p>The selected firmware files are not compatible with your controller model.</p>	<p>Use Module Profile → Controller details to check your controller type and model (see Figure 13 on page 39), and obtain the correct firmware version files.</p>
<p>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.</p>	<p>Use Module Profile → Controller details to determine how many controllers the module has (see Figure 13 on page 39).</p> <ul style="list-style-type: none"> • 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. Recovery").
<p>An unknown failure occurred.</p>	<p>Use the Status Application to select Message Log for component information (see "Viewing the Log" on page 80).</p>

Running a Parity Check

The RAID Controllers 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 can not be checked. A non-optimal RAID 1, 3, or 5 logical unit (such as a degraded units) also can not be checked.

If a parity error is found, the Parity Check/Repair application fixes the parity. 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 Parity Check/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, you should set your RAID Controllers 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.

Setting Automatic Parity Time

To ensure data integrity on your RAID Controllers, 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 Controllers, 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, do the following:

1. Start the Maintenance and Tuning application.
2. Select any RAID Controller or All RAID Controllers. The automatic time always applies to all RAID 1, 3, and 5 logical units on all the RAID Controllers 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 56](#) 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.

IMPORTANT Automatic parity check runs only once a day. If you modify its run-time setting after it has already run for the day, it will not run again until the next day. For example, if the automatic parity check ran at 11:00 AM on Tuesday morning and you change the setting at 1:00 PM Tuesday afternoon so that it will run at 9:00 PM, automatic parity check will not run until 9:00 PM on Wednesday. This does not affect your ability to run a manual parity check.

Figure 56. Setting the Time for Automatic Parity Check/Repair

Click here to enable or disable Automatic Parity Check/Repair.

Set the time for automatic check/repair here. Pick a time of low system usage. The current setting is displayed.

Select the day of the week you want automatic parity check to run, or select Daily to run it every day.

Click here to save the values currently displayed.

Click here to go back to the application without changing the current settings.

Running a Manual Parity Check

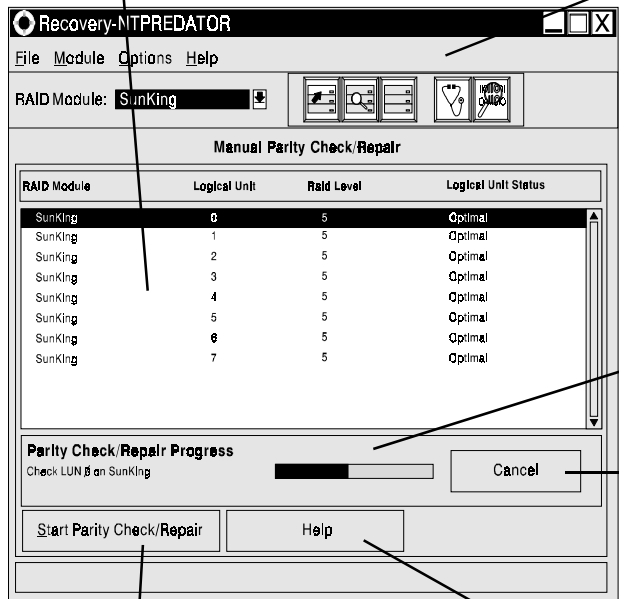
Normally, you should set your RAID Controllers so that Parity Check is run automatically on a regular basis (see ["Setting Automatic Parity Time" on page 122](#) for information on how to set the time for Automatic Parity Check). However, during recovery, or if you suspect a problem, you may need to run manual Parity Check/Repair to find and correct parity errors immediately.

1. Use the following procedure to run Parity Check/Repair manually.
2. Start the Recovery application.

3. Select the RAID Module containing the LUNs you want to check (or select ALL RAID Controllers).
4. Click the Manual Parity Check/Repair button or select Options → Manual Parity Check/Repair from the drop-down menus. A screen like Figure 57 on page 124 is displayed.
5. 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.
6. As each LUN is checked, a histogram bar appears on the screen indicating the Parity Check/Repair progress on that LUN.
7. When Parity Check/Repair is completed, you will see a message indicating if any errors were found.

Select the logical units you want to check here. You can not select a RAID 0 logical unit, or a logical unit with a status other than Optimal.

Click any button here to leave this screen without performing a check.



A histogram appears here to show the progress of Parity Check/Repair on each LUN as it is checked.

Click here to cancel Manual Parity Check.

Click here to start the Parity Check/Repair operation. This box is grayed out if you select a RAID 0 or a non-Optimal logical unit.

Click here to view the help.

Figure 57. Manual Parity Check Repair Screen

Notes:

- 1 You can not run Parity Check/Repair on a RAID 0 logical unit or a non-optimal RAID 1, 3, or 5 logical unit.
- 2 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.
- 3 While parity check/repair is in progress, you can not perform other Recovery tasks.

