

Sun™ StorEdge™ RAID Manager 6.2 User's Guide



THE NETWORK IS THE COMPUTER™

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

Part No. 805-6887-10
September 1998, Revision A

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

Copyright 1998 Sun Microsystems, Inc., 901 San Antonio Road • Palo Alto, CA 94303 USA. All rights reserved.
Portions copyright 1998 Symbios Logic, Inc. All rights reserved.

This product or document is protected by copyright and distributed under licenses restricting its use, copying, distribution, and decompilation. No part of this product or document may be reproduced in any form by any means without prior written authorization of Sun and its licensors, if any. Third-party software, including font technology, is copyrighted and licensed from Sun suppliers.

Parts of the product may be derived from Berkeley BSD systems, licensed from the University of California. UNIX is a registered trademark in the U.S. and other countries, exclusively licensed through X/Open Company, Ltd.

Sun, Sun Microsystems, the Sun logo, AnswerBook, Java, the Java Coffee Cup, and Solaris are trademarks, registered trademarks, or service marks of Sun Microsystems, Inc. in the U.S. and other countries. All SPARC trademarks are used under license and are trademarks or registered trademarks of SPARC International, Inc. in the U.S. and other countries. Products bearing SPARC trademarks are based upon an architecture developed by Sun Microsystems, Inc.

The OPEN LOOK and Sun™ Graphical User Interface was developed by Sun Microsystems, Inc. for its users and licensees. Sun acknowledges the pioneering efforts of Xerox in researching and developing the concept of visual or graphical user interfaces for the computer industry. Sun holds a non-exclusive license from Xerox to the Xerox Graphical User Interface, which license also covers Sun's licensees who implement OPEN LOOK GUIs and otherwise comply with Sun's written license agreements.

RESTRICTED RIGHTS: Use, duplication, or disclosure by the U.S. Government is subject to restrictions of FAR 52.227-14(g)(2)(6/87) and FAR 52.227-19(6/87), or DFAR 252.227-7015(b)(6/95) and DFAR 227.7202-3(a).

DOCUMENTATION IS PROVIDED "AS IS" AND ALL EXPRESS OR IMPLIED CONDITIONS, REPRESENTATIONS AND WARRANTIES, INCLUDING ANY IMPLIED WARRANTY OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE OR NON-INFRINGEMENT, ARE DISCLAIMED, EXCEPT TO THE EXTENT THAT SUCH DISCLAIMERS ARE HELD TO BE LEGALLY INVALID.

Copyright 1998 Sun Microsystems, Inc., 901 San Antonio Road • Palo Alto, CA 94303 Etats-Unis. Tous droits réservés.
Des portions de ce produit sont protégées par un copyright (1998) de Symbios Logic, Inc. Tous droits réservés.

Ce produit ou document est protégé par un copyright et distribué avec des licences qui en restreignent l'utilisation, la copie, la distribution, et la décompilation. Aucune partie de ce produit ou document ne peut être reproduite sous aucune forme, par quelque moyen que ce soit, sans l'autorisation préalable et écrite de Sun et de ses bailleurs de licence, s'il y en a. Le logiciel détenu par des tiers, et qui comprend la technologie relative aux polices de caractères, est protégé par un copyright et licencié par des fournisseurs de Sun.

Des parties de ce produit pourront être dérivées des systèmes Berkeley BSD licenciés par l'Université de Californie. UNIX est une marque déposée aux Etats-Unis et dans d'autres pays et licenciée exclusivement par X/Open Company, Ltd.

Sun, Sun Microsystems, le logo Sun, AnswerBook, Java, le logo Java Coffee Cup, et Solaris sont des marques de fabrique ou des marques déposées, ou marques de service, de Sun Microsystems, Inc. aux Etats-Unis et dans d'autres pays. Toutes les marques SPARC sont utilisées sous licence et sont des marques de fabrique ou des marques déposées de SPARC International, Inc. aux Etats-Unis et dans d'autres pays. Les produits portant les marques SPARC sont basés sur une architecture développée par Sun Microsystems, Inc.

L'interface d'utilisation graphique OPEN LOOK et Sun™ a été développée par Sun Microsystems, Inc. pour ses utilisateurs et licenciés. Sun reconnaît les efforts de pionniers de Xerox pour la recherche et le développement du concept des interfaces d'utilisation visuelle ou graphique pour l'industrie de l'informatique. Sun détient une licence non exclusive de Xerox sur l'interface d'utilisation graphique Xerox, cette licence couvrant également les licenciés de Sun qui mettent en place l'interface d'utilisation graphique OPEN LOOK et qui en outre se conforment aux licences écrites de Sun.

CETTE PUBLICATION EST FOURNIE "EN L'ETAT" ET AUCUNE GARANTIE, EXPRESSE OU IMPLICITE, N'EST ACCORDEE, Y COMPRIS DES GARANTIES CONCERNANT LA VALEUR MARCHANDE, L'APTITUDE DE LA PUBLICATION A REPONDERE A UNE UTILISATION PARTICULIERE, OU LE FAIT QU'ELLE NE SOIT PAS CONTREFAISANTE DE PRODUIT DE TIERS. CE DENI DE GARANTIE NE S'APPLIQUERAIT PAS, DANS LA MESURE OU IL SERAIT TENU JURIDIQUEMENT NUL ET NON AVENU.



Contents

- 1. Getting Started 1**
 - About This User's Guide 2
 - About The *Installation And Support Guide* 3
 - About The Hardware Guides 3
 - About Software Installation 4
 - About RAID Manager For DOS 4
 - What To Do Next 5

- 2. Overview 7**
 - Common Definitions 8
 - Cache Memory 9
 - Device Name 9
 - Drive Group 10
 - Drive Group Numbering 11
 - Drive Names 13
 - Hot Spare 13
 - Logical Unit (LUN) 13
 - Parity 14
 - Parity Check/Repair 14
 - RAID Level 15

RAID Module	15
RDAC Driver	15
Reconstruction	16
Redundant Controllers	16
Segment Size	17
SNMP	17
Module Configurations	18
Single-Host Configuration	19
Multi-Host Configuration	20
Independent Controller Configuration	21
Special Network Considerations	24
Component Failures	25
Protecting Yourself From Component Failures	25
Application Summary Charts	26
Task List	31
What's New In 6.2?	34
3. Common Application Features	35
Navigation	36
Starting and Exiting An Application	37
Starting An Application	37
Exiting An Application	37
Using Online Help	38
Selecting a Module	41
RAID Module Selection Screen	42
Locating a Module	44
Locating Drives	45
Viewing A Module Profile	47

The Controller Profile Screen	50
The Drive Profile Screen	52
The LUN Profile Screen	54
Saving A Module Profile	56
4. Configuring RAID Modules	57
Configuration Overview	58
Types Of Logical Unit (LUN) Creation	58
Types Of Logical Unit (LUN)/Drive Group Modification	59
The Module Information Display	60
Creating Logical Units (LUNs)	62
Creating A Logical Unit From Unassigned Drives	62
Creating A LUN From Remaining Capacity	65
Using the Options Screens	67
Modifying Drive Groups and LUNs	73
The Defragment Option	74
Changing RAID Level	76
Adding Drives	78
Modifying Segment Size	80
Defragmenting The Drive Group	82
Deleting LUNs	83
Creating Hot Spares	85
Deleting A Hot Spare	88
Resetting the Configuration	89
5. Checking Status	91
Status Overview	92
Background Polling	92
The Message Log File	92

Message Types	93
The Status Log And Health Check	93
Event Notification	94
Viewing The Log	94
Changing Log File Settings	99
Clearing The Message File	101
Using The Performance Monitor	103
Using Performance Monitor Data	107
6. Recovery	109
Recovery Overview	110
Recovery Procedures	110
General Recovery Steps	111
Recovering From Drive Failures	112
Recovering From Controller/Connection Failures	116
Recovering From Fan/Power Supply Failures	117
Running Recovery Guru/Health Check	118
Manual Recovery Options	121
Manual Recovery/Drives	122
Manual Recovery/Logical Units	124
Manual Recovery/Controllers	126
Monitoring/Changing Reconstruction	128
7. Maintenance and Tuning	131
Maintenance And Tuning Overview	132
Upgrading Controller Firmware	133
Controller Firmware Files	133
General Procedure	134
Online And Offline Method	134

Selecting Controllers	135
Download Procedure	136
Confirming The Download	140
Running Parity Check	142
Setting Automatic Parity Time	143
Running Manual Parity Check	145
Changing The LUN Reconstruction Rate	147
Balancing LUNs	149
Automatic LUN Balancing	149
Manual LUN Balancing	151
Changing Controller Mode	153
Changing Cache Parameters	155
A. Command Line Interface	159
Command Line Commands	160
Index	163

Figures

FIGURE 2-1	Drive Groups and Logical Units (LUNs)	11
FIGURE 2-2	Single-Host To RAID Module Configuration	19
FIGURE 2-3	Multi-host To RAID Module Configuration	21
FIGURE 2-4	Independent Controller Configuration	23
FIGURE 2-5	Configuration Application Summary	27
FIGURE 2-6	Status Application Summary	28
FIGURE 2-7	Recovery Application Summary	29
FIGURE 2-8	Maintenance And Tuning Application Summary	30
FIGURE 3-1	Application Icons	37
FIGURE 3-2	Main Online Help Screen	38
FIGURE 3-3	RAID Module Selection Screen	42
FIGURE 3-4	List/Locate Drives Screen	46
FIGURE 3-5	Module Profile Screen	48
FIGURE 3-6	Controller Profile Screen (SCSI)	50
FIGURE 3-7	Drive Profile Screen	52
FIGURE 3-8	LUN Profile Screen	54
FIGURE 3-9	Save Module Profile Screen	56
FIGURE 4-1	Configuration Application Module Information Display	60
FIGURE 4-2	Create LUN Main Screen (Unassigned Drives)	64

FIGURE 4-3	Create LUN Main Screen (Existing Drive Group)	66
FIGURE 4-4	LUN Capacity Screen	68
FIGURE 4-5	Drive Selection Screen	69
FIGURE 4-6	Cache Parameters Screen	70
FIGURE 4-7	Segment Size Screen	71
FIGURE 4-8	Set LUN Assignment Screen	72
FIGURE 4-9	Defragment Logical Units #1	74
FIGURE 4-10	Defragment Logical Units #2	74
FIGURE 4-11	Defragment Logical Units #3	75
FIGURE 4-12	Defragment Logical Units #4	76
FIGURE 4-13	Modify RAID Level Screen	77
FIGURE 4-14	Modify Group Capacity Screen	79
FIGURE 4-15	Modify Segment Size Screen	81
FIGURE 4-16	Delete LUNs Screen	84
FIGURE 4-17	Create Hot Spare Screen	86
FIGURE 4-18	Select Hot Spare Screen	87
FIGURE 4-19	Delete Hot Spares Screen	88
FIGURE 5-1	Selecting Message Type	95
FIGURE 5-2	Message Log Screen	96
FIGURE 5-3	Message Details Screen	98
FIGURE 5-4	Log Settings Screen	100
FIGURE 5-5	Message Limit Warning Message	102
FIGURE 5-6	Message Log Save Message	102
FIGURE 5-7	Performance Monitor Settings Screen	103
FIGURE 5-8	Performance Monitor Start-Up Screen	104
FIGURE 5-9	Performance Monitor Data Screen	105
FIGURE 6-1	Recovery Main Screen	118
FIGURE 6-2	Recovery Screen/Exceptions Found	119

FIGURE 6-3	Recovery Procedure Screen	120
FIGURE 6-4	Manual Recovery/Drives Screen	123
FIGURE 6-5	Manual Recovery/Logical Units Screen	125
FIGURE 6-6	Manual Recovery/Controllers Screen	127
FIGURE 6-7	Viewing/Changing Reconstruction Screen	129
FIGURE 7-1	Firmware Upgrade Screen	137
FIGURE 7-2	Setting The Time For Automatic Parity Check/Repair	144
FIGURE 7-3	Manual Parity Check Repair Screen	146
FIGURE 7-4	Changing The Reconstruction Rate Screen	148
FIGURE 7-5	Automatic LUN Balancing Screen	150
FIGURE 7-6	Manual LUN Balancing Screen	152
FIGURE 7-7	Change Controller Mode Screen	154
FIGURE 7-8	Cache Parameters Screen	156

Tables

TABLE 2-1	Storage Management Task List	31
TABLE 3-1	Using A Mouse And Keyboard	36
TABLE 3-2	Online Help Screen Elements	39
TABLE 3-3	Select Module Screen Elements	43
TABLE 3-4	Module Profile Screen Contents	49
TABLE 3-5	Controller Profile Screen Contents (SCSI)	51
TABLE 3-6	Drive Profile Screen Contents	53
TABLE 3-7	LUN Profile Screen Contents	55
TABLE 4-1	Configuration Display Columns	61
TABLE 4-2	Adding Capacity Formulas	78
TABLE 5-1	Message Types	93
TABLE 5-2	Message Log Explanation	97
TABLE 5-3	Performance Monitor Screen Contents	105
TABLE 6-1	Drive Failures And Recovery (Summary)	113
TABLE 6-2	Recovering From Failed Controllers (Summary)	117
TABLE 7-1	Firmware Confirmation Screen Elements	138
TABLE 7-2	Controller Firmware Upgrade: Failed Statuses	140
TABLE 7-3	Cache Parameter Interdependencies	157
TABLE A-1	Command Line Utilities And Program Description	160

Getting Started

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

- About This User's Guide—page 2
- About The Installation And Support Guide—page 3
- About The Hardware Guides—page 3
- About Software Installation—page 4
- About RAID Manager For DOS—page 4
- What To Do Next—page 5

About This User's Guide

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

Note – Screen and menu illustrations throughout this document are provided only as examples that simulate the Sun™ StorEdge™ RAID Manager graphical user interface.

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

About The *Installation And Support Guide*

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

The *Installation and Support Guide* contains:

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

You must refer to the *Installation and Support Guide* for installation instructions before doing anything else on the RAID Modules as described in this book.

About The Hardware Guides

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

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

Always follow the software recovery procedures given in this *User's Guide*, 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 Support Guide* that you received along with this book. This document is specific to the operating system you are using. Refer to the *Installation and Support Guide* for complete information on installation, then return to this *User's Guide* for information on how to run the installed software.

About RAID Manager For DOS

Along with the RAID Manager software, you received a program called RAID Manager for DOS. RAID Manager for DOS (DRM) communicates with the RAID Module controllers directly over an RS232 serial link. Use RAID Manager for DOS if:

- You want to create your boot device on a RAID Module.
- You need to connect directly to the RAID Module controller for troubleshooting.

See the *RAID Manager for DOS User Guide* for information on how to run this program.

What To Do Next

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

- 1. Install the RAID Module hardware, using the information given in the hardware guides.**
- 2. Install the RAID Manager software, using the information in the *Installation and Support Guide*.**
- 3. If necessary, read Chapter 2 to become acquainted with the RAID Manager 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 Modules.**
- 4. If necessary, read Chapter 3 to familiarize yourself with the common program features.**
- 5. Go to Chapter 4 to configure the RAID Modules into logical units and hot spares, according to your system needs.**

Overview

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 Modules.

- Common Definitions—page 8
- Module Configurations—page 18
- Component Failures—page 25
- Application Summary Charts—page 26
- Task List—page 31
- What's New In 6.2?—page 34

Common Definitions

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

- Cache Memory—page 9
- Device Name—page 9
- Drive Group—page 10
- Drive Group Numbering—page 11
- Drive Names—page 13
- Hot Spare—page 13
- Logical Unit (LUN)—page 13
- Parity—page 14
- Parity Check/Repair—page 14
- RAID Level—page 15
- RAID Module—page 15
- RDAC Driver—page 15
- Reconstruction—page 16
- Redundant Controllers—page 16
- Segment Size—page 17
- SNMP—page 17

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 155).

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 command line `rdacutil` command. See Appendix A.

Device Name

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

Device names are determined by the location of the subsystem hardware and vary according to the operating system you are using.

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

Drive Group

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

- Unassigned drive group — the drives in the RAID Module that have not been configured into logical units or hot spares. This drive group is only displayed in the Configuration Application.
- Hot spare drive group — drives that have been assigned as hot spares. This drive group is only displayed in the Configuration Application.
- Configured drive group — drives that have been configured into one or more logical units 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 2-1 Drive Groups and Logical Units (LUNs)

Drive Group Numbering

Each drive group in a RAID Module 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.

Tip – 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

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

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

Hot 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 Module. 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, but instead can be used for any failed drive in the RAID Module with the same or smaller capacity.

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

See “Creating Hot Spares” on page 85 for information on how to create a hot spare. See “Protecting Yourself From Component Failures” on page 25 for more information on the advantages of hot spares.

Logical Unit (LUN)

A logical unit (called a “LUN” for logical unit number) is the basic structure you create on the RAID Module to retrieve and store your data. A LUN is a set of physical drives (that 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 2-1 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 Parity Check” on page 142 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.
- 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 MB drives has a data capacity of 72 MB).

RAID Module

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

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

RDAC Driver

The Redundant Disk Array Controller (RDAC) driver is part of the storage management software package and manages the I/O data connection for RAID Modules with redundant controllers. If a component (cable, controller, host adapter,

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 Support Guide* for your operating system for specific information.

Reconstruction

Reconstruction is the process of using data and parity on all operational drives within a drive group to 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

Most RAID Modules contain 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 Module continues to operate. You can then replace the failed controller, often without shutting down the RAID Module, to resume normal operation.

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

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

Whether or not your RAID Modules have the redundant controller feature depends on a number of things:

- Whether the hardware supports it. See the hardware documentation for your RAID Modules to determine whether or not your hardware supports the redundant controller feature.
- Whether your operating system supports RDAC. See the *Installation and Support Guide* for your operating system to determine if your operating system supports the redundant controller feature.
- How the RAID Modules are connected. See “Module Configurations” on page 18 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 Support Guide* for your operating system for details on enabling this notification option.