Changing the LUN Reconstruction Rate

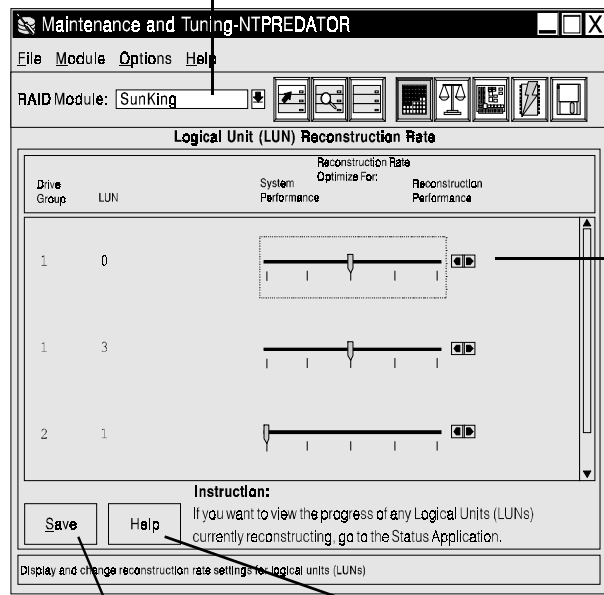
The Reconstruction Rate is the rate at which the controller rebuilds data on a replaced drive. The faster the reconstruction rate, the more time the controller spends rebuilding data and the less time it spends servicing I/O requests from the host. The slower the rate, the faster system I/O occurs, but the longer reconstruction takes to finish. See ["Reconstruction" on page 12](#) for more information.

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/Changing Reconstruction" on page 109](#).

To change the LUN reconstruction, perform the following steps.

1. Start the Maintenance and Tuning application.
2. Select the RAID module you want to change reconstructions rates on. Do not select All RAID Controllers.
3. Click on the LUN Reconstruction button or select Options → LUN Reconstruction from the drop-down menus. A screen like [Figure 58 on page 126](#) 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 will take, but slows system performance.
5. After you have made all the changes you want, click Save to implement the changes.

You must have an individual RAID Controller selected



This display shows all the logical units in the selected RAID Controller and the current reconstruction rate settings for them.

This setting is shown as a slider bar. If you move the tab to the left, reconstruction will be slower, but system performance will be faster. If you move the tab to the right, reconstruction will occur at a faster rate, at the expense of system performance.

Click here to implement the changes.

Click here for help.

Figure 58. Changing the Reconstruction Rate Screen

Notes:

- 1 You can not perform this option if you select All RAID Controllers.
- 2 The new reconstruction rate will go into effect the next time the selected logical units are reconstructed.
- 3 You can use this option to change the reconstruction rate of LUN currently being reconstruction; however, it is best to use the procedure given in "[Monitoring/ Changing Reconstruction](#)" on page 109.
- 4 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.
 - a Slow — 1024 blocks/0.8 second delay
 - b Medium slow — 1024 blocks/0.4 second delay
 - c Medium — 1024 blocks/0.2 second delay
 - d Medium fast — 1024 blocks/0.1 second delay
 - e Fast — reconstruct with no delays.

Balancing LUNs

Balancing LUNs involves assigning drive groups/LUNs to controllers in an active/active redundant controller pair (see ["Redundant Controllers" on page 12](#) for a definition of redundant controllers). Some LUN 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 logical units to change this default balancing, or if you need to reassign logical units after a controller failure.

Logical units are balanced by *drive group*. That is, all 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 1 logical unit each and drive group 3 has 4 logical units, relying on automatic balancing will result in one controller servicing 5 logical units while the other controller services only one.

Using the LUN balancing option, you can either:

- Automatically balance drive groups/logical units (see ["Automatic LUN Balancing" on page 127](#)).
- Manually balance the logical units in one RAID Controller (see ["Manual LUN Balancing" on page 129](#)).

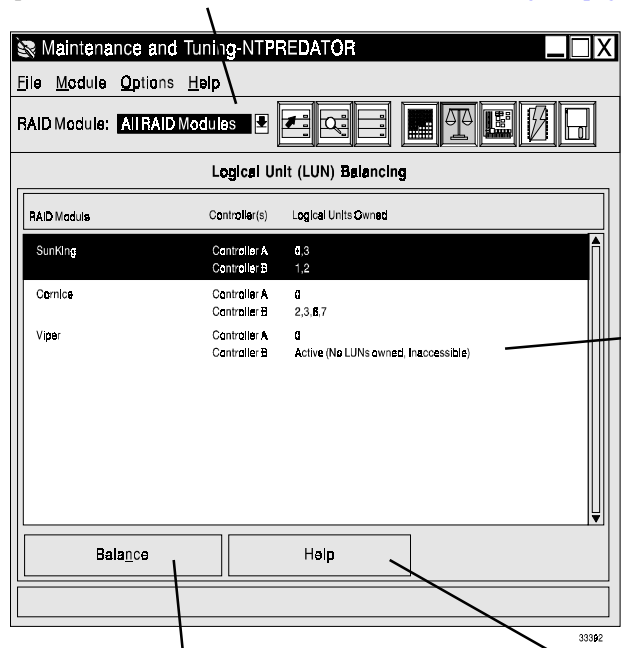
Automatic LUN Balancing

IMPORTANT If a RAID Controller 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 Controllers are eligible for balancing, the Balance button is grayed out.

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

1. Start the Maintenance and Tuning application.
2. Select All RAID Controllers. (If you select a single RAID Controller, you can only balance the logical units manually; see ["Manual LUN Balancing" on page 129](#)).
3. Click on the LUN Balancing button or select Options→ LUN Balancing. A screen like [Figure 59 on page 128](#) is displayed.
4. Select the RAID Controllers you want to balance, then click Balance to automatically balance them.

You must have All RAID Controllers selected here. If you have selected a specific RAID Controller, see ["Manual LUN Balancing"](#) on page 129.



This display shows all the RAID Controllers on your system and the current LUN assignments on them. Select the RAID Controllers 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 Controllers. Click here for help.

Figure 59. Automatic LUN Balancing Screen

Notes:

- 1 You can only select RAID Controllers with active/active controller pairs.
- 2 You can not select controllers in an independent configuration (you must use manual LUN balancing to reassign LUNs in this configuration; see ["Manual LUN Balancing"](#) on page 129).
- 3 Remember that 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.
- 4 Odd-numbered drive groups are assigned to one controller and even numbered drive groups are assigned to the second controller.
- 5 Logical units/drive groups will not be balanced on RAID Controllers without RDAC protection unless all I/O to that module is stopped.
- 6 If none of the displayed RAID Controllers are eligible for balancing, the Balance button is grayed out.

Manual LUN Balancing

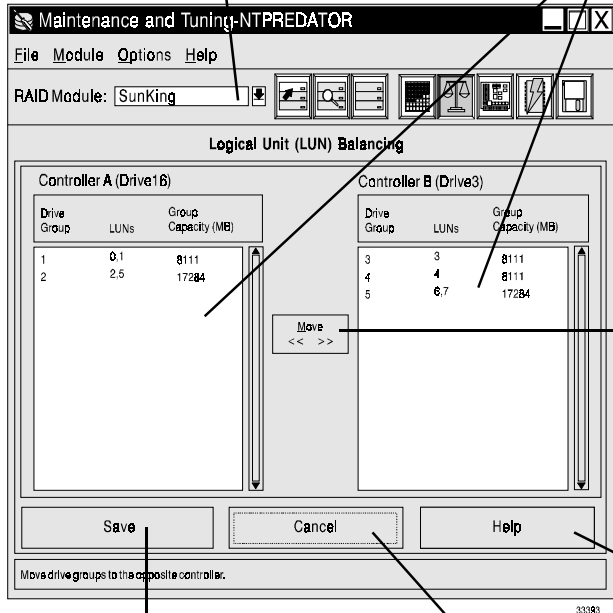
To manually balance the logical units/drive groups in a single RAID Controller, perform the following steps.

CAUTION If the RAID Controller does not have RDAC fail-over protection, you must stop all I/O to that module before balancing the LUNs.

1. Start the Maintenance and Tuning application.
2. Select the individual RAID Controller containing the logical units/drive groups you want to reassign. (If you select All RAID Controllers, you can only automatically balance all the drive groups in all your RAID Controllers; see "[Automatic LUN Balancing](#)" on page 127).
3. Click on the LUN Balancing button or select Options→ LUN Balancing. A screen like [Figure 60 on page 130](#) 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.

You must have a specific RAID Controller selected here. If you have selected ALL RAID Controllers, see "[Automatic LUN Balancing](#)" on page 127.

These two windows show the current logical unit/drive group assignments. To move a drive group from one controller to the other, highlight the drive group and use the Move button.



Use the Move button to move drive groups between screens. Highlight a drive group and then click on the button to move the selected drive group to the other window.

Click here to make the changes you've entered. The drive groups are then reassigned according to the current display.

Click here to exit without changing drive group assignment.

Click here for help.

Figure 60. Manual LUN Balancing Screen

Notes:

- 1 You can only select RAID Controllers with active/active controller pairs.
- 2 If the RAID Controller does not have RDAC protection, you must stop all I/O to the module before balancing the LUNs.
- 3 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 can not assign LUNs from the accessible controller to the inaccessible controller.

Changing the Controller Mode

IMPORTANT You can use this option only if the redundant controllers in the RAID Controller are active/passive.

Use this option to:

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

Note that you can not use this option to change controllers from active/active to active/passive. This can only be done with the rdacutil utility. See the rdacutil entry on [page 136](#) for information on this utility.

To change the controller mode of your controllers, perform the following steps:

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

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

IMPORTANT You will not be able to select this option unless the selected RAID Controller contains an active/passive controller pair, or, if you selected All RAID Controllers, 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 61 on page 132](#) is displayed.
4. Either:
 - Swap the controllers, making the active one passive and the passive one active.
 - Change the controller mode to active/active. A confirmation message will appear, giving you the chance to automatically balance the drive groups/logical units across the two controllers during the mode change.

The selected module must contain a pair of redundant controllers in active/passive mode. If you selected All RAID Controllers, at least one controller pair must be active/passive.