Module Configurations

The storage management software supports three main configurations from the host systems to the RAID Modules. Be sure to consult the *Installation and Support Guide* specific to your operating system for details on hardware requirements required to use this software with that operating system.



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

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

Use the following figures and information to review the three configurations supported:

- Single-Host Configuration—page 19
- Multi-Host Configuration—page 20
- Independent Controller Configuration—page 21
- Special Network Considerations—page 24

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 Module. The two connections are required for maximum RDAC failover support for redundant controllers (see FIGURE 2-2).

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.

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

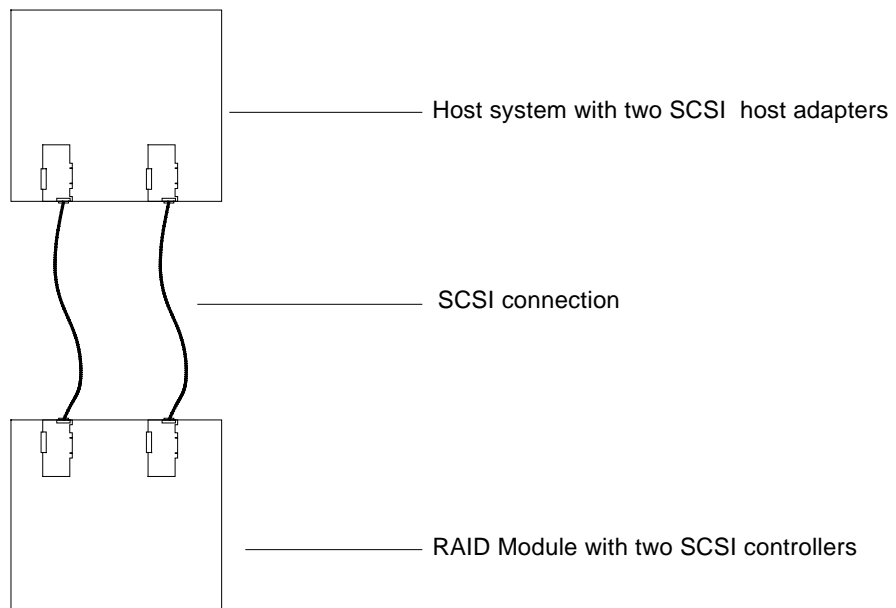


FIGURE 2-2 Single-Host To RAID Module Configuration

Multi-Host Configuration

In a Multi-host configuration, two host machines are each connected by two connections to both of the controllers in a RAID Module (see FIGURE 2-3).



Caution – Not every operating system supports this configuration. Be sure to consult the restrictions in the *Installation and Support Guide* 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 Module, plus RDAC failover support for the redundant controllers. However, in this configuration, you must use caution when performing storage management tasks (especially creation and deletion of LUNs) to ensure the two hosts do not send conflicting commands to the controllers in the RAID Modules.

The following items are unique to this configuration:

- Both hosts *must* have the same operating system and storage management software versions installed.
- Both host machines should have the same LUNs-per-host adapter capacity (that is, either both are limited to eight LUNs or both can have 16-32 LUNs). This is important for RDAC failover situations so that each controller can take over for the other and display all configured drive groups/LUNs.
- If the operating system on the host machine is capable of creating reservations, the storage management software will honor them. This means that each host could have reservations to specified drive groups/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.

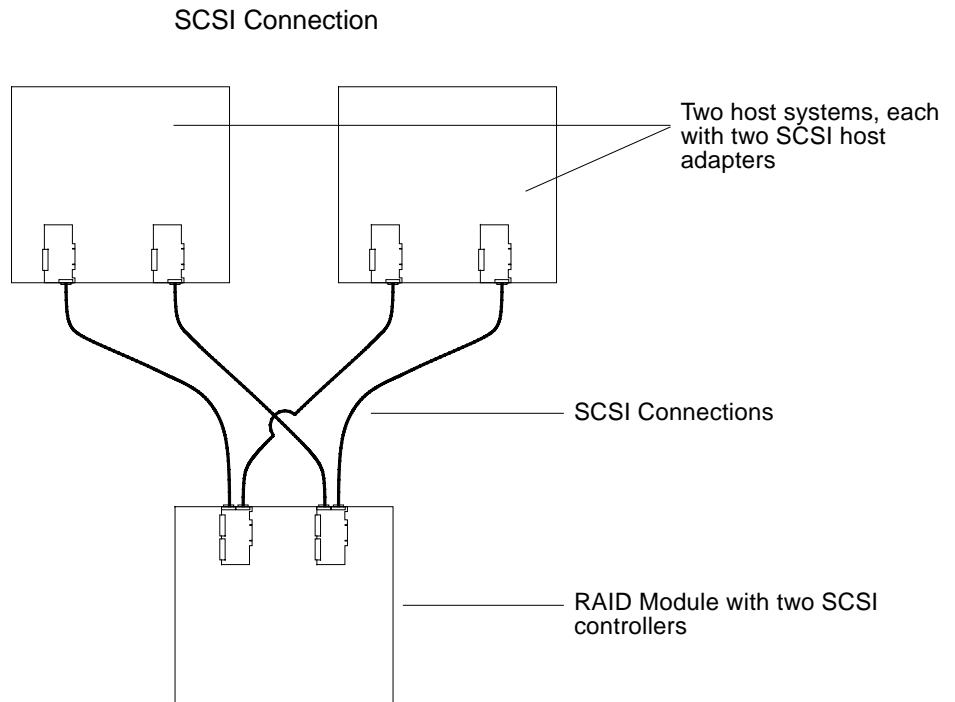


FIGURE 2-3 Multi-host To RAID Module Configuration

Independent Controller Configuration

In the Independent controller configuration, two host machines are connected to a dual-controller RAID Module. One host machine is connected to one RAID Module controller, and a second host machine is connected to the other controller (see FIGURE 2-4.) 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 Module.

Also, when an independent controller is selected with Select Module, the storage management software has knowledge of the alternate controller and displays all configured drive groups/LUNs. It only reports real-time statuses 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.

Note – 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 Module do not have RDAC failover protection.
- The storage management software reports the alternate controller and its drive group/LUNs as “Inaccessible.”
- Recovery Guru/Health Check (Recovery Application) detects data 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.

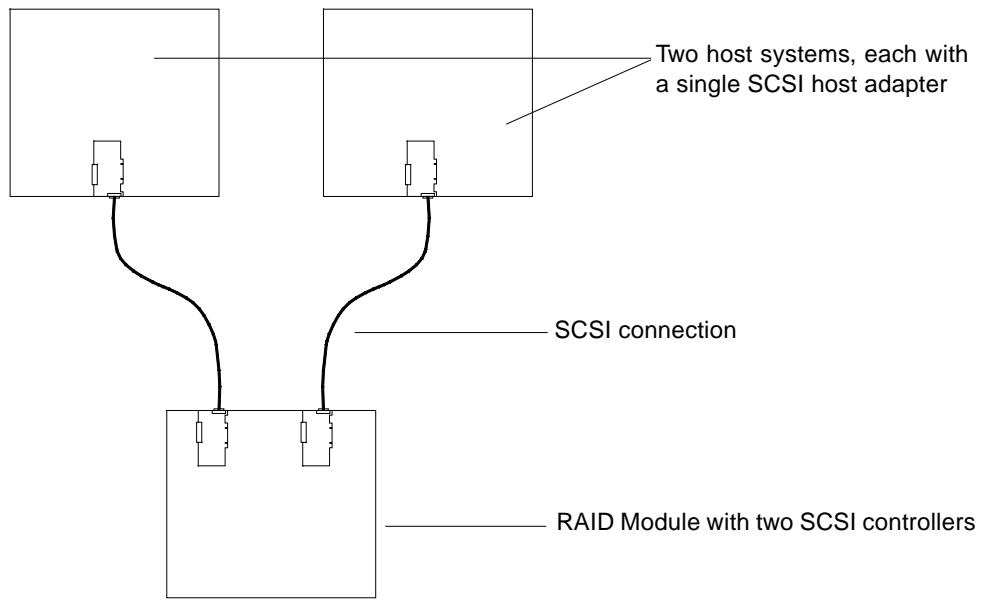


FIGURE 2-4 Independent Controller Configuration

Special Network Considerations

The Networked version of the storage management software always sees both controllers in a dual-controller RAID Module regardless of which configuration you have. However, the Networked version will be able to tell if the currently selected RAID Module 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 Module configurations:

- Your RAID Modules do not have RDAC failover protection unless there is failover protection installed on the host connected to the modules through the SCSI Bus connection.
- This software does *not* provide I/O connection failure detection or recovery. However, any problems with a network 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 Modules serviced by the storage management software are designed to be fault tolerant. All modules provide drive redundancy, so that if a drive fails, the data contained on the drive is still available. In most cases, the modules also feature redundant controllers, fans, and power supplies as well, so that the RAID Modules can keep functioning even if a single component fails. The controllers in some RAID Modules also feature caching, to both speed up I/O operations and to preserve data in the event of a controller failure, and battery backup, to prevent data loss in the event of a power failure. See the hardware books accompanying your RAID Modules for complete information on the redundant features of your hardware.

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

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, if available on the RAID Module you are using. 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.

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

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

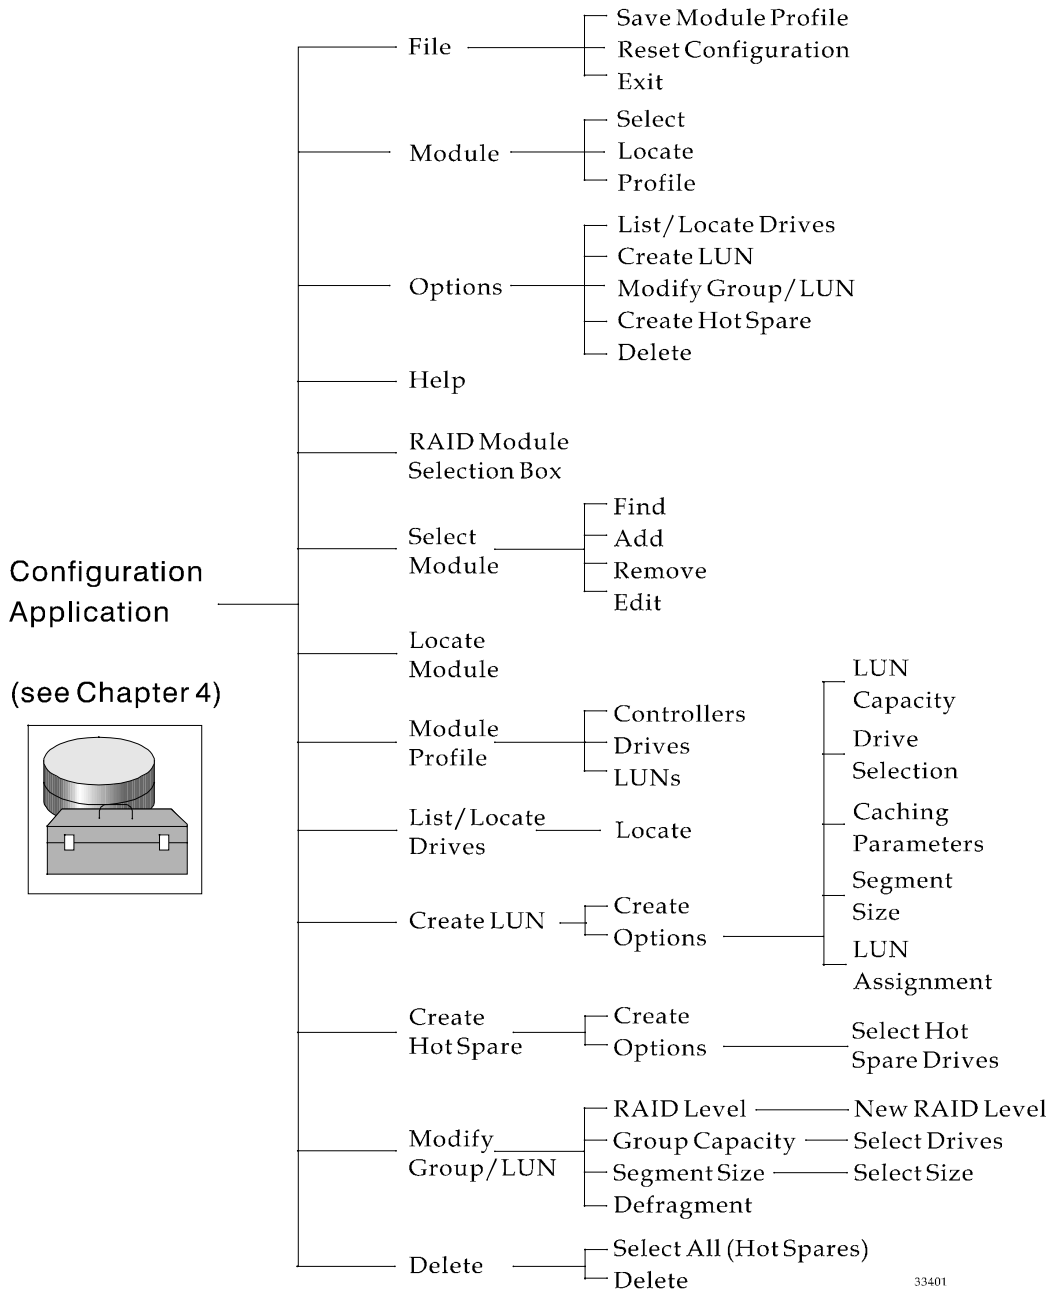
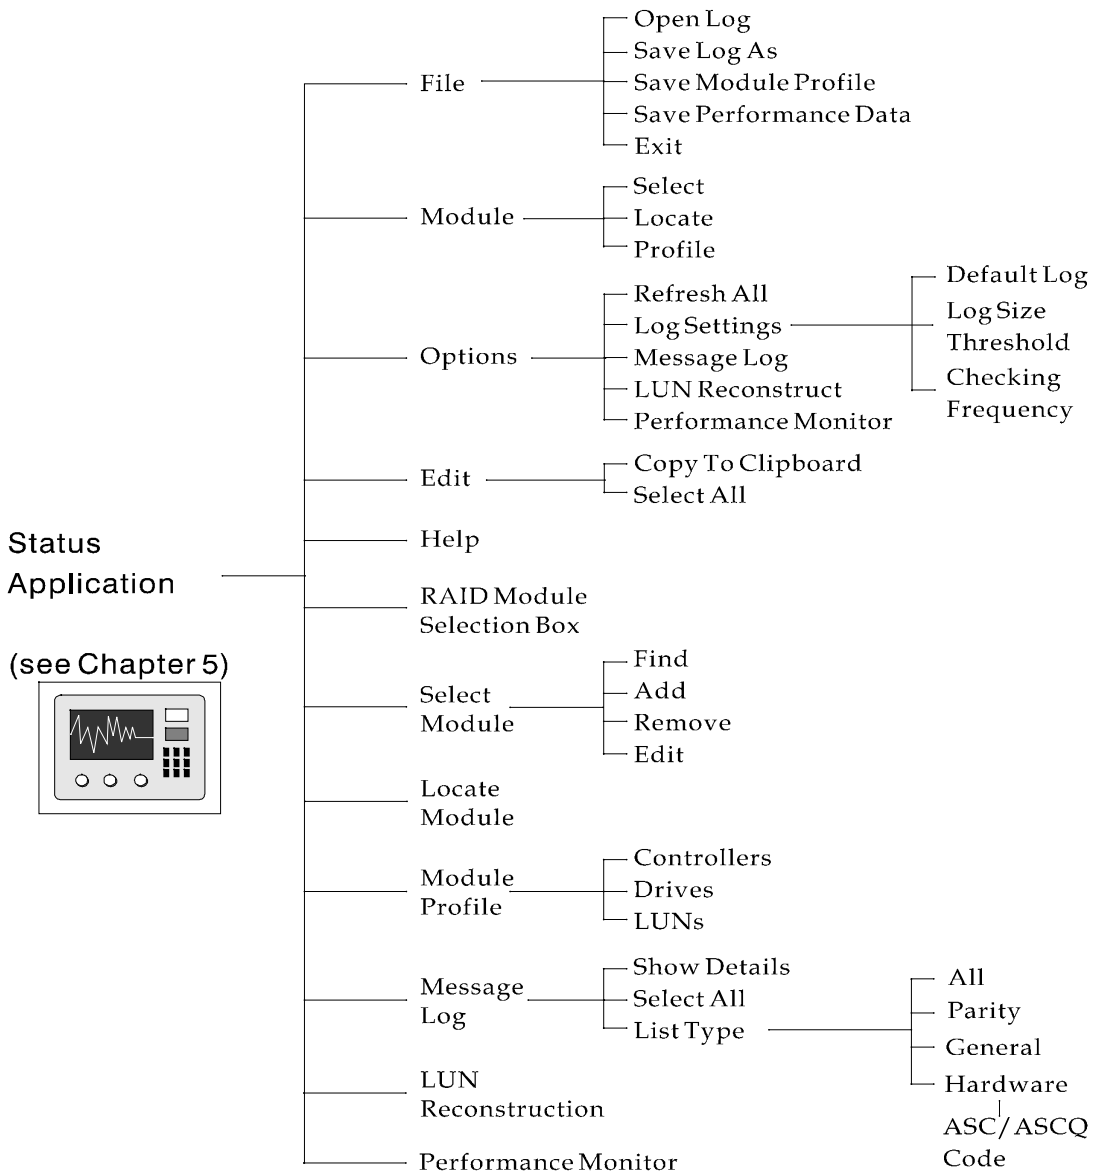
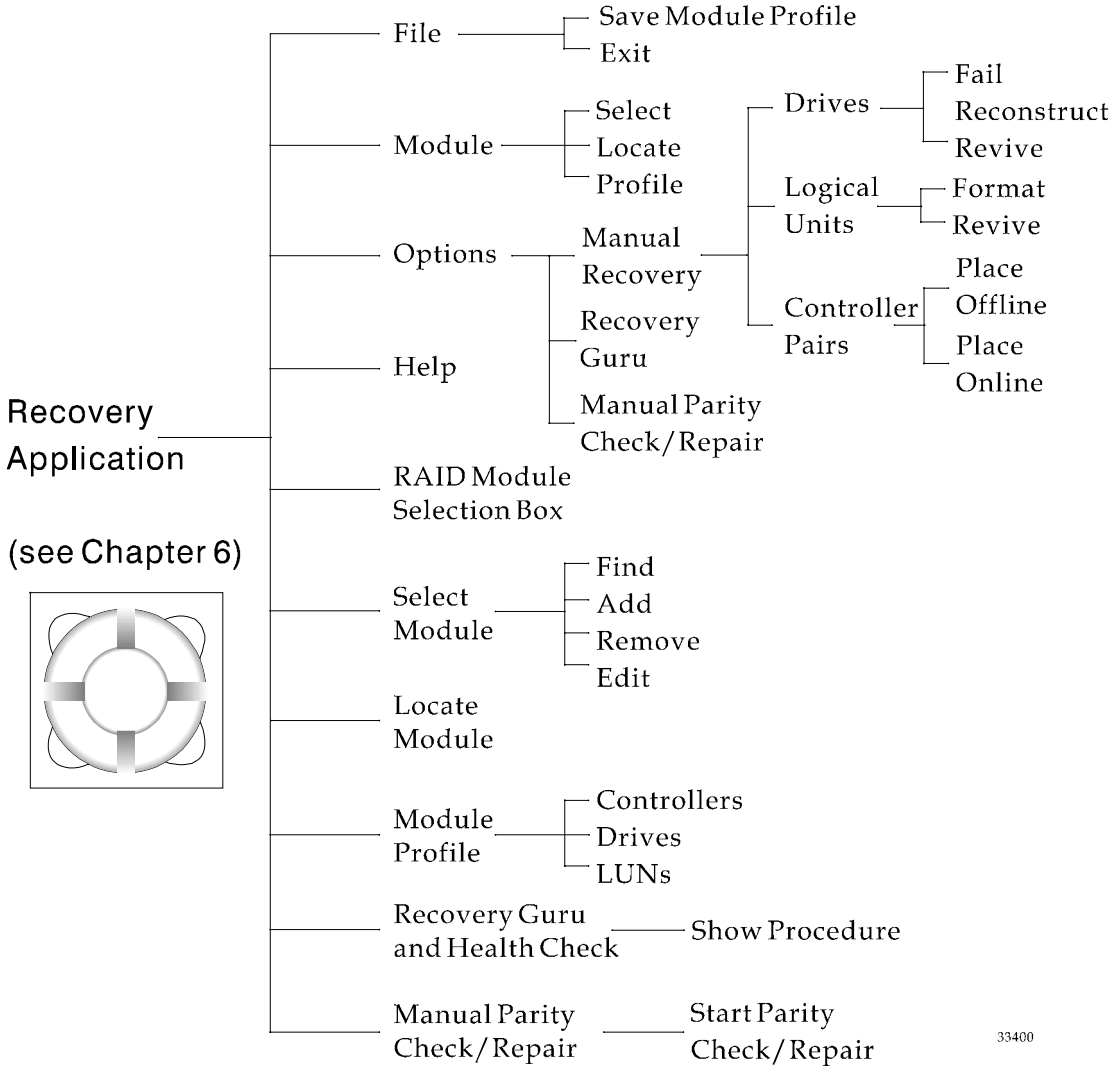