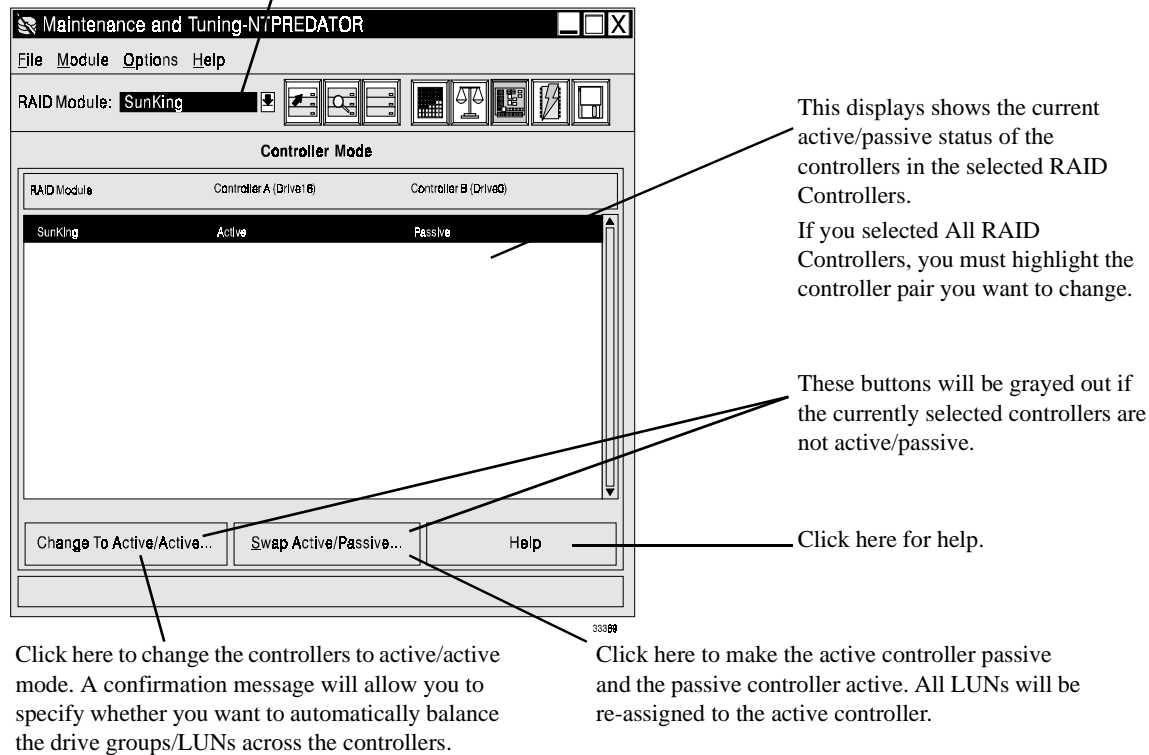


Figure 61. Change Controller Mode Screen

Notes:

- 1 If the RAID Controller does not have RDAC protection, you must stop all I/O to the module before changing controller mode.
- 2 You will not be able to select this option unless the selected RAID Controller contains an active/passive controller pair, or, if you selected All RAID Controllers, at least one controller pair is active/passive.
- 3 You can not change active/active controllers to active/passive using this application. See the rdacutil entry on [page 136](#).
- 4 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 5](#) for more information on cache memory.

You can set three cache memory parameters with this option: write caching, write cache mirroring, and caching without batteries. You can change other cache options using the `raidutil` utility. See the `raidutil` entry on [page 136](#) for information on this utility.

You can only change cache parameters if:

- Both controllers in the RAID Controller (if there are two) have the same size cache memory (at least 128 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 ["Controller Profile Screen" on page 40](#)).

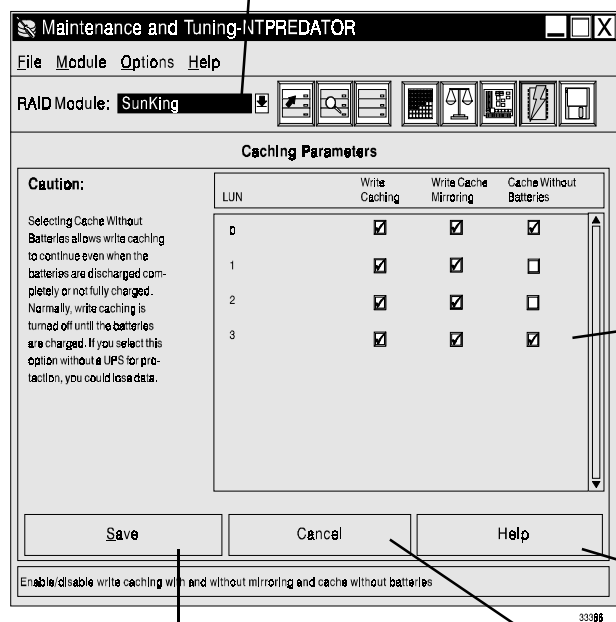
Use the following procedure to change the cache parameters.

1. Start the Maintenance and Tuning application.
2. Select the RAID Controller 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 62 on page 134](#) 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 will automatically select or deselect others. See [Table 19 on page 134](#).
5. After you have made all the changes you want, click Save to apply the changes.

Select an individual RAID Controller here. You can not set cache parameters if All RAID Controllers is selected.



This display shows the cache parameter settings for each logical unit in the selected RAID Controller. If an option is grayed out, it is because the controllers in the RAID Controller do not support that parameter.

Click in the boxes to select or deselect parameters. These parameters are interdependent. Clicking on one may cause others to be enabled or disabled (see Table 19).

Click here to save the current cache parameter settings.

Click here to exit this screen without making any changes.

Click here for help.

Figure 62. Cache Parameters Screen

Table 19: Cache Parameter Interdependencies

If You Select:	The Following Parameters Are Enabled:
Write Caching	Write Cache Mirroring ¹
Write Cache Mirroring ¹	Write Caching
Cache Without Batteries	Write Caching
If You Deselect:	The Following 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 20 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 20: Command Line Utilities and Program Description

Option	Description
Informational	
symism	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 Utility	
drivutil	The drive/LUN utility. This program helps manage 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, fibre channel code, or an NVSRAM file to a specified controller.

Table 20: Command Line Utilities and Program Description (continued)

Option	Description
healthck	The health check utility. This program performs a health check on the indicated RAID Module(s) and displays a report to standard output.
lad	The list array devices utility. This program identifies what RAID controllers and logical units 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.
nvutil	The NVSRAM display/modification utility. This program permits the viewing and changing of 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. (While correct parity is vital to the operation of the array, the possibility of damage to parity is extremely unlikely.)
raidutil	The RAID configuration utility. This program is the command line counterpart to the graphical configuration application. It permits RAID logical unit and hot spare creation and deletion to be performed 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 (Networked version only)	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's names, and clear information in the host store region.
symping	The network connection verification utility. This program verifies that the network connection between the Networked storage management software's host and a RAID Module's controllers is operational. If a failure occurs, the symping utility will display possible reasons.
Background Process Programs and Driver Modules	
arraymon	The array monitor background process. The primary function of the array monitor is to watch for the occurrence of exception conditions in the array and provide administrator notification when they happen.

Table 20: Command Line Utilities and Program Description (continued)

Option	Description
Customizable Elements	
rmparams	<p>This 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, etc. The applications read this file on startup or at select times during their execution. A subset of the parameters in rmparams are changeable under the graphical interface.</p> <p>For more information on this file, see the <i>Installation and Support Handbook</i>.</p>
rmscript	<p>The notification script. A program that 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 e-mail to the superuser/administrator, and, in some cases, sending an SNMP trap. Although you can edit rmscript, make certain that you do not disturb any of the standard actions.</p> <p>For more information on this file, see the <i>Installation and Support Handbook</i>.</p>

Appendix B. Notices

Notices

References in this publication to IBM products, programs, or services do not imply that IBM intends to make these available in all countries in which IBM operates. Any reference to an IBM product, program, or service is not intended to state or imply that only that IBM product, program, or service may be used. Subject to IBM's valid intellectual property or other legally protectable rights, any functionally equivalent product, program, or service may be used instead of the IBM product, program, or service. The evaluation and verification of operation in conjunction with other products, except those expressly designated by IBM, are the responsibility of the user.

IBM may have patents or pending patent applications covering subject matter in this document. The furnishing of this document does not give you any license to these patents. You can send license inquiries, in writing, to:

IBM Director of Licensing
IBM Corporation
N. Castle Drive
Armonk, NY 10504-1785
U.S.A.

Licensees of this program who wish to have information about it for the purpose of enabling: (1) the exchange of information between independently created programs and other programs (including this one) and (ii) the mutual use of the information which has been exchanged, should contact IBM Corporation, Department 80D, P.O. Box 12195, 3039 Cornwallis, Research Triangle Park, NC 27709, U.S.A. Such information may be available, subject to appropriate terms and conditions, including in some cases, payment of a fee.

Trademarks

The following terms are trademarks of the IBM Corporation in the United States or other countries or both:

IBM
Netfinity

Microsoft, Windows, and Windows NT are trademarks of Microsoft Corporation in the United States, other countries, or both.

Intel is a registered trademark of Intel.

Other company, product, and service names may be trademarks or service marks of others.

Index

A

- Active/active controllers 13
 - balancing 127, 130
 - changing to active/passive 132
- Active/passive controllers 12
 - changing mode 26, 130
 - LUN ownership after swap 132
- Add (module) 36
- Appware 114
 - see also Firmware
- arraymon utility 136
- ASC/ASCQ codes 80, 82
- Automatic LUN balancing 127
 - see also Balancing Logical units
- Automatic parity check/repair
 - see Parity check/repair
- autosynch 98
 - controller fault LED and 99

B

- background format 51
- background formatting
 - degraded mode and 51
 - logical unit failure and 51
- Balancing logical units 26, 57, 127
 - LUN creation and 55, 62
 - restrictions 128
- Batteries
 - caching and 5
 - see also Cache without batteries
- batteries
 - caching and 100
 - replacing 100
 - resetting date 100
- Board ID 40
- Board name 40
- Board serial number 40
- Boot level 41
- Bootware 114
 - see also Firmware
- boundary, LUN 50

C

- Cache memory
 - cache options and 133
 - defined 5
 - performance monitor and 89, 90
 - see also Cache options
- Cache options
 - asterisks in display 133
 - cache without batteries 5, 44, 133
 - changing 26
 - interdependencies 134
 - other 133
 - rdacutil utility and 64
 - selecting 57, 60, 133
 - restrictions 133
 - selections for 44
 - write cache mirroring 5, 20, 44, 133
 - restrictions 133
 - write caching 5, 44, 133
- cache parameters
 - asterisks in display 99
 - controller replacement and 99
- Cache parameters option 60
- Cache size 42, 133
- Cache without batteries 133
 - defined 5
 - setting 44, 133
 - see also Cache options
- Channel ID
 - drives 9, 37
- Channels
 - creating LUNs and 20, 60
 - drive 20
 - failures on 20, 94, 106
 - hot spares and 74
 - see also Channel ID
- Command line interface 135
 - see also Command line utilities 135
- Command line utilities
 - drivutil 135
 - fwutil 135
 - healthck 136
 - lad 136
 - logutil 136
 - nvutil 114, 119, 136
 - parityck 136
 - raidutil 133, 136