FIGURE 2-5 Configuration Application Summary



33403

FIGURE 2-6 Status Application Summary

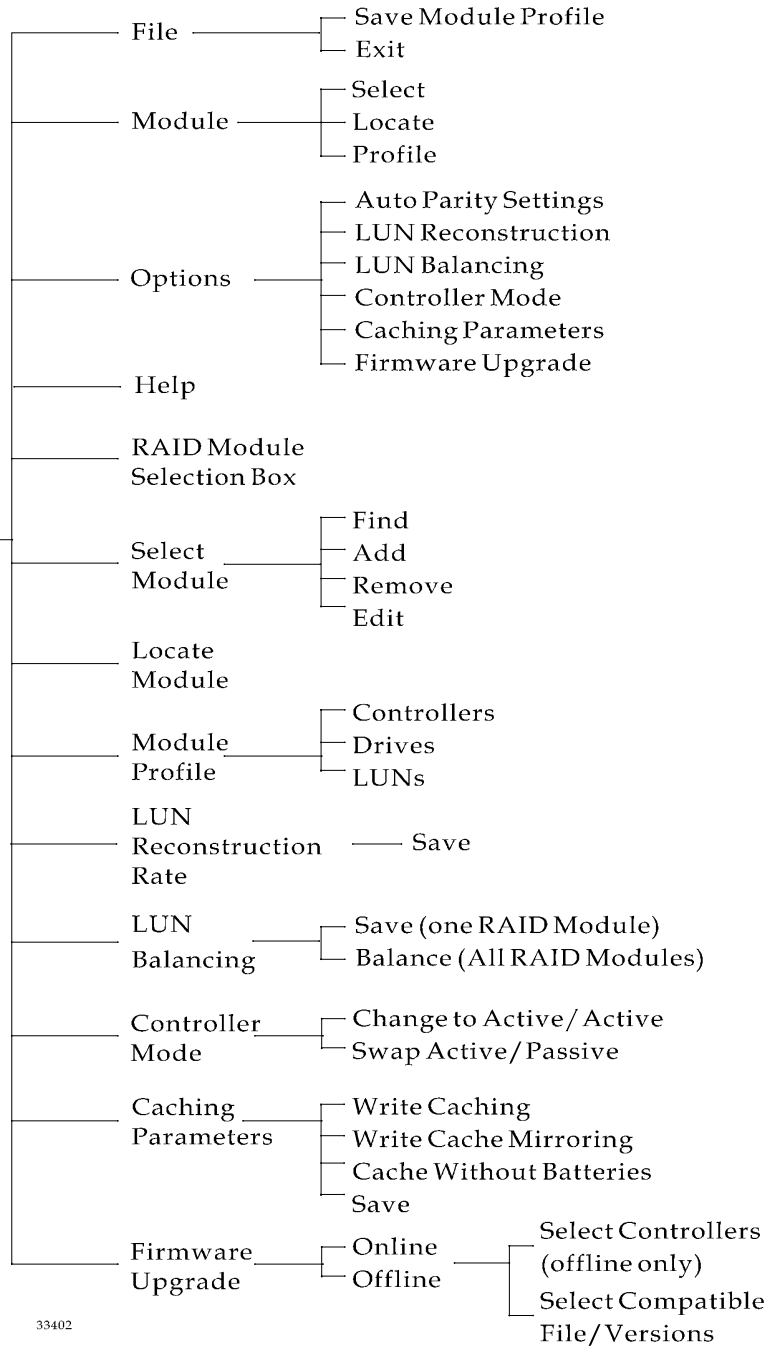



33400

FIGURE 2-7 Recovery Application Summary

Maintenance/Tuning Application

(see Chapter 7)



33402

FIGURE 2-8 Maintenance And Tuning Application Summary

Task List

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

TABLE 2-1 Storage Management Task List

Task	When to Perform	Reference
Informational Tasks (Locating Raid Modules, Checking Configurations)		
Locate a RAID Module	When you want to physically locate a RAID Module. Perform this task when you first install the module and mark the module for future reference.	“Locating a Module” on page 44
Locate A Logical Unit or Specific Drives	When you want to physically locate a drive or a logical unit.	“Locating Drives” on page 45
Assign Names to Modules and Controllers	To make the RAID Module components easily identifiable. Only the networked version allows you to name controllers.	“Selecting a Module” on page 41
View RAID Module Configuration Information	When you need information on your RAID Modules, such as current firmware version or drive information.	“Viewing A Module Profile” on page 47
Check the Status Log	When you want to see if any errors have occurred on the RAID Module.	“Viewing The Log” on page 94
Run the Performance Monitor	To gather information on the performance of the RAID Module, with the goal of discovering any changes that could be made to speed performance.	“Message Log Save Message” on page 102
Configuration Tasks (Creating and Modifying Logical Units)		
Create Logical Units (LUNs)	After you install a RAID Module.	“Creating Logical Units (LUNs)” on page 62
Create a Hot Spare	When you configure a RAID Module. You should always have at least one hot spare on your system to protect against failures.	“Creating Hot Spares” on page 85
Add Drives to an Existing Drive Group	When you want to add capacity in a Drive Group	“Modifying Drive Groups and LUNs” on page 73

TABLE 2-1 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 Module after you have created a LUN.	“Modifying Drive Groups and LUNs” on page 73
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 73
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 83
Delete a Hot Spare	If you need the hot spare drive for a new LUN.	“Deleting A Hot Spare” on page 88
Maintenance Tasks (Loading new firmware, changing parameters)		
Download Controller Firmware	<ul style="list-style-type: none"> • When you receive new firmware from Customer Services. • When adding RAID Modules with new firmware into an existing site (you will need to upgrade the existing controllers). 	“Upgrading Controller Firmware” on page 133
Download NVSRAM	<ul style="list-style-type: none"> • When you receive new firmware from Customer Services. • When adding RAID Modules with new firmware into an existing site (you will need to upgrade the existing controllers). 	“Upgrading Controller Firmware” on page 133
Reset the Configuration	If all attempts to use the RAID Module have failed. This is a last resort option that will delete all data and configuration information from your RAID Module.	“Resetting the Configuration” on page 89
Change Controller Modes	If the controllers in a RAID Module are currently active/passive and you want to change them to active/active.	“Changing Controller Mode” on page 153
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 149
Change Cache Parameters	To turn cache on and off after you have created a logical unit.	“Changing Cache Parameters” on page 155

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

Task	When to Perform	Reference
Change the Reconstruction Rate	To change the speed at which data is reconstructed after a drive replacement.	“Changing The LUN Reconstruction Rate” on page 147
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 143
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 94
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 118
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
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 Parity Check” on page 142

What's New In 6.2?

RAID Manager 6.2 contains the following new features:

- **Completely different user interface** – The user interface is considerably changed from previous versions of the storage management/RAID Manager software.
- **Performance Monitor** – A software option is provided to track LUN and controller usage.
- **Dynamic RAID changes** – You can now change the capacity of a Drive Group and the RAID level or segment size of an existing LUN without deleting LUNs (with controller firmware 3.0 or higher).
- **Changed Recovery Functions** – The Recovery Guru has changed. The function is now tied with the Health Check function and provides recovery information without providing the step by step procedures found in earlier versions of the software. In addition, Recovery Guru does not send any commands to the controller or initiate any action automatically. All Recovery operations are performed by the user.

Common Application Features

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

- Navigation—page 36
- Starting and Exiting An Application—page 37
- Using Online Help—page 38
- Selecting a Module—page 41
- Locating a Module—page 44
- Locating Drives—page 45
- Viewing A Module Profile—page 47
- Saving A Module Profile—page 56

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 3-1 describes how to use the mouse and keyboard.

TABLE 3-1 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:<ol style="list-style-type: none">1. Single-click to highlight a single item.2. 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.3. 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

Starting An Application

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

- The first screen you see will prompt you to select the RAID Module you want to access (see “Selecting a Module” on page 41).
- If the message log file is currently over its maximum size limit, you will be prompted to rewrite the file (see “Clearing The Message File” on page 101).




FIGURE 3-1 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's Guide*. If you have questions concerning a specific procedure, check the online help before coming back to this manual.

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

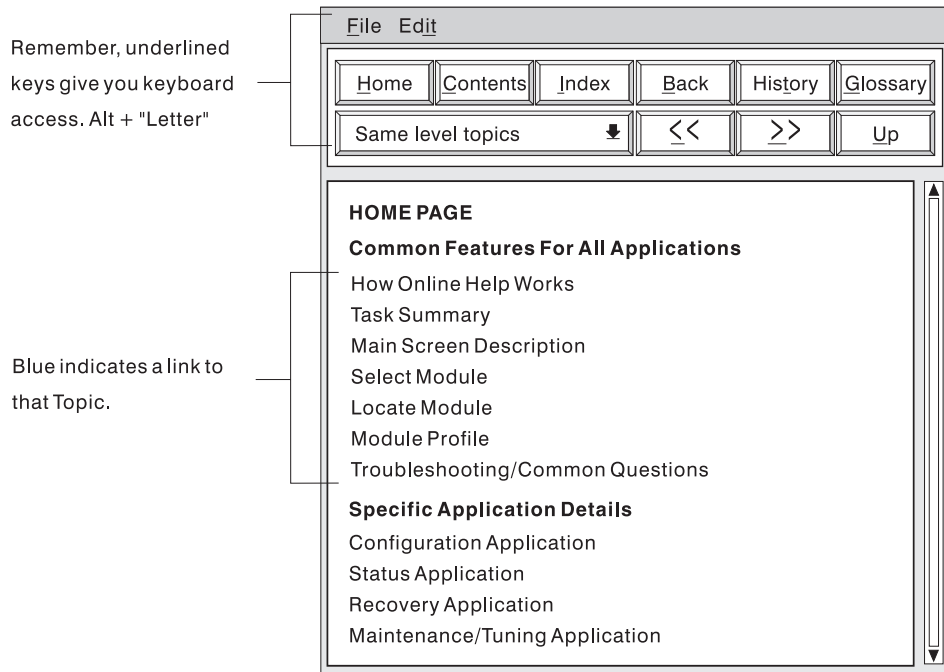


FIGURE 3-2 Main Online Help Screen

TABLE 3-2 Online Help Screen Elements

Selection	Description
File	<p>Allows you to do the following:</p> <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	<p>Copies text to a clipboard.</p> <p>From the top menu, select Edit → Copy to Clipboard to copy the topic in the window you are viewing.</p>
Home	<p>Returns you to the Home Page.</p> <p>This screen appears whenever you select Help from the top menu in an application.</p>
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>

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

Selection	Description
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	Moves you to the next higher level of topics.

Selecting a Module

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

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

RAID 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.

There are several ways to select a RAID 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 Module when you first start an application (see FIGURE 3-3). Either highlight a module and select OK or double-click on the module you want.
- Once in an application, you can:
 - Select the RAID Module you want from the RAID Module drop-down list box next to the tool bar.
 - Click on the Module Select button.
 - Select Module → Select from the drop-down menus.
 - Right click on any item in the Module Information tree (in the Configuration application only).

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

RAID Module Selection Screen

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

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

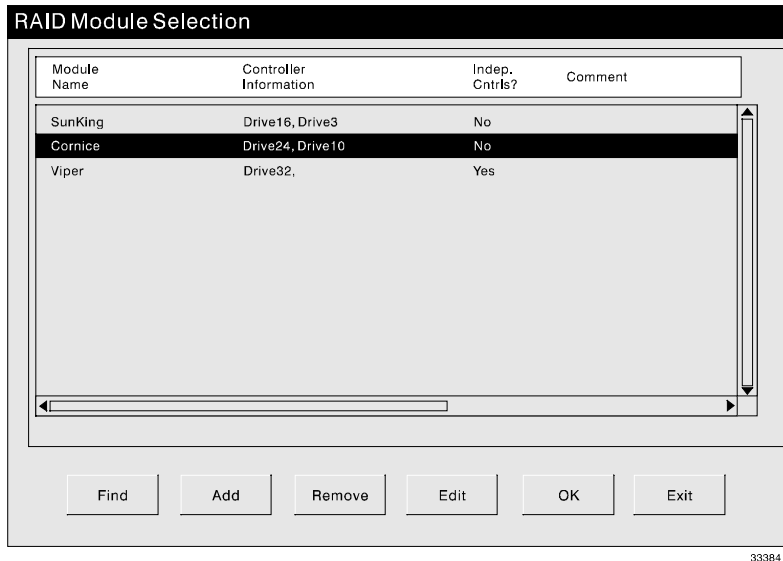


FIGURE 3-3 RAID Module Selection Screen

Note – The display gives the module name, the names of the controllers in the module, indicates whether the controllers have an independent controller configuration (see “Independent Controller Configuration” on page 21), and displays any comments you may have entered with the Edit option.

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

TABLE 3-3 Select Module Screen Elements

Screen Element	Description
Find	<p>Allows you to quickly locate a RAID Module. 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 Support Guide</i> specific to your operating system for details on adding new modules to your system.</p>
Remove	<p>Allows you to remove RAID Modules from your system.</p> <p>You can only remove one module at a time. This option will be grayed out if All RAID Modules 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 Modules 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 Module and to enter comments about it. Use the comments area to provide detailed information about the RAID Module (such as location information, independent controllers, etc.) to help you identify it.</p> <p>On Networked versions, you can also assign names to the RAID Module controllers.</p> <p>You can only edit information for one module at a time. This option will be grayed out if All RAID Modules is selected.</p>

Locating a Module

The storage management software allows you to physically locate a RAID 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 45).

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

To locate a RAID Module, do the following.

1. **Start the application you want to perform.**
2. **Select the RAID Module you want to locate. Do not select All RAID 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.**

End of Procedure

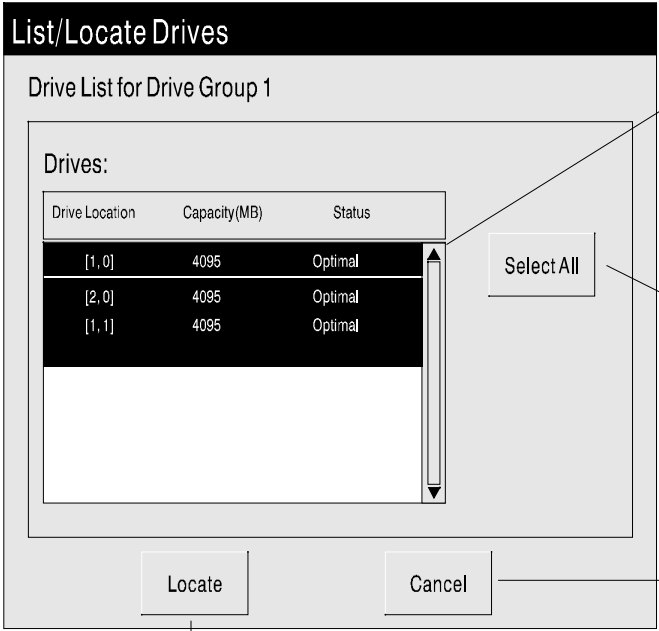
Locating Drives

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 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 3-4 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.**

End of Procedure



This display shows all the drives in the selected drive group, along with individual drive capacity and status. The first number given in Drive Location is the channel number of the drive; the second number is the SCSI ID. Highlight the drives you want to locate.

Click here to locate all the drives listed. This locates the entire drive group.

Click here to exit this screen without locating any drives.

Click here to locate the selected drives.

FIGURE 3-4 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 Modules.

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 RAID Modules (the Module Profile option is grayed out if All RAID 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 3-5).**
- 4. After you have viewed all the profile information you want, click OK on the RAID Module Profile screen to return to the application.**
- 5. After exiting Module Profile, you can save the profile information to a file. See “Saving A Module Profile” on page 56.**

End of Procedure

See TABLE 3-4 for an explanation of the items in this display.

Module Profile

Profile for SunKing

Summary Information:

Controllers: Name	Serial Number	Mode	Number of LUNs
A (Drive16)	1T75053813	Active	4
B (Drive1)	1T75053857	Active	4

Disk Drives:

Number of Drives = 20

Detailed Information:

- Controllers
- Drives
- LUNs

Close

33380

Click here to view detailed information on the module controllers (see "The Controller Profile Screen" on page 50).

Click here to view detailed information on the module drives (see FIGURE 3-7).

Click here to view detailed information on the module logical units (see FIGURE 3-8).

Click here to exit this screen

FIGURE 3-5 Module Profile Screen


TABLE 3-4 Module Profile Screen Contents

Screen Element	Explanation
Name	Identifies the controllers in the selected RAID Module as A or B and, where applicable, includes a system device name. The A and B correspond to the hardware controller IDs.
Serial Number	Identifies the controller by a number assigned by the manufacturer.
Mode	Identifies the operating state of the controller. Possible modes are Active, Passive, or Offline. You could also see Inaccessible with these statuses if the RAID Module has an independent controller configuration. Important If you do not see Mode information or other information in this screen 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 Module.

The Controller Profile Screen

SCSI Controller Profile

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



33361

FIGURE 3-6 Controller Profile Screen (SCSI)

TABLE 3-5 Controller Profile Screen Contents (SCSI)

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

The Drive Profile Screen

FIGURE 3-7 shows the Drive Profile screen. See TABLE 3-6 for an explanation of the items in this display.

Module Profile

Detailed Drive Information for SunKing

Drives:

Location	Capacity (MB)	Status	Vendor	Product ID	Firmware Version	Serial Number	Date Code
[1,0]	4095	Optimal	SEAGATE	ST34501W	8604	LG455691	303130
[2,0]	4095	Optimal	SEAGATE	ST34501W	8604	LG454814	303130
[1,1]	4095	Optimal	SEAGATE	ST34501W	8604	LG451206	313233
[2,1]	4095	Optimal	SEAGATE	ST34501W	8604	LG455289	303130
[1,2]	4095	Optimal	SEAGATE	ST34501W	8604	LG454930	303130
[2,2]	4095	Optimal	SEAGATE	ST34501W	8604	LG455061	303130
[1,3]	8682	Optimal	SEAGATE	ST19171W	8602	LAK39380	303932
[2,3]	0	Failed					
[1,4]	8682	Optimal	SEAGATE	ST19171W	8602	LAJ57877	303931
[2,4]	8682	Spare[2,3]	SEAGATE	ST19171W	8602	LAJ99452	303931

Close

33383

FIGURE 3-7 Drive Profile Screen

TABLE 3-6 Drive Profile Screen Contents

Screen Element	Explanation
Location	Designation indicating the unique location of the drive in the selected RAID 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.

The LUN Profile Screen

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

Module Profile

Detailed LUN Information for SunKing

Logical Units (LUNs):

LUN	Controller	Capacity (MB)	RAID Level	Segment Size (Blocks)	Write Cache	Cache Mirroring	Cache Without Batteries	Status
0	Drive16	100	5	32	Enabled	Enabled	Disabled	Optimal
1	Drive1	150	3	128	Enabled	Enabled	Disabled	Optimal
2	Drive18	8642	1	128	Enabled	Enabled	Disabled	Optimal
3	Drive3	125	5	32	Enabled	Enabled	Disabled	Optimal
4	Drive20	75	3	128	Enabled	Enabled	Disabled	Optimal
5	Drive21	8642	1	128	Enabled	Enabled	Disabled	Optimal
6	Drive6	500	5	32	Enabled	Enabled	Disabled	Optimal
7	Drive7	150	5	32	Enabled	Enabled	Disabled	Optimal

Close

33382

FIGURE 3-8 LUN Profile Screen

TABLE 3-7 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.

Saving A Module Profile

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 Module you want to save profile information for. Do not select All RAID Modules (this option is grayed out if All RAID Modules is selected).**
3. **Select File → Save Module Profile from the drop-down menus. A screen like FIGURE 3-9 is displayed.**




FIGURE 3-9 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.**

End Of Procedure

Configuring RAID Modules

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

- Configuration Overview—page 58
- The Module Information Display—page 60
- Creating Logical Units (LUNs)—page 62
- Modifying Drive Groups and LUNs—page 73
- Deleting LUNs—page 83
- Creating Hot Spares—page 85
- Deleting A Hot Spare—page 88
- Resetting the Configuration—page 89

Configuration Overview

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

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

1. **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 Module. This also creates a new drive group. When you create a LUN from unassigned drives, you assign the RAID level and number of drives. See “Creating A Logical Unit From Unassigned Drives” on page 62.
- Creating a new LUN from remaining capacity in a drive group. In this configuration mode, you can not 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 LUN From Remaining Capacity” on page 65.

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 RAID Level” on page 76.
- You can modify LUN parameters (such as segment size, caching, reconstruction rate and so on). See:
 - “Modifying Segment Size” on page 80
 - “Changing Cache Parameters” on page 155
 - “Changing The LUN Reconstruction Rate” on page 147

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 78.
- You can change the controller assignment of the LUN/drive group. See “Balancing LUNs” on page 149.

The Module Information Display


The Configuration application displays module information as shown in FIGURE 4-1.

This is the currently selected RAID Module.

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

This gives information on the drive groups and logical units in the RAID Module. 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 Module itself is selected. See TABLE 4-1 for a description of the elements in the display.



This 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 on the right of the screen.

FIGURE 4-1 Configuration Application Module Information Display

Click on the “+” or “-” signs to expand or collapse the tree.

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.

You can not click on items in the right-hand display.

TABLE 4-1 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	The name or ID of the controller, logical unit, or drive, depending on the display. <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 Module Selection Screen” on page 42).• Drive IDs are the channel number and SCSI ID of the drives. See “Drive Names” on page 13 for more information.
RAID Level	The RAID level of the logical unit.
Capacity (MB)	The capacity of the logical unit.
Status	The current status of the logical unit. If a status other than Optimal appears here, run Recovery Guru to determine what the problem is, and, if necessary, what action to take.

Creating Logical Units (LUNs)

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” on page 62).
- By creating a new LUN from capacity remaining in a drive group (see “Creating A LUN From Remaining Capacity” on page 65).

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 Module on which you want to create the new LUNs.**
- 3. Select the RAID Module or the Unassigned Drive group and click on the Create LUN button. A screen like FIGURE 4-2 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.**
- 6. Set the number of LUNs you want to create on the new drive group. This may be limited by the total number of LUNs you can create on your RAID Module. Each LUN created gets an equal amount of the 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 Module 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 149.

- 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 67 for a description of those screens.**

Tip – If your RAID Module 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.

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 start LUN creation with the current parameters.

Click here to view or change LUN options (capacity, drives, caching, segment size, and controller assignment). See FIGURE 4-4 through FIGURE 4-8 for examples of those options.

FIGURE 4-2 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.

Note – 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 Support Guide* for your operating system for more information.**

End of Procedure

Creating A LUN 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 Module 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 4-3 is displayed.**
4. **Set the number of LUNs you want to create from the unassigned capacity. This may be limited by the total number of LUNs you can create on the RAID Module. By default, each LUN created 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 67).**
6. **After you have set all the LUN parameters you want, start LUN creation by clicking Create in the Create LUN screen.**

Note – If you change any parameters in the Create LUN screen after making changes in the Options screens, any Options changes are 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 Support Guide* for your operating system for more information.

End of Procedure

This area displays the RAID level, number of drives, and available group capacity for the drive group. You can not 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.

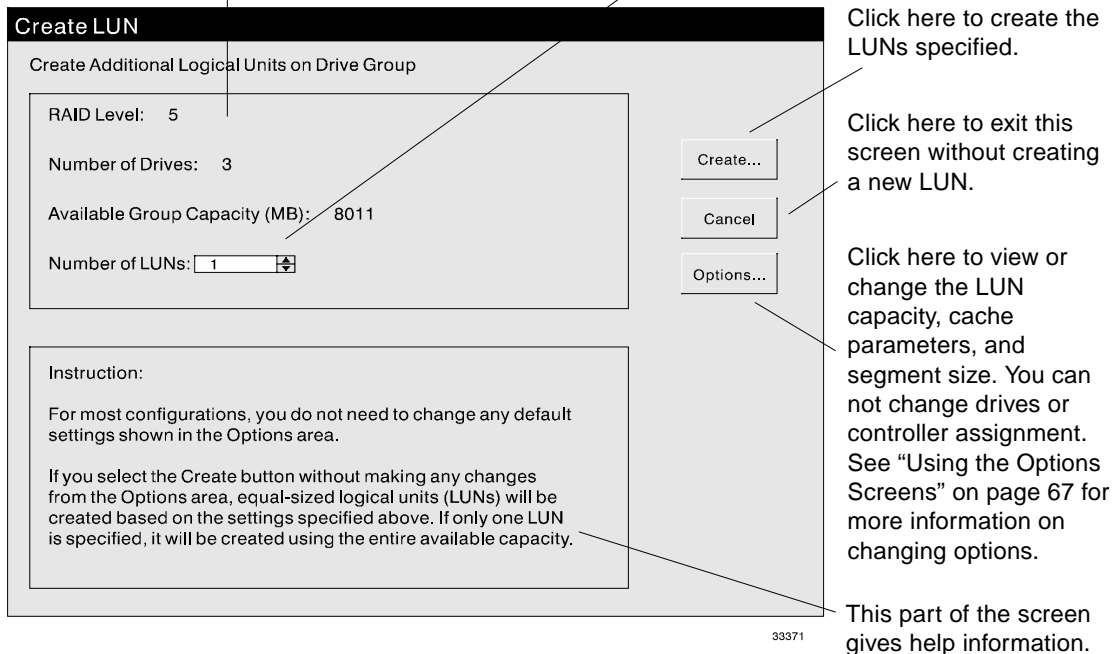


FIGURE 4-3 Create LUN Main Screen (Existing Drive Group)

You cannot change the RAID level or drives in the LUN.

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

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.

See FIGURE 4-5. See FIGURE 4-6. See FIGURE 4-7. See FIGURE 4-8.

Logical Unit	Capacity (MB)
2	8642
5	8642

Click here to set the capacity for the specified LUN or enter a value manually. The value in the Remaining Group Capacity box is updated to reflect the new setting.

Click here to return to the original settings.

This box displays the capacity remaining in the drive group.

Click here to go back to the Create LUN screen using the new values.

Click here to return to the Create LUN screen without changing the capacity. Values set in other options screens are retained.

FIGURE 4-4 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.

See FIGURE 4-4. See FIGURE 4-6. See FIGURE 4-7. See FIGURE 4-8.

This box displays the drives currently assigned and their capacities. If the RAID level is RAID 1, this box also shows the mirrored pairing. To move a drive from this box to the unassigned box, highlight the drive and click on the Move button.

Click here to return to the original values.

This box displays the drives available for assignment and their capacities. To assign a drive, highlight it and click the Move button.

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

Click here to return to the Create LUN screen assigning any drives. Values set in other options screens are retained.

33367

FIGURE 4-5 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 4-2).

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 9 and “Changing Cache Parameters” on page 155 for more information on the cache parameters.

See FIGURE 4-4. See FIGURE 4-5. See FIGURE 4-7. See FIGURE 4-8.

The cache parameters are listed here by LUN. Click the parameters you want to assign to the LUN. Note that setting or deleting certain values will set and delete other related values.

Click here to return to the original values.

Information on the cache without batteries parameter is displayed here.

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

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

FIGURE 4-6 Cache Parameters Screen

Note – See TABLE 7-3 for a description of cache parameter interdependencies.

Segment Size Option

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

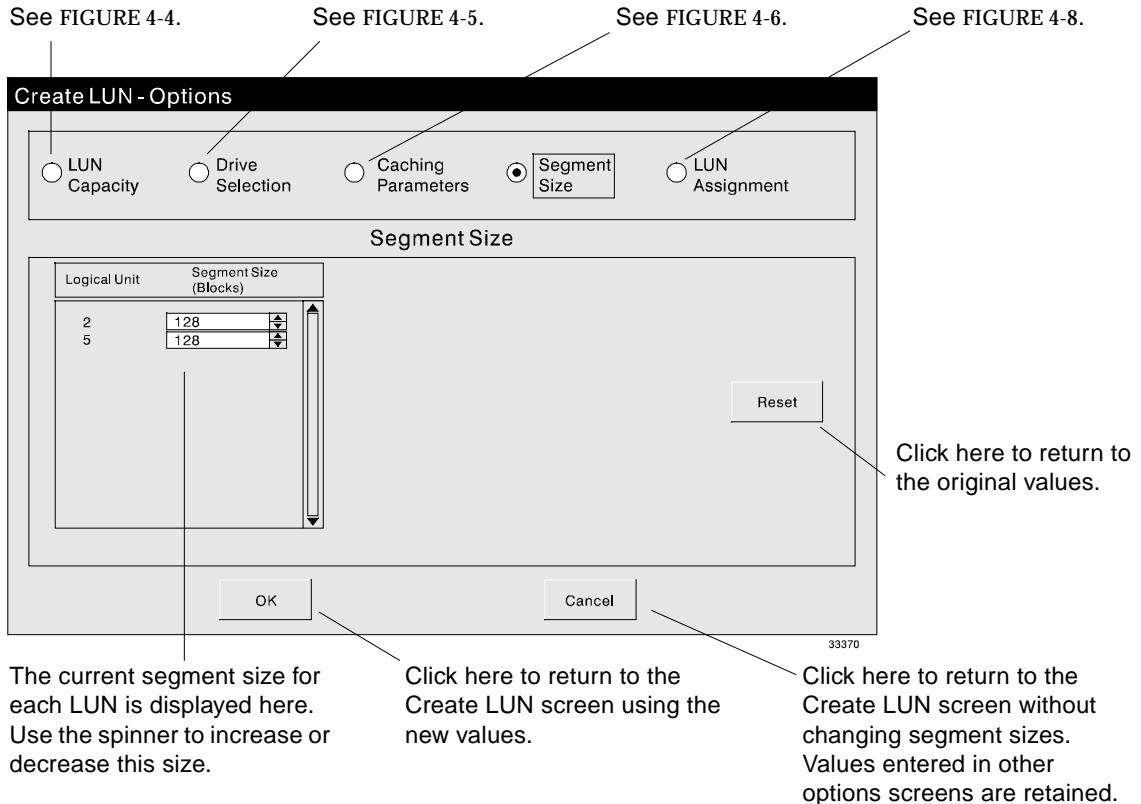


FIGURE 4-7 Segment Size Screen

You can set the segment size for individual LUNs; LUNs in a drive group do not need to have the same segment size value.

You can change the segment size of a LUN after you create it; see “Modifying Segment Size” on page 80.

LUN Assignment Option

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

See FIGURE 4-4. See FIGURE 4-5. See FIGURE 4-6. See FIGURE 4-7.

Create LUN - Options

LUN Capacity Drive Selection Caching Parameters Segment Size LUN Assignment

LUN Assignment

Current Drive Group/LUN Assignment:

Controller	Drive Groups Owned	Logical Units Owned	Total Group Capacity (MB)
Drive16	1, 3, 4	0, 1, 3, 6, 7	33506
Drive2	2	2, 5	17204

Assign New Group/LUN to Controller:

Controller A (Drive 16)
 Controller B (Drive 2)

33369

This display shows the current drive group and LUN assignment for the two controllers.

Click here to return to the original values.

This button shows the current controller assignment for the new drive group. Click the controller you want to own the new drive group.

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

Click here to return to the Create LUN screen without reassigning the drive group. Values entered in other options screens are retained.

FIGURE 4-8 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 can not change a 5-drive RAID 5 to RAID 1 because RAID 1 must have an even number of drives. See “Changing RAID Level” on page 76.
- 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 4-2). See “Adding Drives” on page 78.
- 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 80.
- Defragment the logical units on the drive group — this may increase total available capacity. See “The Defragment Option” on page 74 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 155. You can change other cache parameters using the `rdacutil` command line command. See the `rdacutil` command entry in Appendix A.
- Change the reconstruction rate — you can change the reconstruction rate assigned to individual logical units. See “Changing The LUN Reconstruction Rate” on page 147.

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

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

The Defragment Option

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

For example, FIGURE 4-9 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 4-9 Defragment Logical Units #1

Now assume LUN 2 (the middle LUN) is deleted. This is shown in FIGURE 4-10.

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 4-10 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 4-11.

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 4-11 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 4-11, the resulting drive group will look like FIGURE 4-12.

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 4-12 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” on page 82 for procedural information.