- rdacutil 6, 64, 131, 132
- storutil 136
- symping 136
- Component failures
 - see Failures
- Configuration
 - fibre channel 13
 - independent controller 17
 - downloading firmware 115
 - multi-host 15
 - network 18
 - overview 47
 - recommended 14
 - resetting 26, 76
 - single-host 14
- Configuration application 55, 72
 - description of 21
 - main screen 52
 - module information display 48
- Configuration Application Main Screen 52
- Configured drives
 - as spares 96
- controller fault LED
 - autosynch and 99
- Controller firmware download utility (fwutil)
 - 135
- Controller ID 41
- Controllers
 - active/active 13, 127, 131
 - active/passive 12, 130
 - assigning 62
 - automatic LUN assignment 55, 58, 127
 - balancing 26, 127
 - changing mode 26, 130
 - firmware 113
 - downloading 116
 - inaccessible status 17
 - manual LUN assignment 127
 - manual recovery options 108
 - names of 6, 25
 - naming 36
 - placing offline 108
 - placing online 108
 - same host bus
 - restrictions 14
 - selecting for download 115
 - unknown status 42

- see also Redundant Controllers
- controllers
 - autosynch of 98
 - fault LED and 99
 - replacing
 - cache parameters and 99
- creating a LUN 54
- Creating logical units
 - see Logical units, creating

D

- Daemon processes
 - arraymon 136
- Data connection
 - failures 12, 97
 - see also Redundant controllers
- Data path
 - see Data connection
- date
 - resetting battery 100
- Date code 43
- Date of manufacture 41
- Dead logical unit 93
- Dead mode 93
- Defragment option
 - defined 64
- defragment option
 - procedure 66
 - restrictions 66
- Degraded logical unit 93
 - defined 12
 - parity check and 122
- degraded logical unit
 - background formatting and 51
- Degraded mode 93
 - failing drives and 105
- Deleting logical units 72
 - network configuration 73
 - restrictions 72
- Device names
 - controllers 6
 - drives 8
 - logical units 6
 - user defined 6
- Drive groups
 - adding drives to 25, 69
 - capacity added 69

- balancing 127
- configured drives 6
- creating 54, 57
- defined 6
- defragmenting 64
- hot spare 6
- locating 37
- modifying 48, 69
- numbering 7
- selecting drives 60
 - restrictions 60
- space allocation on 64
- unassigned 6, 72
- viewing 48
- drive groups
 - defragmenting 66
- Drive selection option 60
- Drive/LUN utility (drivutil) 135
- Drives
 - adding to drive group 69
 - channel number 9, 37
 - channels 20, 94
 - different sizes 55, 60, 73
 - drive groups 6
 - numbering 7
 - failing 104
 - failures on 92
 - firmware 43, 113
 - hot spares 9, 20
 - locating 37
 - manual recovery options 104
 - names 8
 - numbering 37
 - profile information 42
 - reconstructing 104
 - replacing 94
 - reviving 95, 104
 - SCSI ID 8, 37
 - serial number 43
 - status 43
 - using configured drives as spares 96
- drives
 - PFA flags 97
- drivutil utility 135

E

- Edit (module) 36

- Exiting the software 31

F

- Failing drives 104
- Failures
 - accidental 95
 - channel 20, 60, 94, 106
 - controllers 12, 97
 - data connection 12
 - drives 92
 - fans 99
 - general information 19
 - message log and 80
 - power supplies 99
 - preventing 19
 - recovering from 91
 - recovery guru/health check and 19
 - recovery information for 102
 - redundant controllers 12
 - replacing components 27
- failures
 - battery 100
- fault LED
 - autosynch and 99
- features, new software 29
- Fibre channel
 - configuration 13
 - controller ID 41
 - firmware 3.0 and 114, 118
 - firmware file 114
 - host-side ID 41
 - independent controller configuration 17, 18
 - link speed 41
 - multi-host configuration 15, 16
 - NL_Port 41
 - physical connection 41
 - preferred address 41
 - preferred loop ID 41
 - restarting after firmware download 119
 - single-host configuration 15
 - topology 41
 - version 42
 - worldwide port name 41
- Fibre channel level 42
- Find (module) 36
- Firmware
 - boot level 114

- controller 113
- downloading 26
 - confirming download 120
 - general steps 114
 - online/offline method 115
 - procedure 116
 - progress of 118
- drives 113
- fibre channel 114, 118
- firmware level 41, 114
- fwcompat.def file 114
- redundant controller requirements 116
- versions 116
- firmware
 - autosynch of 98
- firmware autosynch 29
- Firmware level 41
 - drives 43
- Formatting logical units 106
- formatting restrictions 51
- formatting, logical unit 49
- fwcompat.def file 114, 116
- fwutil utility 135

G

- GUI
 - navigating in 30
 - online help 31
 - overview 135

H

- Hardware guides
 - when to use 2, 91, 99
- Health Check
 - see Recovery guru/health check
- Health check utility (healthck) 136
- Host bus
 - caution with dual controllers 14
 - see also SCSI bus, Fibre Channel
- Host interface 41
- Host store utility (storutil) 136
- Host-side ID 41
- hot spare drive 9
- Hot spares 20
 - creating 73
 - defined 9

- deleting 26, 75
- drive group 6
- drive sizes to use 73
- independent controllers and 18
- LUN assignment 9
- number allowed 74, 75
- recovery and 94
- restrictions on creating 73
- viewing 48

I

- Inaccessible status 17
- Independent controller configuration 17, 18
 - hot spares and 18
 - inaccessible status and 17
 - rdac support for 17
 - recovery and 17
 - see also RAID Modules, configuration
- Installation and support guide
 - description of 2
 - software installation and 2
 - troubleshooting information in 2