Changing RAID Level

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

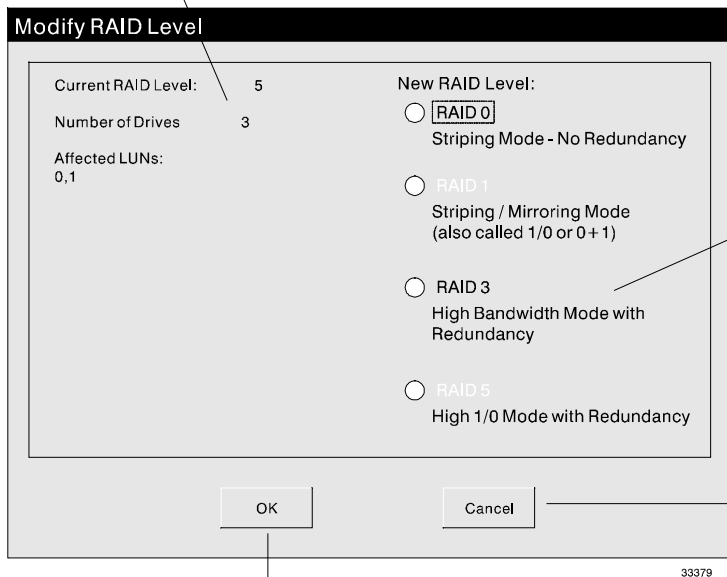
- 1. Start the Configuration application.**
- 2. Select the RAID Module containing the logical units you want to modify.**
- 3. Select the Drive Group you want to modify in the Module directory tree.**
- 4. Click on the Modify Group/LUNs button, then select Modify RAID Level, or select the option from the drop-down menus. A screen like FIGURE 4-13 is displayed.**
- 5. Select the RAID level you want for the drive group.**
Keep in mind that you may be limited by the number of drives in the group. For example, if the drive group has an odd number of drives, you will not be able to select RAID 1.
- 6. After you have selected the new RAID level, click OK to make the change.**

7. 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.

End of Procedure

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



Select the new RAID level here. If the RAID level is grayed out, you can not change the drive group to that RAID level. This is usually due to the number of drives in the current group.

Click here to exit without changing the RAID level.

Click here to change the RAID level.

FIGURE 4-13 Modify RAID Level Screen

Adding Drives

Use this procedure to increase the capacity of a drive group by adding drives to the group. TABLE 4-2 gives the formulae for calculating how much space is added to the drive group.

TABLE 4-2 Adding Capacity Formulas

RAID Level	The Drive Group Capacity Is <i>Increased</i> By:
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 x 18 or 36 GB. The drive group would then have a capacity of 72 + 36 or 108 GB.

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

Perform the following procedure:


1. **Start the Configuration application.**
2. **Select the RAID Module 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 4-14 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.**

7. 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.

End of Procedure

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 4-14 Modify Group Capacity Screen

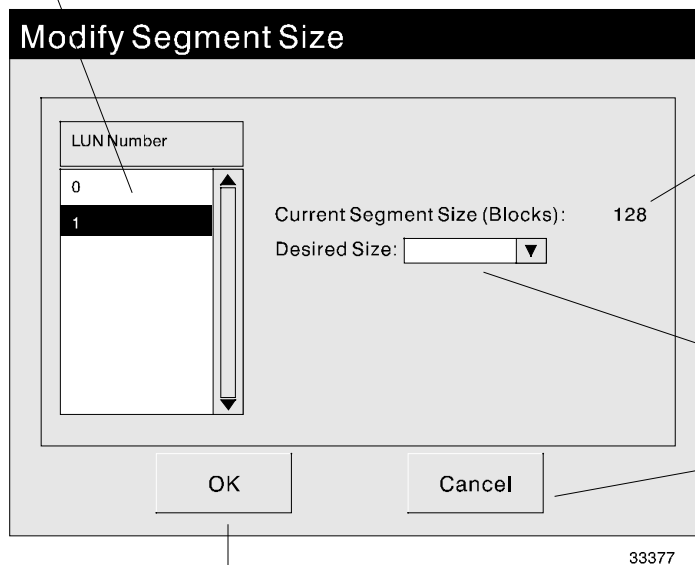
Modifying Segment Size

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 Module 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 4-15 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.

End of Procedure

Select the logical units you want to modify here.



When you select a logical unit, the current segment size setting is displayed here.

Select the new size here.

Click here to exit without changing the segment size.

Click here to change the segment size as specified.

FIGURE 4-15 Modify Segment Size Screen

Defragmenting The Drive Group

Defragmenting the drive group can free up space after the deletion of a logical unit from the drive group. See “The Defragment Option” on page 74 for more information.

To defragment the drive group, perform the following procedure.

- 1. Start the Configuration application.**
- 2. Select the RAID Module 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 LUNs button, then select Defragment, or select the option from the drop-down menus.**
- 5. A message is displayed, stating:**
 - That 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.
 - That 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.
- 6. 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 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.

End of Procedure

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 Module and exit configuration without creating at least 1 LUN. Doing so could prevent you from accessing the RAID Module.

Before you delete a LUN, check your *Installation and Support Guide* to see if there are any special restrictions or special requirements on LUN deletion. For example, you may be required to first delete partitions, 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 84.

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 Support Guide* for information.
- 2. Start the Configuration application.**
- 3. Select the RAID Module containing the LUN you want to delete.**
- 4. Select the drive group containing the LUN you want to delete in the Module directory tree.**
- 5. Do one of the following, and a screen similar to FIGURE 4-16 appears:**
 - Press the Delete key.
 - Select Options → Delete from the drop-down menu.
 - Right click on the drive group containing the LUN, then select Delete.
- 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 Support Guide* for your operating system for more information.

End of Procedure




FIGURE 4-16 Delete LUNs Screen

Deleting LUNs In The Network Configuration

If you are using the Networked version or a host-RAID Module configuration with more than one host (see "Module Configurations" on page 18), 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 Hot Spares

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

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

To create a hot spare drive, do the following.

- 1. Start the Configuration application.**
- 2. Select the RAID Module you want to create hot spares for.**
- 3. Highlight the Unassigned drive group or the Hot Spare drive group.**
- 4. Click on the Hot Spare button or select Options → Create Hot Spare from the drop-down menus. A screen like FIGURE 4-17 is displayed.**
- 5. Enter the number of hot spares you want to create.**

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

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

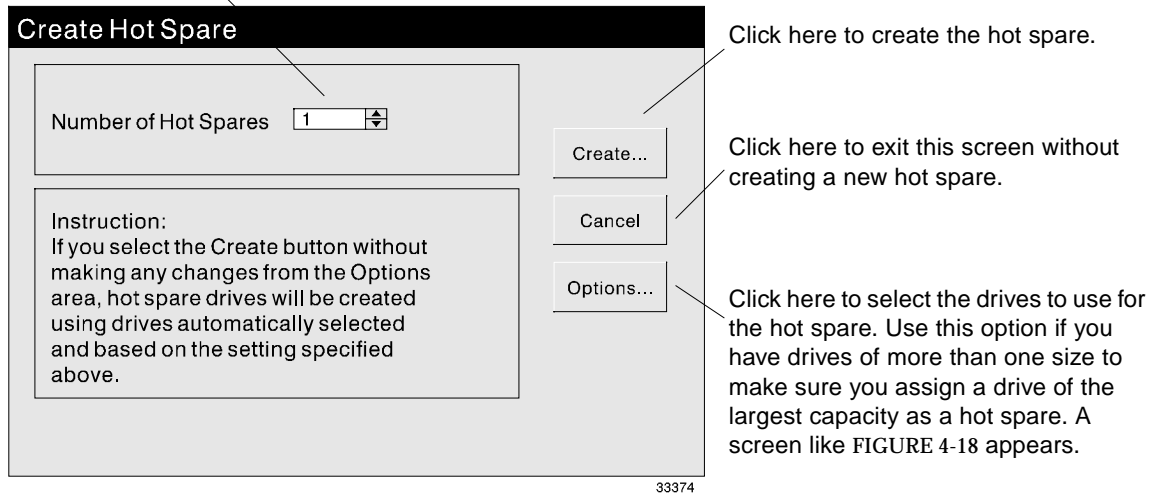


FIGURE 4-17 Create Hot Spare Screen


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.

End of Procedure

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.



This box shows the drives to be made into hot spares. Highlight a drive and click on the Move button to move the drive back to the unassigned group.

Use this button to move drives between the two boxes.

Click here to return to the Create Hot Spares screen.

Click here to exit this screen without creating any hot spares.

FIGURE 4-18 Select Hot Spare Screen

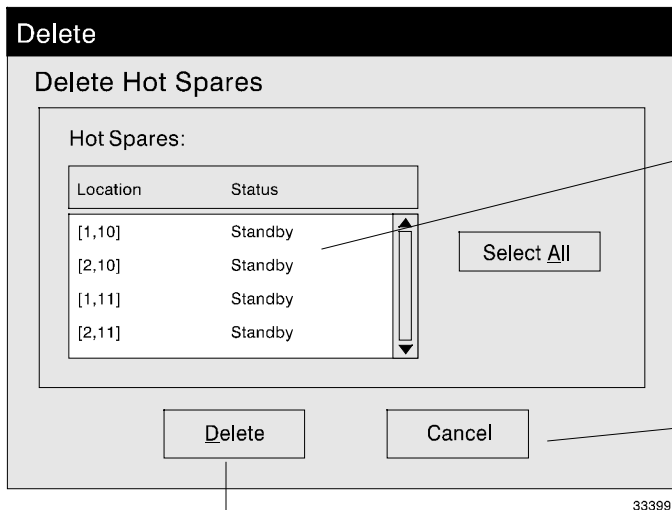
You can create as many hot spares on the RAID Module as there are drive channels. Always use the largest capacity drives for hot spares.

Deleting A Hot Spare

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 Module containing the hot spare you want to delete.**
3. **Select the Hot Spare drive group.**
4. **Press the Delete key or select Options → Delete from the drop-down menus. A screen like FIGURE 4-19 is displayed, listing all the drives currently assigned as hot spares.**
5. **Select the hot spares you want to delete, then click Delete.**
6. **The selected hot spares are deleted and returned to the Unassigned drive group.**

End of Procedure



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 4-19 Delete Hot Spares Screen

Resetting the Configuration



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

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

- Because deleting LUNs causes data loss, back up data on *all* the drive group/LUNs in the RAID Module. This operation also deletes any filesystems mounted on the LUNs.
- You must first stop I/Os to the affected RAID Module and you should ensure no other users are on the system.
- You cannot reset the configuration on a RAID Module with an independent controller configuration. For a description of this configuration, see “Independent Controller Configuration” on page 21.
- This operation could also fail if the storage management software cannot gain exclusive access to the drive groups/LUNs, such as if filesystems are mounted.

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

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 Module! Step 3 is your last chance to Cancel!

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

A default configuration appears in the main Configuration screen.

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

4. Select OK at the “Reset Was Successful” confirmation screen.
5. 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 Support Guide* specific to your operating system (restrictions and/or troubleshooting sections) and to the appropriate system documentation for specific details.

End of Procedure

Checking Status

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

- Status Overview—page 92
- Viewing The Log—page 94
- Changing Log File Settings—page 99
- Clearing The Message File—page 101
- Using The Performance Monitor—page 103

Status Overview

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

- Background Polling—page 92
- The Message Log File—page 92
- Message Types—page 93
- The Status Log And Health Check—page 93
- Event Notification—page 94

Background Polling

The storage management software polls all RAID Modules 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 99). 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.

The Message Log File

The message log is the file the application uses to store event messages. The default file is defined by your operating system (see the *Installation and Support Guide*), but can be changed to any file (see “Changing Log File Settings” on page 99).

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 99). 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 101 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 Module or message types are currently selected.

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

Message Types

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

TABLE 5-1 Message Types

Type	Description
Parity	Messages related to parity check/repair (either manual or automatic). The message 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.

The Status Log And Health Check

Note that in some cases, you may discover a problem when running Recovery Guru/Health Check that is not 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 Support Guide* for your operating system for details.

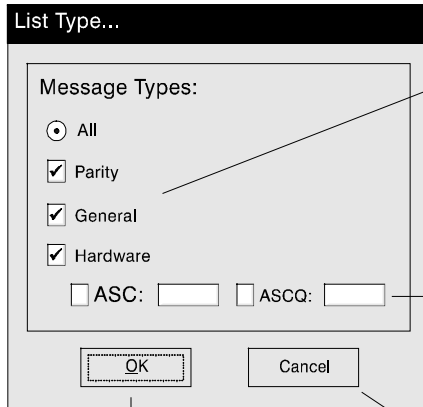
Viewing The Log

To view the status log, perform the following steps.

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

The contents of the currently selected message log is displayed (if you didn't select another message log file, this is the default message log). FIGURE 5-2 shows an example of the screen displayed.

4. **If you want to view a log other than the current default message log, go to File → Open Log and specify the log file you want to view.**
5. **If desired, click on List Type to filter the types of messages displayed (FIGURE 5-1).**
6. **Continue with Step 7.**



Select the types of message you want displayed by clicking in the appropriate box. After you have selected the message types you want, click OK. See TABLE 5-1 for a description of message types.


If you select Hardware messages, you can further filter the messages by entering the ASC/ASCQ codes you want to view.

Click here to display the selected message types.

Click here to exit this screen without changing the messages displayed.

FIGURE 5-1 Selecting Message Type

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



This display gives summary information on the messages in the currently selected message log. TABLE 5-2 explains the entries shown.

Highlight the messages you want to view, then click Show Details.

Click here for Help.

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

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

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

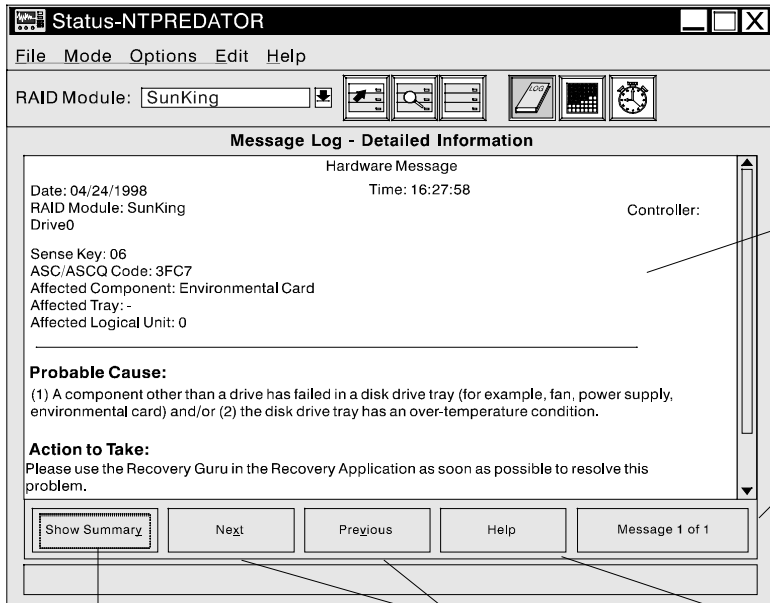
Click here to specify the types of messages you want to display (see FIGURE 5-1).

FIGURE 5-2 Message Log Screen

TABLE 5-2 Message Log Explanation

Column Heading	Explanation
Date & Time	Indicates when the detected event was logged. <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 Module event; either Hardware, Parity, or General (see “Message Types” on page 93).
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.

7. Highlight a displayed message or messages and click Show Details to get more information on the message. A screen like FIGURE 5-3 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 5-2).

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

Click here for help.

FIGURE 5-3 Message Details Screen

8. View the details and take the appropriate action.

9. You can also perform the following actions:

- Save the message log file to another file — select File Save → Log As from the drop-down menus. See “The Message Log File” on page 92.
- 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 (FIGURE 5-3).
- Refresh the screen — select Options → Refresh All from the drop-down menus. This updates the message log with any events that occurred since you opened the log. You can only perform this option from the Summary screen (FIGURE 5-2).

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

End of Procedure

Changing Log File Settings

You can change the following log settings (see “The Message Log File” on page 92):

- **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 101). Note that setting a size does not limit the size of the file. The file will 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 Modules. 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 Module (this change affects all RAID Modules).**
3. **Select Options → Log Settings from the drop-down menus. A screen like FIGURE 5-4 is displayed.**
4. **Enter the values you want, then click Save (or view the current values and click cancel to exit without changing values).**

End of Procedure

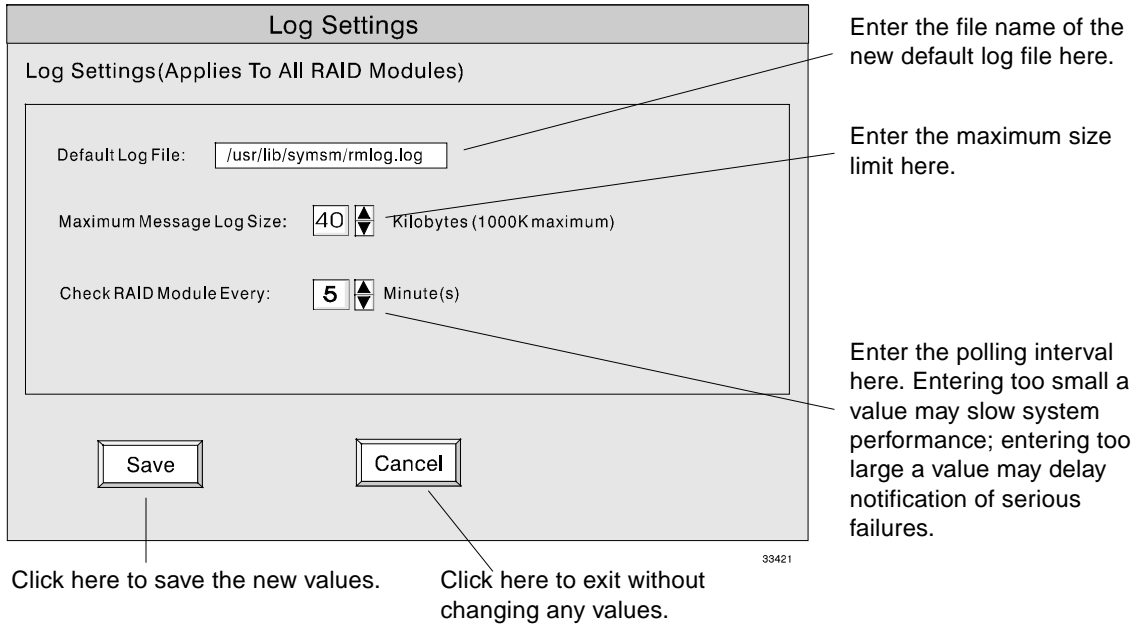


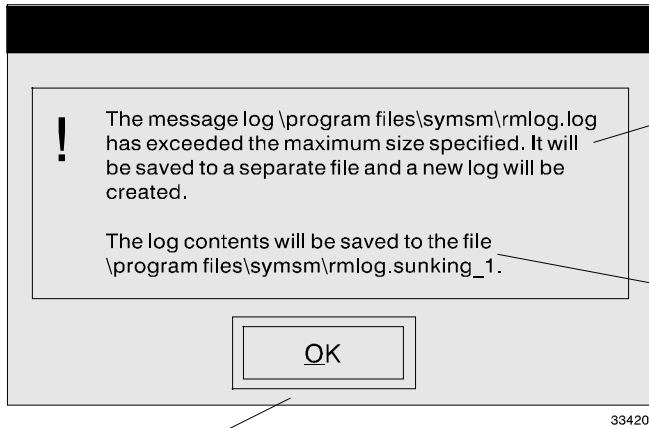
FIGURE 5-4 Log Settings Screen

Clearing The Message File

If you open a storage management application and the message log file currently exceeds the defined maximum size limit (see “Changing Log File Settings” on page 99), the message log file is saved and cleared. The following sequence occurs:

- 1. When you open the storage management application, a warning message like FIGURE 5-5 is displayed.**
- 2. Click OK.**
 - If the 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 5-6 is displayed.
- 3. Click Yes if you want to write over the current `rmlog.<hostname>_1` file. Click No if you want to create the next file in the sequence (`_2`, `_3`, `_4`, and so on).**
- 4. The first message in the new log file will notify you that this backup occurred. Highlight the message and click Show Details to find the name of the backup file used.**

End of Procedure

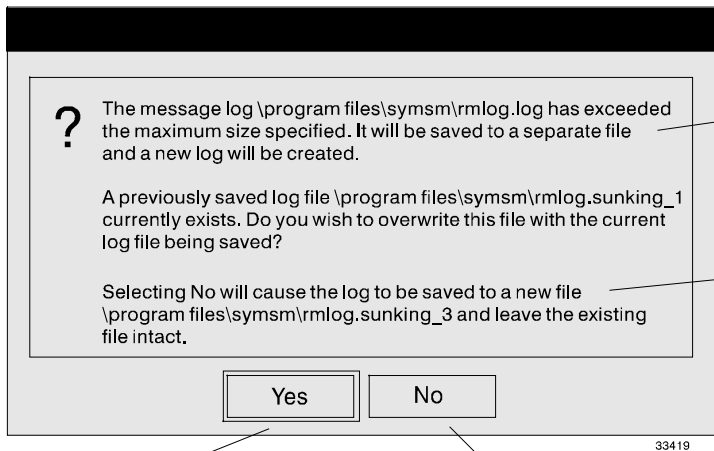


This screen indicates that the default message file has reached or exceeded its defined maximum size limit.

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

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

FIGURE 5-5 Message Limit Warning Message



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

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

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

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

FIGURE 5-6 Message Log Save Message

Using The Performance Monitor

Note – Before you start the Performance Monitor, make sure you have I/O going to the RAID Modules you want to monitor.

To run the Performance Monitor, perform the following steps:

1. Start the Status application.
2. Select the RAID Module you want to monitor. (Do not select All RAID Modules.)
3. Click on the Performance Monitor button or select Options → Performance Monitor from the drop-down menus. A screen like FIGURE 5-7 is displayed.




FIGURE 5-7 Performance Monitor Settings Screen

4. Select the Controllers/LUNs you want to monitor and click OK. A screen like FIGURE 5-8 is displayed.

5. Click Start to start the Performance Monitor. A screen like FIGURE 5-9 is displayed.

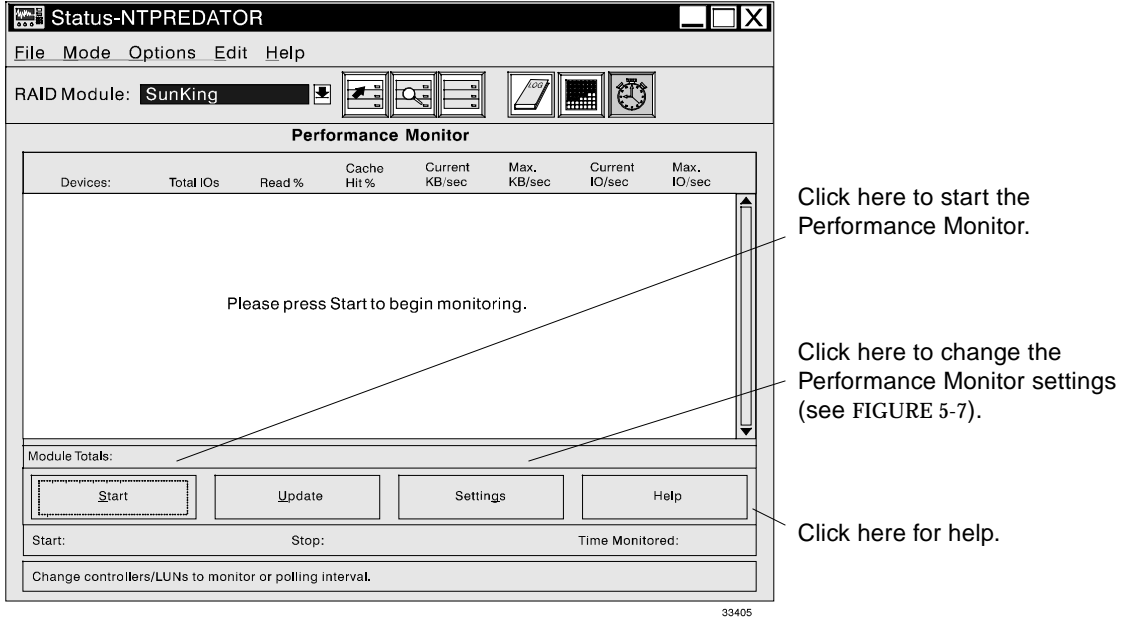



FIGURE 5-8 Performance Monitor Start-Up Screen



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

This bar gives the totals for the selected controllers.

Click here for help.

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

Click here to stop the Performance Monitor.

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

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

FIGURE 5-9 Performance Monitor Data Screen

TABLE 5-3 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 IOs	The number of total I/Os performed by this device since the “Start” button of the performance monitor was last clicked. When this number exceeds 99,999, it 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.

TABLE 5-3 Performance Monitor Screen Contents (*Continued*)

Column	Explanation
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.**

End of Procedure

Using Performance Monitor Data

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

Total I/Os — This field is useful for monitoring the I/O traffic to a specific controller and a specific LUN. This is needed to identify possible I/O “hot spots.” It is also useful to be able to identify actual I/O patterns to the individual LUNs and compare those with the expectations based on the application. If a particular LUN has considerably more I/O activity, consider moving the LUN to the other controller in the RAID Module. If a particular array controller has considerably more I/O activity, consider moving 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.

Recovery

This chapter describes how to recover from component failures.

- Recovery Overview—page 110
- Running Recovery Guru/Health Check—page 118
- Manual Recovery Options—page 121
- Monitoring/Changing Reconstruction—page 128

Recovery Overview

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

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

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

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

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

- General Recovery Steps—page 111
- Recovering From Drive Failures—page 112
- Recovering From Controller/Connection Failures—page 116
- Recovering From Fan/Power Supply Failures—page 117

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 Module
 - SNMP notification
 - Console messages from your operating system
 - Results of a Recovery Guru/Health Check
 - A non-optimal status appearing for any RAID Module component in the Configuration screen or in Module Profile
2. **You run Recovery Guru/Health Check to get information on the failure (see “Running Recovery Guru/Health Check” on page 118).**
3. **You should view the recommended procedure in Recovery Guru/Health Check for detailed information on the steps you need to take to recover from the failure (see FIGURE 6-2).**
4. **You take the steps indicated, which often involve replacing the failed component.**

In some cases, you may need to refer to your hardware 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 121, but do not do so unless instructed by Recovery Guru.

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

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

End of General Procedure Description

Recovering From Drive Failures



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

- Degraded And Dead Modes—page 112
- Recovery And RAID Level—page 113
- Multiple Drive Failures And Channel Failures—page 114
- Reviving Drives—page 114
- Monitoring Reconstruction Progress—page 115
- Using Configured Drives As Spares—page 115

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 6-1 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 Module to recover from the failure. TABLE 6-1 gives general information only; *always* run Recovery Guru/Health Check before attempting to replace drives.

TABLE 6-1 Drive Failures And Recovery (Summary)

RAID Level	Number of Failed Drives	Explanation
0	1 or more	A single drive failure in a RAID 0 logical unit results in the loss of all data on the logical unit, because RAID 0 does not provide data redundancy. To recover from a drive failure in a RAID 0 logical unit, you need to replace the failed drive, reformat the logical unit, and restore data from backup.
1, 3, or 5	1	A single drive failure in a RAID 1, 3, or 5 logical unit does not result in the loss of data. The logical unit 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	<p>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.</p> <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.

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 Module has 1 hot spare drive, *three* drives must fail before data is lost in a RAID 5 logical unit. See “Hot Spare” on page 13 and “Protecting Yourself From Component Failures” on page 25 for more information.

Multiple Drive Failures And Channel Failures

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

To prevent channel failures from causing multiple drive failures, 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 121). As always, however, follow the steps given in the Recovery Guru procedure screen before attempting any kind of manual recovery.

Reviving Drives

Under limited circumstances, you can “revive” a failed drive and return it to an optimal status using the Revive Drive function in Manual Recovery. In general, you can only revive a drive if you accidentally removed it or failed it, not if the controller failed it. See “Manual Recovery/Drives” on page 122 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 16 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 128 for more information.

Using Configured Drives As Spares

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

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

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

If this is not possible, you can 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 114 *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.

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 Module becoming unusable unless *all* of the following conditions are met:

- The RAID Module has redundant controllers (see “Redundant Controllers” on page 16).
- The RAID Module is configured in the single host configuration using two host busses (see “Module Configurations” on page 18).
- The RAID Module has RDAC protection (see “RDAC Driver” on page 15).

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

TABLE 6-2 Recovering From Failed Controllers (Summary)

Fail-Over Protection?	Recovery Procedure
Yes	<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 Module. Follow the instructions on the Recovery Guru procedure screen. In general, you must manually fail the failed controller, remove it, and insert a new one.</p>
No	<p>If the RAID Module does not have fail-over protection, data will not flow along the failed connection, and any data currently in the cache of the failed controller is lost. Data on the drives is not necessarily destroyed, but is inaccessible.</p> <p>Recovery: You will need to shut down the RAID Module to replace the controller. Follow the instructions on the procedure screen.</p>

Recovering From Fan/Power Supply Failures



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

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

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.


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 Modules you want to check, or select All RAID Modules to check all the modules on your system. A screen like FIGURE 6-1 is displayed.**

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.

33987

FIGURE 6-1 Recovery Main Screen

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

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



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

Click here for help.

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

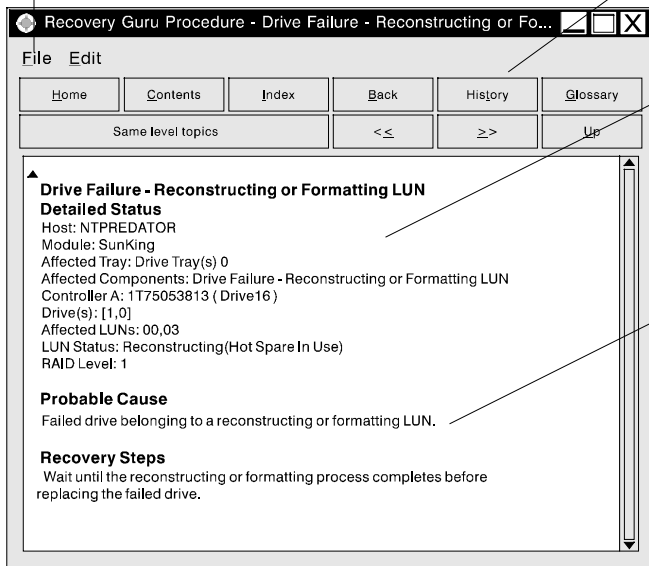
FIGURE 6-2 Recovery Screen/Exceptions Found

4. If exceptions are found, highlight the exception and click on Show Procedures for recommendations on how to fix the problem (see FIGURE 6-3).
5. Print out the procedure, if desired, using File → Print Topic.
6. Follow the procedure given to recover from the component failure. You may need to use the procedures given in “Manual Recovery Options” on page 121.

End of Procedure

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 6-3 Recovery Procedure Screen

Manual Recovery Options



Caution – Take care when using any of these options. In general, do not use these options unless specifically directed by the procedure screen of Recovery Guru or by 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 114).
- You want to force a drive failure rather than wait for the controller to fail it.
- You want to remove a controller from the RAID Module.
- A channel failure caused multiple drives in a logical unit to fail.

There are three Manual recovery options:

- Manual recovery of drives — See “Manual Recovery/Drives” on page 122.
- Manual recovery of logical units — See “Manual Recovery/Logical Units” on page 124.
- Manual recovery of controllers — See “Manual Recovery/Controllers” on page 126.

Manual Recovery/Drives



Caution – Take care when using manual recovery. In general, do not use these options unless specifically directed by the procedure screen of Recovery Guru or by 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:

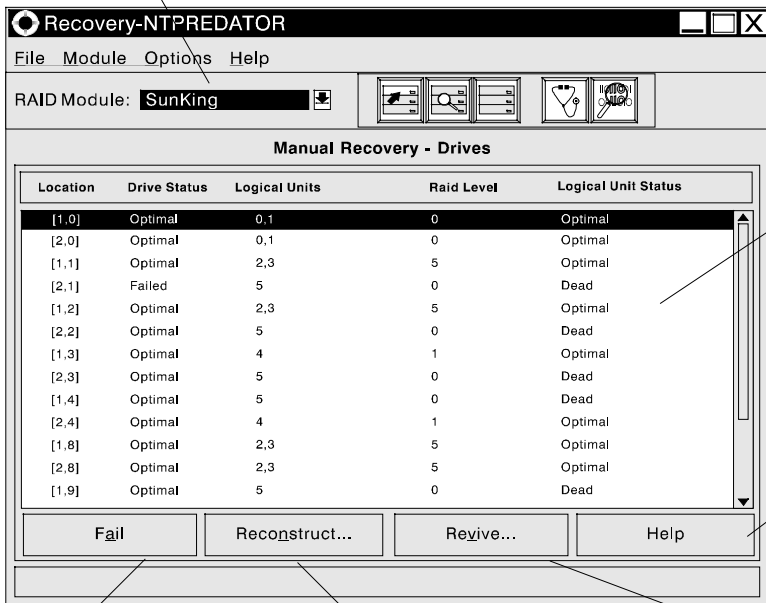
- Fail a drive — you may need to do this if the drive is causing problems but the controller won't fail it.
- Reconstruct data on a drive — you may need to do this if reconstruction doesn't occur automatically when you replace a failed drive.
- Revive a drive — you may need to do this if you accidentally fail or remove a drive. See “Reviving Drives” on page 114 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 Module containing the drives you want to recover.**
3. **Select Options → Manual Recovery → Drives from the drop-down menus. A screen like FIGURE 6-4 is displayed.**
4. **Select the individual drive you want to fail, reconstruct, or revive, then click on the appropriate button to perform the operation.**

End of Procedure

Select the RAID Module 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 114 for a description of when to use this option.

FIGURE 6-4 Manual Recovery/Drives Screen

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

Failing a drive in a Degraded RAID 3 or 5 logical unit will result in the loss of data.

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.

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

Reviving a drive should be done only under the circumstances described in "Reviving Drives" on page 114.

Manual Recovery/Logical Units



Caution – Take care when using manual recovery. In general, do not use these options unless specifically directed by the procedure screen of Recovery Guru or by 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 RAID Module containing the logical units you want to format or revive.**
3. **Select Options → Manual Recovery → Logical Units from the drop-down menus. A screen like FIGURE 6-5 is displayed.**
4. **Select the logical units you want to format or revive, then click on the appropriate button to perform the operation.**

End of Procedure

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



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

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 6-5 Manual Recovery/Logical Units Screen

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

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

Reviving a logical unit should be performed only if the logical unit failure was caused by a channel failure. Do not attempt to revive a logical unit unless instructed to by Recovery Guru. Doing otherwise can result in corrupted data or another failure. See “Multiple Drive Failures And Channel Failures” on page 114 for more information.

Manual Recovery/Controllers



Caution – Take care when using manual recovery. In general, do not use these options unless specifically directed by the procedure screen of Recovery Guru or by 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 RAID Module containing the controllers you want to place online or offline.**
- 3. Select Options → Manual Recovery → Controller Pairs from the drop-down menus. A screen like FIGURE 6-6 is displayed.**
- 4. Select the controller you want to place online or offline and then click on the appropriate button to perform the operation.**

End of Procedure

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



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

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 Controller Mode" on page 153 for more information).

FIGURE 6-6 Manual Recovery/Controllers Screen

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 16).

If you place an offline controller online, you may need to go into the Maintenance and Tuning application to reset the controller modes and assign LUNs to the online controller (see "Changing Controller Mode" on page 153).

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 147.