K

- Keyboard 30

L

- lad utility 136
- LED
 - controller fault LED and autosynch 99
- Link speed 41
- List array devices utility (lad) 136
- Locating
 - drive groups 37
 - drives 37
 - logical units 37
 - RAID modules 36
- Log format utility (logutil) 136
- Logical Unit Formatting 49
- logical unit formatting 49
- Logical units
 - adding drives to 69
 - cache options
 - selecting 57, 60, 133
 - capacity
 - defaults 60

- setting 57, 59
- changing RAID level 67
- checking parity on 122
- controller ownership 13, 26
 - changing 127
- creating 25, 47
 - controller assignment 62
 - drive group numbering 7
 - from unassigned drives 54
 - in existing drive group 57
 - operating system and 47, 57, 58
 - overview 47
 - RAID level 55
 - segment size 57, 61
 - selecting drives 60
 - space available for 65
- dead 93
- defined 9
- degraded 12, 93
- deleting 26, 72
 - drive group numbering 7
 - operating system and 72
- drive groups 6
- formatting 106
- locating 25, 37
- manual recovery options 106
- modifying 47
- names of 6
- number allowed 14, 17
- profile information 43
- reconstruction rate of 109, 125
- recovering 93
- reviving 106
 - when 107
- segment size 62
 - changing 71
- status 44
- viewing 48

logical units

- dead
 - formatting and 51
- degraded
 - formatting and 51

logutil utility 136

LUN

- see Logical units

LUN assignment option 62

LUN boundary 50

LUN capacity option 59

LUN creation 29

LUN creation and 57

M

Maintenance and Tuning application 116, 133

- description of 24
- overview 113

Manual LUN balancing 129

- see also Balancing Logical units

Manual parity check/repair

- see Parity check/repair

Manual recovery options 103

- controllers 108
- drives 104
- logical units 106
- when to use 92

Maximum size limit

- effects when exceeded 31, 85
- message file and 79, 84
- resetting file 85
- warning message 85

Message log file 79, 92

- backup file names 85
- changing settings 84
- clearing 31, 85
- copying 83
- default 80
- maximum size limit 31, 79, 84, 85
- message types 80
- recovery guru and 80
- saving 83
- viewing 80

message log file

- PFA flags and 96

Module information display 48

module information display 52

Module profile

- drives 42
- logical units 43
- saving 44
- viewing 38
- see also RAID modules

Mouse 30

Multi-host configuration 15, 16

- number of LUNs on 16

- operations on 15
- rdac support for 15
- reservations and 16
- see also RAID Modules, configuration

Multiple drive failures 93

N

Navigation 30

- keyboard 30
- mouse 30
- online help 31

Network configuration 18

- adding modules 36
- deleting logical units 73
- exclusive access and 19
- naming controllers 36
- RDAC and 18
- storutil utility 136
- symping utility 136
- see also RAID Modules, configuration

Network connection verification utility (symping) 136

Network management station 13

new software features 29

NL_Port 41

NMS 13

Number of drives

- setting 55

Number of LUNs

- setting 55, 57

NVSRAM 116

- autosynch and 99
- defined 114
- downloading 26, 115, 116
- nvutil utility 136
- resetting 76
- restart after downloading 115, 119

NVSRAM utility (nvutil) 136

nvutil utility 114, 119, 136

O

Offline controller 108

Offline method (firmware download) 115

Online controller 108

Online help 31

Online method (firmware download) 115

Operating system

- dead mode and 93
- degraded mode and 93
- deleting logical units and 72
- device names on 6
- installation and support guide 2
- logical unit creation and 47, 57, 58
- logical unit deletion and 72
- messages from 92
- RDAC on 11, 12
- recovery and 92

P

Parity 122

- defined 10
- how used 122
- RAID 1 and 10, 122
- RAID 3 10, 11, 122
- RAID 5 10, 11, 122
- reconstruction and 12
- see also Parity Check/Repair

Parity check/repair

- automatic 122
 - current setting 123
 - disabling 123
 - setting time 27, 122
- corrupted data and 122
- defined 10
- eligible LUNs 125
- manual 122
 - restrictions 124
 - running 123
 - when to run 27, 123
- parityck utility 136
- RAID 1 and 10
- RAID 3 10
- RAID 5 10
- restrictions 122, 125
- running 122
- viewing results of 27, 80
- what is repaired 122