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

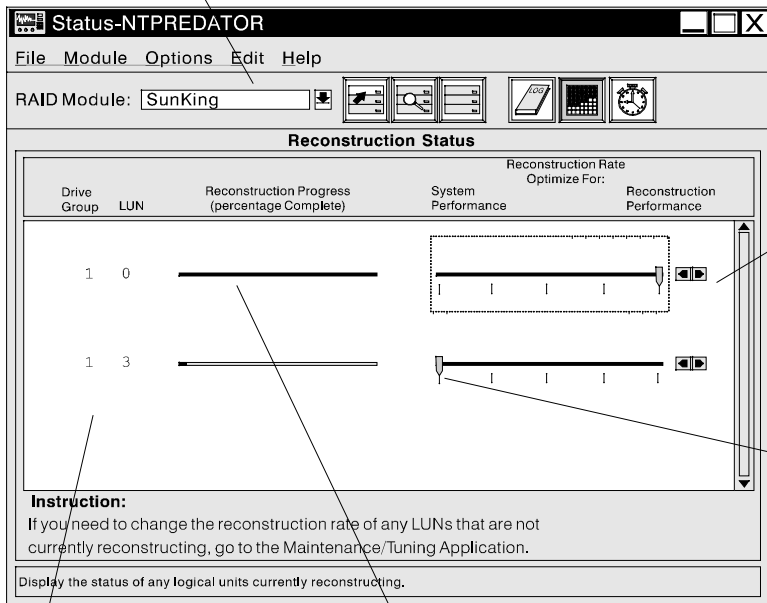
- 1. Start the Status Application.**
- 2. Select the RAID Module containing the logical units currently being reconstructed. Do not select All RAID Modules.**
- 3. Click on the LUN Reconstruction button or select Options → LUN Reconstruction from the drop-down menus. A screen like FIGURE 6-7 is displayed.**
- 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.**

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

End of Procedure

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



Each row displays information on a single logical unit.

Move the slider bar here to change the reconstruction rate. Moving to the left slows reconstruction but increases system performance. Moving to the right speeds reconstruction at the expense of system performance.

These columns show the drive group and LUNs currently being reconstructed.

This histogram shows the amount of reconstruction accomplished as a percentage. If you exit this screen and then return, any LUNs that have completed reconstruction will no longer be displayed.

FIGURE 6-7 Viewing/Changing Reconstruction Screen

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

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

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

- Slow — 1024 blocks/0.8 second delay
- Medium slow — 1024 blocks/0.4 second delay
- Medium — 1024 blocks/0.2 second delay
- Medium fast — 1024 blocks/0.1 second delay
- Fast — reconstruct with no delays.

Maintenance and Tuning

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

- Maintenance And Tuning Overview—page 132
- Upgrading Controller Firmware—page 133
- Running Parity Check—page 142
- Changing The LUN Reconstruction Rate—page 147
- Balancing LUNs—page 149
- Changing Controller Mode—page 153
- Changing Cache Parameters—page 155

Maintenance And Tuning Overview

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

- Download new controller firmware and NVSRAM files (see “Upgrading Controller Firmware” on page 133).
- Manually run Parity Check (see “Running Parity Check” on page 142).
- Set the time for Automatic Parity Check (see “Running Parity Check” on page 142).
- Change the Reconstruction Rate for logical units (see “Changing The LUN Reconstruction Rate” on page 147).
- Balance drive groups/logical units across active/active redundant controller pairs (see “Balancing LUNs” on page 149).
- Swap active/passive controllers or change them to active/active (see “Changing Controller Mode” on page 153).
- Change the write cache, write cache mirroring, and cache without batteries parameters for logical units (see “Changing Cache Parameters” on page 155).

Upgrading Controller Firmware

- Controller Firmware Files—page 133
- General Procedure—page 134
- Online And Offline Method—page 134
- Selecting Controllers—page 135
- Download Procedure—page 136
- Confirming The Download—page 140

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

Tip – Controller firmware is different from the drive firmware. Use this option only to upgrade controller firmware. If you need to upgrade drive firmware, call your Customer Services representative.

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 `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 and Firmware version levels through the Controller Profile screen; see “The Controller Profile Screen” on page 50.

General Procedure

In general, use the following steps to 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 Support Guide* for information on which directory to use).**

Note – If you do not copy the `fwcompat.def` file to the correct subdirectory, the software is unable to check the files for compatibility. Although you can still enter firmware 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. **Turn the power to the RAID module 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 in “Download Procedure” on page 136.

Online And Offline Method

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

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

- Select Online to upgrade firmware while the selected RAID Module receives I/O. You may only use the Online method if:
 - The controllers have the (RDAC) driver installed for redundant controller support.

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

The software 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 Module 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 Support Guide*).**
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 Modules 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 Modules to download firmware to all controllers in all RAID Modules.
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 Module receives I/O.
 - Offline to upgrade firmware when the selected RAID Module 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 Modules for restrictions based on the type of firmware upgrade you selected. If there are no restrictions, a screen like FIGURE 7-1 is displayed.**
9. **Depending on whether you are downloading NVSRAM files or upgrading controller firmware, do one of the following:**

Tip – 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.
 - 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.

Offline Firmware Upgrade

Current Firmware Version(s):

RAID Module	Controller	Firmware Level	Boot Level	Fibre Channel Level
star_a_001	A (Drive16)	02.04.03	02.04.03	-
star_a_001	B (Drive1)	02.04.03	02.04.03	-

Instruction:
 Highlight files to download. These files are listed in compatible groups. If the files are not displayed or you want to download an NVSRAM file, then enter a path and/or file(s) in the Path area. For firmware files, you can specify multiple files in the same path by using a "+" between the files.

Compatible Files/Versions:

Firmware Level (AP)	Boot Level (BW)	Fibre Channel Level (FC)
03.00.00.00	03.00.00.00	

Path:

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.

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

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

Click here to cancel without downloading any new files.

FIGURE 7-1 Firmware Upgrade Screen

Note – 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 7-1 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 140.

TABLE 7-1 Firmware Confirmation Screen Elements

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

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 Modules 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 module off and then on again to establish the new settings.**
- 16. Go back to Step 9 and download any new firmware files.**

End of Procedure

Confirming The Download

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

- If you see that the upgrade was Successful, you should still verify that all the logical units (LUNs) are *not* assigned to only one controller. See “Balancing LUNs” on page 149 for the procedures to balance the LUNs across both controllers.
- If you see “Failed” for any module, you should fix the specified failure and try the firmware upgrade procedure again. See TABLE 7-2 for possible actions to take to correct a failed upgrade.

TABLE 7-2 Controller Firmware Upgrade: Failed Statuses

Reasons for Failed Status	Actions
The selected module had I/O activity occurring or filesystems 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 6.
Upgrading to the selected firmware version requires that you use the Offline method.	Try to upgrade the firmware again and this time be sure to select Offline.
The current firmware version is unable to upgrade to the files you selected.	Most likely, you need to upgrade to an intermediate version of firmware. Try to upgrade to a version earlier than the one you selected. If that upgrade is successful, perform a second upgrade for this latest firmware version.

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

Reasons for Failed Status	Actions
The files you selected are not compatible with the current firmware version on the selected module's controllers.	Most likely, the current directory does not contain <i>all</i> the necessary firmware files. Copy the firmware files and the <code>fwcompat.def</code> file to the correct subdirectory in the installation directory and try again. Be sure the version you select has both Firmware Level and Bootware Level versions specified. If the upgrade fails a second time, obtain a new copy of the firmware upgrade files.
The software was unable to access the controllers during the upgrade process.	Use Recovery Guru/Health Check in the Recovery Application to determine if the module has a failure. See Chapter 6. <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.
You tried to load a pre-2.04 firmware version, which is not supported by this software or the redundant controller configuration.	Do <i>not</i> try to load this firmware version again.
The selected firmware files are not compatible with your controller model.	Use Module Profile → Controller details to check your controller type and model (see FIGURE 3-6), and obtain the correct firmware version files.
The online upgrade cannot be performed because either the selected module has only one controller or one of the controllers in the pair is not accessible.	Use Module Profile → Controller details to determine how many controllers the module has (see FIGURE 3-6). <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).
An unknown failure occurred.	Use the Status Application to select Message Log for component information (see “Viewing The Log” on page 94).

Running Parity Check

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

RAID 0 logical units do not have parity and therefore 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 Modules to perform Automatic Parity Check/Repair on a regular basis. You can set it to run either daily or once a week. At times, you may need to run Parity Check/Repair manually as part of the recovery process. The two following sections describe how to set the time for Automatic Parity Check and how to run a Manual Parity Check.

- Setting Automatic Parity Time—page 145
- Running Manual Parity Check—page 145

Setting Automatic Parity Time

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

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

To set the time for Automatic Parity Check/Repair, do the following:

1. **Start the Maintenance and Tuning application.**
2. **Select any RAID Module or All Raid Modules. The automatic time always applies to all RAID 1, 3, and 5 logical units on all the RAID Modules on your system, even if you select an individual module.**
3. **Select Option → Auto Parity Settings from the drop-down menus. A screen like FIGURE 7-2 is displayed.**
4. **Enter the settings you want, then click Save to make the settings effective, or press Cancel to return to the application without changing the settings.**

End of Procedure.

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

Automatic Parity Check/Repair Settings

Parity Settings (Applies To All RAID Modules)

Enable Automatic Parity Check/Repair

Start Time (24-Hour Format): 02 00 Range 00:00 to 23:59
00:00 = 12:00 am
23:59 = 11:59 pm

Day of Week Sunday

Save Cancel

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.

33418

FIGURE 7-2 Setting The Time For Automatic Parity Check/Repair

Running Manual Parity Check

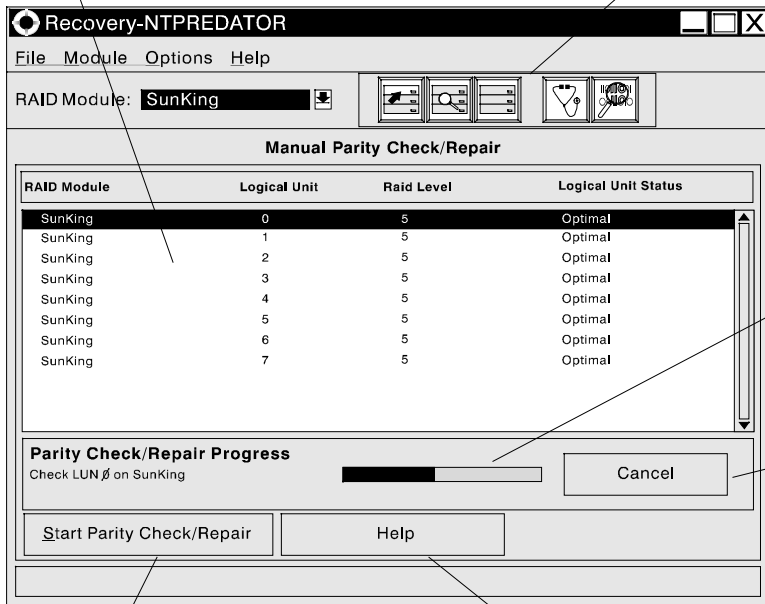
Normally, you should set your RAID Modules so that Parity Check is run automatically on a regular basis (see “Setting Automatic Parity Time” on page 143 for information on how to set the time for Automatic Parity Check). However, during recovery, or if you suspect a problem, you may need to run manual 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 Modules).**
- 4. Click the Manual Parity Check/Repair button or select Options → Manual Parity Check/Repair from the drop-down menus. A screen like FIGURE 7-3 is displayed.**
- 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.**

End of Procedure

Select the logical units you want to check here. You cannot 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 7-3 Manual Parity Check Repair Screen

You cannot run Parity Check/Repair on a RAID 0 logical unit or a non-optimal RAID 1, 3, or 5 logical unit.

Parity repair fixes parity, not data. If the parity inconsistencies were caused by corrupted data, the data is still corrupted, even though the parity is correct.

While parity check/repair is in progress, you cannot perform other Recovery tasks.

Changing The 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 16 for more information.

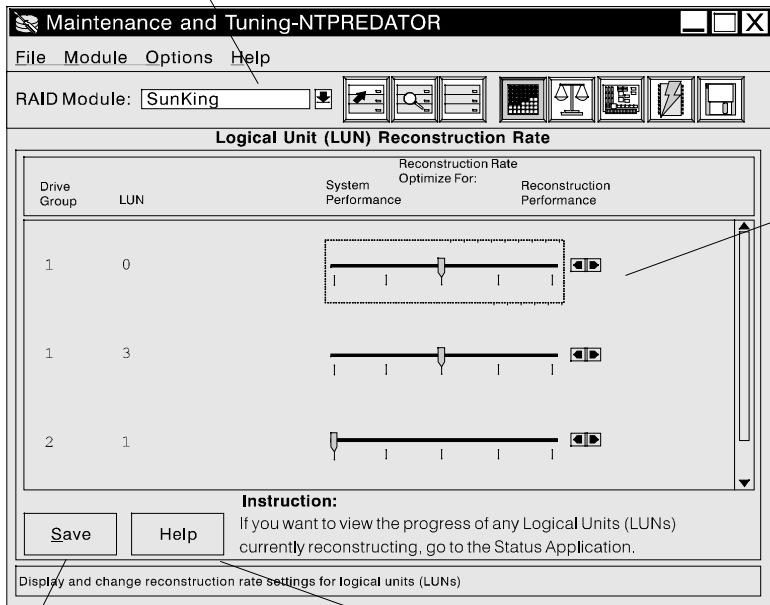
Use this option to change the LUN Reconstruction Rate of logical units that are not currently being reconstructed. To change the rate of logical units that are currently being reconstructed, see “Monitoring/Changing Reconstruction” on page 128.

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 Modules.**
- 3. Click on the LUN Reconstruction button or select Options → LUN Reconstruction from the drop-down menus. A screen like FIGURE 7-4 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.**

End of Procedure

You must have an individual RAID Module selected here.



This display shows all the logical units in the selected RAID Module 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 7-4 Changing The Reconstruction Rate Screen

You cannot perform this option if you select All RAID Modules.

The new reconstruction rate will go into effect the next time the selected logical units are reconstructed.

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 128.

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

- Slow — 1024 blocks/0.8 second delay
- Medium slow — 1024 blocks/0.4 second delay
- Medium — 1024 blocks/0.2 second delay
- Medium fast — 1024 blocks/0.1 second delay
- Fast — reconstruct with no delays.

Balancing LUNs

Balancing LUNs involves assigning drive groups/LUNs to controllers in an active/active redundant controller pair (see “Redundant Controllers” on page 16 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 149).
- Manually balance the logical units in one RAID Module (see “Manual LUN Balancing” on page 151).

Automatic LUN Balancing

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

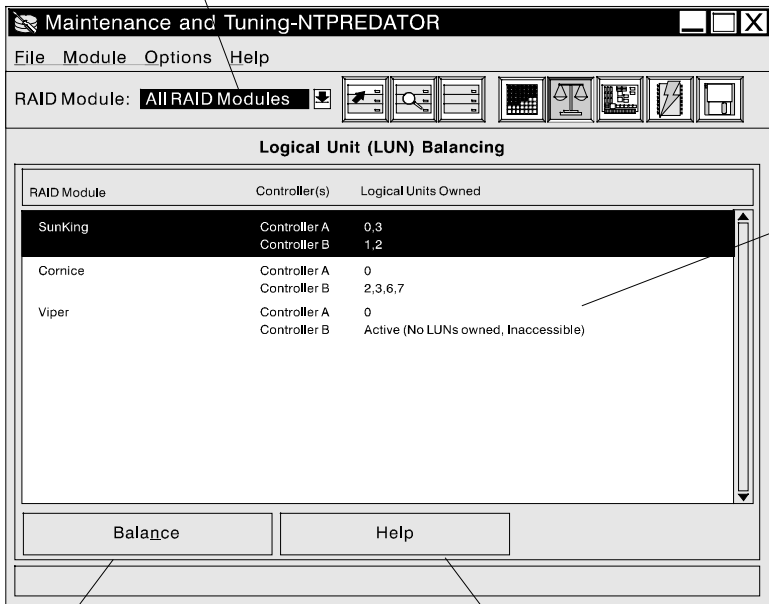
To automatically balance drive groups/logical units in RAID Modules, use the following procedure.

1. **Start the Maintenance and Tuning application.**
2. **Select All RAID Modules. (If you select a single RAID Module, you can only balance the logical units manually; see “Manual LUN Balancing” on page 151).**

3. Click on the LUN Balancing button or select Options → LUN Balancing. A screen like FIGURE 7-5 is displayed.
4. Select the RAID Modules you want to balance, then click Balance to automatically balance them.

End of Procedure

You must have All RAID Modules selected here. If you have selected a specific RAID Module, see “Manual LUN Balancing” on page 151.



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

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

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

Click here for help.

33392

FIGURE 7-5 Automatic LUN Balancing Screen

You can only select RAID Modules with active/active controller pairs.

You cannot select controllers in an independent configuration (you must use manual LUN balancing to reassign LUNs in this configuration; see “Manual LUN Balancing” on page 151).

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.

Odd-numbered drive groups are assigned to one controller and even numbered drive groups are assigned to the second controller.

Logical units/drive groups will not be balanced on RAID Modules without RDAC protection unless all I/O to that module is stopped.

If none of the displayed RAID Modules are eligible for balancing, the Balance button is grayed out.

Manual LUN Balancing

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




Caution – If the RAID Module 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 Module containing the logical units/drive groups you want to reassign.**
(If you select All RAID Modules, you can only automatically balance all the drive groups in all your RAID Modules; see “Automatic LUN Balancing” on page 149).
- 3. Click on the LUN Balancing button or select Options → LUN Balancing. A screen like FIGURE 7-6 is displayed.**
- 4. Make the re-assignments you want by highlighting the drive groups and then clicking on the Move button.**
- 5. After you have made all the changes you want to make, click Save to make the changes.**

End of Procedure

You must have a specific RAID Module selected here. If you have selected ALL RAID Modules, see "Automatic LUN Balancing" on page 149.

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 7-6 Manual LUN Balancing Screen

You can only select RAID Modules with active/active controller pairs.

If the RAID Module does not have RDAC protection, you must stop all I/O to the module before balancing the LUNs.

In an independent controller configuration, you can only reassign LUNs *from* the inaccessible controller *to* the accessible controller (the controller connected to the host system you are using). You can not assign LUNs from the accessible controller to the inaccessible controller.