Parity check/repair utility (parityck) 136

Performance monitor 86

- polling interval 87
- saving data from 89
- screen elements 89
- starting 86

- using data from 90
 - when to run 25
 - PFA 96
 - PFA flags 96
 - fixing drives 97
 - message log and 96
 - message type 96
 - Physical connection 41
 - Polling interval
 - performance monitor 87
 - status 79
 - changing 84
 - Predictive Failure Analysis 96
 - Preferred address 41
 - Preferred loop ID 41
 - Processor size 42
 - Product ID 40
 - Product revision 41
 - Product serial number 40
- R**
- RAID 0
 - dead mode 93
 - defined 10
 - failing drives and 105
 - parity check and 122
 - recovery on 93
 - reviving drives and 95
 - see also RAID level
 - RAID 1
 - dead mode 93
 - defined 11
 - degraded mode 93
 - failing drives and 105
 - parity and 10
 - parity check/repair and 10, 122
 - recovery on 93
 - reviving drives and 95
 - when to use 19
 - see also RAID level
 - RAID 3
 - dead mode 93
 - defined 11
 - degraded mode 93
 - failing drives and 105
 - parity 10
 - parity check/repair 10, 122
 - recovery on 93
 - reviving drives and 95
 - when to use 19
 - see also RAID level
 - RAID 5
 - dead mode 93
 - defined 11
 - degraded mode 93
 - failing drives and 105
 - parity 10
 - parity check/repair 10, 122
 - recovery on 93
 - reviving drives and 95
 - when to use 19
 - see also RAID level
 - RAID configuration utility (raidutil) 136
 - RAID level
 - changing 26, 67
 - restrictions 67
 - RAID 0 10
 - RAID 1 11
 - RAID 3 11
 - RAID 5 11
 - setting 55
 - when to use 19
 - see also individual RAID level entries
 - RAID modules
 - adding 36
 - configuration
 - independent 17
 - multi-host 15
 - network 18
 - overview 47
 - single-host 14
 - failure recovery 91
 - finding 36
 - identifying 34
 - locating 36
 - module information display 48, 52
 - naming 25, 36
 - profile information 25, 38
 - saving 44
 - removing 36
 - selecting 34
 - tasks 25
 - raidcode.txt file 135
 - raidutil utility 133, 136

- resetting battery date 100
 - RDAC
 - controller failures and 97
 - controller mode changes and 132
 - defined 11
 - failover conditions 98
 - firmware download and 115
 - independent configuration and 17
 - LUN balancing and 127, 128, 129, 130
 - multi-host configuration and 15
 - overview 135
 - redundant controllers and 12
 - single-host configuration and 14
 - rdac text file 135
 - rdacutil utility 6, 64, 131, 132
 - Reconstructing drives 104
 - Reconstruction 104, 105
 - defined 12
 - monitoring 96, 109
 - rate of 12
 - Reconstruction rate
 - changing 27, 96, 109
 - restrictions 126
 - defined 12
 - settings 111, 126
 - Recovery
 - component failures and 19
 - failed controllers 97
 - failed drives 92
 - failed fans 99
 - failed power supplies 99
 - independent controllers and 17
 - instructions for 102
 - manual recovery options 103
 - overview 91
 - reconstruction rate 109
 - tasks 27
 - recovery
 - batteries 100
 - LUN formatting and 51
 - Recovery application 100, 104
 - description of 23
 - overview 91
 - Recovery Guru/Health Check
 - PFA flags and 96
 - Recovery guru/health check 19
 - independent configuration and 17
 - message log and 80
 - show procedures 102
 - starting 100
 - viewing procedures 91
 - when to use 27, 92
 - Redundant controller management utility (raid-util) 136
 - Redundant controllers
 - active/active 13, 131
 - active/passive 12, 131
 - changing mode of 130
 - changing status of 109
 - described 12
 - failover support 15, 17, 18
 - firmware requirements 116
 - LUN ownership and 127
 - recommended configuration 13
 - see also Controllers
 - see also RDAC
 - Remove (module) 36
 - Reservations 16
 - Resetting configuration 76
 - Reviving drives 95, 104
 - Reviving logical units 106
 - rmevent file 135
 - rmparams file 137
 - rmscript
 - description 135
 - file 137
- ## S
- SCSI bus
 - independent controller configuration 17, 18
 - multi-host configuration 15, 16
 - single-host configuration 14
 - see also Host Bus
 - SCSI ID
 - drives 8, 37
 - see also Host-side ID
 - Segment size
 - changing 26, 71
 - defined 13
 - selecting 57, 61
 - Segment size option 61
 - Self Monitoring and Reporting Technology 96
 - Serial number
 - board 40

- controller 39
- drives 43
- product 40
- Show procedures 102
 - see also Recovery guru/health check
- Simple Network Management Protocol
 - see SNMP
- Single drive failures 93
- Single-host configuration 14
 - rdac support for 14
 - see also RAID Modules, configuration
- SMART 96
- SMART PFA reporting 96
- SMART technology 96
- SNMP 80, 92
 - defined 13
- Software
 - exiting 31
 - installing 2
 - online help for 31
 - starting 30
- Starting the software 30
- Status 79, 92
 - dead mode 93
 - degraded mode 93
 - drives 43
 - inaccessible 17
 - logical units 44
 - message file 25, 79, 92
 - backup file name 85
 - changing settings 84
 - clearing 31, 85
 - copying 83
 - default 80
 - maximum size limit 31, 79, 85
 - message types 80
 - opening 83
 - viewing 80
 - polling 79, 84
 - unknown 42
- Status application 86
 - description of 22
- Status log 25, 79

- clearing 31, 85
 - see also Message log file
- Stopping the software 31
- storutil utility 136
- symping utility 136
- symism file 135

T

- Tasks 25
- Topology 41
- Troubleshooting 2
 - see also Installation and Support Guide

U

- Unassigned drive group 6, 72
 - creating LUNs from 54
 - viewing 48
- Unknown status 42
- User guide
 - description of 1

V

- Vendor
 - drives 43
- Vendor ID
 - fibre channel controllers 41

W

- Worldwide port name 41
- Write cache mirroring 133
 - defined 5
 - failures and 98
 - restrictions 133
 - setting 44, 133
 - when to use 20
 - see also Cache options
- Write caching 133
 - defined 5
 - failures and 98
 - setting 44, 133
 - see also Cache options