Changing Controller Mode

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

Use this option to:

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

Note that you can not use this option to change controllers from active/active to active/passive. This can only be done with the `rdacutil` command. See the `rdacutil` command entry in Appendix A for information on this command.

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



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


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

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

3. **Click on the Controller Mode button or select Option → Controller Mode from the drop-down menu. A screen like FIGURE 7-7 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.

End of Procedure

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



This displays shows the current active/passive status of the controllers in the selected RAID Modules.

If you selected All RAID Modules, you must highlight the controller pair you want to change.

These buttons will be grayed out if the currently selected controllers are not active/passive.

Click here to change the controllers to active/active mode. A confirmation message will allow you to specify whether you want to automatically balance the drive groups/LUNs across the controllers.

Click here to make the active controller passive and the passive controller active. All LUNs will be re-assigned to the active controller.

FIGURE 7-7 Change Controller Mode Screen

If the RAID Module does not have RDAC protection, you must stop all I/O to the module before changing controller mode.

You will not be able to select this option unless the selected RAID Module contains an active/passive controller pair, or, if you selected All RAID Modules, at least one controller pair is active/passive.

You can not change active/active controllers to active/passive using this application. See the `rdacutil` command entry in Appendix A.

When you swap active/passive controllers, the logical units are automatically switched to the other controller.

Changing Cache Parameters

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

You can set three cache memory parameters with this option: write caching, write cache mirroring, and caching without batteries. You can change other cache options using the command line `raidutil` command. See the `raidutil` command entry in Appendix A for information on this command.

You can only change cache parameters if:

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

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

Use the following procedure to change the cache parameters.

1. **Start the Maintenance and Tuning application.**
2. **Select the RAID Module containing the logical units you want to set cache parameters for.**
3. **Click on the Cache Parameters button or select Option → Cache Parameters from the drop-down menus. A screen like FIGURE 7-8 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.

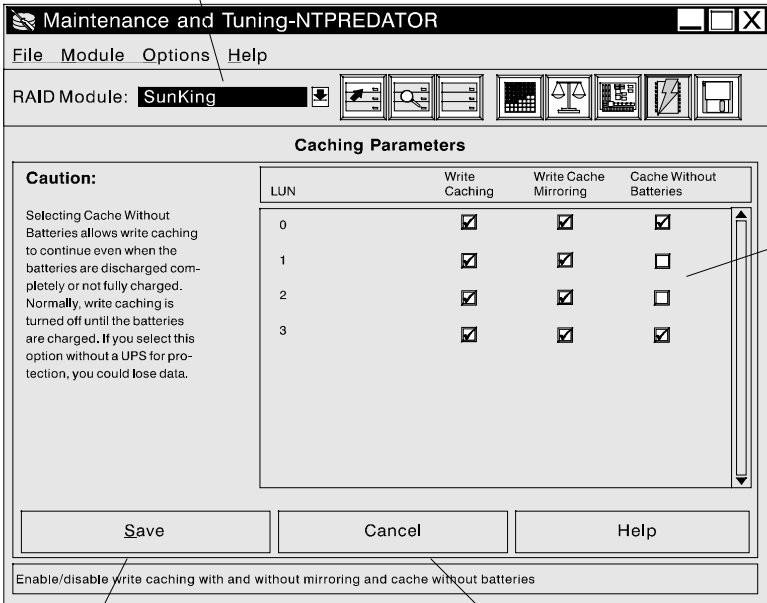
4. **Click in the boxes to select or deselect individual cache parameters for the logical units.**

The parameters are interdependent; clicking on some will automatically select or deselect others. See TABLE 7-3.

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

End of Procedure

Select an individual RAID Module here. You cannot set cache parameters if All RAID Modules is selected.



This display shows the cache parameter settings for each logical unit in the selected RAID Module. If an option is grayed out, it is because the controllers in the RAID Module 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 7-3).

Click here to save the current cache parameter settings.

Click here to exit this screen without making any changes.

FIGURE 7-8 Cache Parameters Screen

TABLE 7-3 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.

Command Line Interface

This Appendix describes the command line interface to the storage management software.

- Command Line Commands—page 160

Command Line Commands

TABLE A-1 summarizes the Command Line commands available in the storage management software. All of these commands have associated Help files available.

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

TABLE A-1 Command Line Utilities And Program Description

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

TABLE A-1 Command Line Utilities And Program Description (*Continued*)

Option	Description
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 <code>symping</code> 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.
rdaemon (UNIX only)	The redundant I/O path error resolution daemon. The primary function of <code>rdaemon</code> is to receive and react to redundant controller exception events and to participate in the application-transparent recovery of those events through error analysis and, if necessary, controller failover.
rdriver (Solaris only)	The redundant I/O path routing driver. The <code>rdriver</code> module works in cooperation with <code>rdaemon</code> in handling the transparent recovery of I/O path failures. Its primary responsibilities include routing I/Os down the proper path and communicating with the <code>rdaemon</code> about errors and their resolution.

TABLE A-1 Command Line Utilities And Program Description (*Continued*)

Option	Description
symsmdisk (HP-UX only)	The redundant I/O path routing driver. The <code>symsmdisk</code> module works in cooperation with <code>rdaemon</code> in handling the transparent recovery of I/O path failures. Its primary responsibilities include routing I/Os down the proper path and communicating with the <code>rdaemon</code> about errors and their resolution.
Customizable Elements	
rmparams	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 <code>rmparams</code> are changeable under the graphical interface. For more information on this file, see the <i>Installation and Support Guide</i> .
rmscript	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 (<code>rmllog.log</code>), sending e-mail to the superuser/administrator, and, in some cases, sending an SNMP trap. Although you can edit <code>rmscript</code> , make certain that you do not disturb any of the standard actions. For more information on this file, see the <i>Installation and Support Guide</i> .

Index

A

- Active/active controllers, 15, 17, 18
 - balancing, 149, 152
 - changing to active/passive, 154
- Active/passive controllers, 15, 17
 - changing mode, 32, 153
 - LUN ownership after swap, 154
- Add (module), 43
- Adding a drive, 78
- Appware, 133
 - see also Firmware
- arraymon utility, 161
- ASC/ASCQ codes, 93, 97
- Automatic LUN balancing, 149
 - see also Balancing Logical units
- Automatic parity check/repair
 - see Parity check/repair

B

- Balancing logical units, 32, 65, 149
 - LUN creation and, 63, 72
 - restrictions, 151
- Batteries
 - caching and, 9
 - see also Cache without batteries
- Board ID, 51
- Board name, 51
- Board serial number, 51
- Boot level, 51
 - see also Firmware, 51

- Bootware, 133
 - see also Firmware

C

- Cache memory
 - cache options and, 155
 - defined, 9
 - performance monitor and, 105, 107
 - size, 51
 - see also Cache options
- Cache options
 - asterisks in display, 155
 - cache without batteries, 9, 55, 155
 - changing, 32
 - interdependencies, 157
 - other, 155
 - rdacutil command and, 73
 - selecting, 65, 70, 155
 - restrictions, 155
 - selections for, 55
 - write cache mirroring, 9, 25, 55, 155
 - restrictions, 155
 - write caching, 9, 55, 155
- Cache parameters option, 70
- Cache size, 155
- Cache without batteries, 155
 - defined, 9
 - setting, 55, 155
 - see also Cache options
- Cache/Processor size, 51
- Capacity of drive group

- changing, 78
- Channel ID
 - drives, 13, 45
- Channels
 - creating LUNs and, 26, 69
 - drive, 26
 - failures on, 26, 114, 124
 - hot spares and, 85
 - see also Channel ID
- Command line commands
 - drivutil, 160
 - fwutil, 160
 - healthck, 160
 - lad, 160
 - logutil, 160
 - nvutil, 134, 138, 161
 - parityck, 161
 - raidutil, 155, 161
 - rdacutil, 9, 73, 153, 154
 - storutil, 161
 - symping, 161
- Command line interface, 159
 - see also Command line commands, 159
- Component failures
 - see Failures
- Configuration
 - independent controller, 21
 - downloading firmware, 135
 - multi-host, 20
 - network, 24
 - overview, 58
 - recommended, 19
 - resetting, 32, 89
 - single-host, 19
- Configuration application, 62, 83
 - description of, 27
 - module information display, 60
- Configured drives
 - as spares, 115
- Controller firmware download utility (fwutil), 160
- Controllers
 - active/active, 17, 149, 153
 - active/passive, 17, 153
 - assigning, 72
 - automatic LUN assignment, 63, 66, 149
 - balancing, 32, 149
 - changing mode, 32, 153

- firmware, 133
 - downloading, 136
 - upgrading, 133
- inaccessible status, 21
- manual LUN assignment, 149
- manual recovery options, 126
- names of, 9, 31
- naming, 43
- placing offline, 126
- placing online, 126
- profile information, 50
- same host bus
 - restrictions, 19
- selecting for download, 135
- see also Redundant Controllers

Creating logical units

- see Logical units, creating

D

- Daemon processes
 - arraymon, 161
 - rdaemon, 161
 - rdriver, 161
 - symsmdisk, 162
- Data connection
 - failures, 16, 116
 - see also Redundant controllers
- Data path
 - see Data connection
- Date code, 53
- Date of manufacture, 51
- Dead logical unit, 112
- Dead mode, 112
- Defragment option
 - defined, 74
 - procedure, 82
- Degraded logical unit, 112
 - defined, 16
 - parity check and, 142
- Degraded mode, 112
 - failing drives and, 123
- Deleting logical units, 83
 - network configuration, 84
 - restrictions, 83
- Device names

- controllers, 9
- drives, 13
- logical units, 9
- user defined, 9

Downloading controller firmware, 136

Drive groups

- adding drives to, 31, 78
- balancing, 149
- configured drives, 10
- creating, 62, 65
- defined, 10
- defragmenting, 74, 82
- hot spare, 10
- locating, 45
- modifying, 59, 78
- numbering, 11
- selecting drives, 69
 - restrictions, 69
- space allocation on, 74
- unassigned, 10, 83
- viewing, 60

Drive selection option, 69

Drive/LUN utility (`drivutil`), 160

Drives

- adding to drive group, 78
- channel number, 13, 45
- channels, 26, 114
- different sizes, 63, 69, 85
- drive groups, 10
 - numbering, 11
- failing, 122
- failures on, 112
- firmware, 53, 133
- hot spares, 13, 25
- locating, 45
- manual recovery options, 122
- names, 13
- numbering, 45
- profile information, 52
- reconstructing, 122
- replacing, 113
- reviving, 114, 122
- SCSI ID, 13, 45
- serial number, 53
- status, 53
- using configured drives as spares, 115

`drivutil` command, 160

DRM (RAID Manager for DOS), 4

E

- Edit (module), 43
- Exiting the software, 37

F

Failing drives, 122

Failures

- accidental, 114
- channel, 26, 69, 114, 124
- controllers, 16, 116
- data connection, 16
- drives, 112
- fans, 117
- general information, 25
- message log and, 93
- power supplies, 117
- preventing, 25
- recovering from, 110
- recovery guru/health check and, 25
- recovery information for, 120
- redundant controllers, 16
- replacing components, 33

Find (module), 43

Firmware

- boot level, 133
- controller, 133
- downloading, 32
 - confirming download, 140
 - general steps, 134
 - online/offline method, 134
 - procedure, 136
 - progress of, 137
- drives, 133
- firmware level, 51, 133
- `fwcompat.def` file, 133
- redundant controller requirements, 135
- upgrading, 133
- versions, 135

Firmware level, 51

- drives, 53
- see also Firmware

Formatting logical units, 124

`fwcompat.def` file, 133, 136

`fwutil` command, 160

G

- GUI
 - navigating in, 36
 - online help, 38
 - overview, 160

H

- Hardware guides
 - redundant controllers, 17
 - when to use, 3, 110, 117
- Health Check
 - see Recovery guru/health check
- Health check command (`healthck`), 160
- Host bus
 - caution with dual controllers, 19
 - see also SCSI bus
- Host interface, 51
- Host store utility (`storutil`), 161
- Host-side ID
 - module profile, 51
- Hot spares, 25
 - creating, 85
 - defined, 13
 - deleting, 32, 88
 - drive group, 10
 - drive sizes to use, 85
 - independent controllers and, 22
 - LUN assignment, 13
 - number allowed, 85, 87
 - recovery and, 113
 - restrictions on creating, 85
 - viewing, 60

I

- Inaccessible status, 21
- Independent controller configuration, 21, 23
 - hot spares and, 22
 - inaccessible status and, 21
 - `rdac` support for, 22
 - recovery and, 22
 - see also RAID Modules, configuration
- Installation and support guide
 - description of, 3

- software installation and, 4
- troubleshooting information in, 3

K

- Keyboard, 36

L

- `lad` command, 160
- List array devices utility (`lad`), 160
- Locating
 - drive groups, 45
 - drives, 45
 - logical units, 45
 - RAID modules, 44
- Log format utility (`logutil`), 160
- Logical units
 - adding drives to, 78
 - cache options
 - selecting, 65, 70, 155
 - capacity
 - defaults, 68
 - setting, 65, 68
 - changing RAID level, 76
 - checking parity on, 142
 - controller ownership, 17, 18, 32
 - changing, 149
 - creating, 31, 58
 - controller assignment, 72
 - drive group numbering, 11
 - from unassigned drives, 62
 - in existing drive group, 65
 - operating system and, 58, 65, 66
 - overview, 58
 - RAID level, 62
 - segment size, 65, 71
 - selecting drives, 69
 - space available for, 75
 - dead, 112
 - defined, 13
 - degraded, 16, 112
 - deleting, 32, 83
 - drive group numbering, 11
 - operating system and, 83
 - drive groups, 10

- formatting, 124
- locating, 31, 45
- manual recovery options, 124
- modifying, 59
- names of, 9
- number allowed, 18, 22
- profile information, 54
- reconstruction rate of, 128, 147
- recovering, 113
- reviving, 124
 - when, 125
- segment size, 71
 - changing, 80
- status, 55
- viewing, 60

logutil command, 160

LUN

- see Logical units

LUN assignment option, 72

LUN capacity option, 68

LUN creation and, 65

M

Maintenance and Tuning application, 136, 155

- description of, 30
- overview, 132

Manual LUN balancing, 151

- see also Balancing Logical units

Manual parity check/repair

- see Parity check/repair

Manual recovery options, 121

- controllers, 126
- drives, 122
- logical units, 124
- when to use, 111

Maximum size limit

- effects when exceeded, 37, 101
- message file and, 92, 99
- resetting file, 101
- warning message, 101

Message log file, 92, 111

- backup file names, 101
- changing settings, 99
- clearing, 37, 101
- copying, 98

- default, 93
- maximum size limit, 37, 92, 99, 101
- message types, 93
- recovery guru and, 93
- saving, 98
- viewing, 94

Module information display, 60

Module profile

- controllers, 50
 - SCSI, 50
- drives, 52
- logical units, 54
- saving, 56
- viewing, 47
- see also RAID modules

Mouse, 36

Multi-host configuration, 20

- number of LUNs on, 20
- operations on, 20
- rdac support for, 20
- reservations and, 20
- see also RAID Modules, configuration

Multiple drive failures, 113

N

Navigation, 36

- keyboard, 36
- mouse, 36
- online help, 38

Network configuration, 24

- adding modules, 43
- deleting logical units, 84
- exclusive access and, 24
- naming controllers, 43
- RDAC and, 24
- storutil command, 161
- symping command, 161
- see also RAID Modules, configuration

Network connection verification utility

- (symping), 161

Network management station, 17

NMS, 17

Number of drives

- setting, 62

Number of LUNs

- setting, 62, 65
- NVSRAM, 136
 - defined, 133
 - downloading, 32, 135, 136
 - nvutil command, 161
 - resetting, 89
 - restart after downloading, 134, 139
- NVSRAM utility (nvutil), 161
- nvutil command, 134, 138, 161

O

- Offline controller, 126
- Offline method (firmware download), 134
- Online controller, 126
- Online help, 38
- Online method (firmware download), 134
- Operating system
 - dead mode and, 112
 - degraded mode and, 112
 - deleting logical units and, 83
 - device names on, 9
 - installation and support guide, 3
 - logical unit creation and, 58, 65, 66
 - logical unit deletion and, 83, 84
 - messages from, 111
 - RDAC on, 16, 17
 - recovery and, 111

P

- Parity, 142
 - defined, 14
 - how used, 142
 - RAID 1 and, 14, 142
 - RAID 3, 14, 15, 142
 - RAID 5, 14, 15, 142
 - reconstruction and, 16
 - see also Parity Check/Repair
- Parity check/repair
 - automatic, 142
 - current setting, 144
 - disabling, 144
 - setting time, 33, 143
 - corrupted data and, 142
 - defined, 14

- eligible LUNs, 146
- manual, 142
 - restrictions, 146
 - running, 145
 - when to run, 33, 145
- parityck command, 161
- RAID 1 and, 14
- RAID 3, 14
- RAID 5, 14
- restrictions, 142, 146
- running, 142
- viewing results of, 33, 93
- what is repaired, 142

- Parity check/repair command (parityck), 161

- Performance monitor, 103
 - polling interval, 103
 - saving data from, 106
 - screen elements, 105
 - starting, 103
 - using data from, 107
 - when to run, 31

- Polling interval
 - performance monitor, 103
 - status, 92
 - changing, 99

- Processor size, 51
- Product ID, 51
- Product revision, 51
- Product serial number, 51

R

- RAID 0
 - dead mode, 112
 - defined, 15
 - failing drives and, 123
 - parity check and, 142
 - recovery on, 113
 - reviving drives and, 114
 - see also RAID level
- RAID 1
 - dead mode, 112
 - defined, 15
 - degraded mode, 112
 - failing drives and, 123
 - parity and, 14
 - parity check/repair and, 14, 142

- recovery on, 113
- reviving drives and, 114
- when to use, 25
- see also RAID level
- RAID 3
 - dead mode, 112
 - defined, 15
 - degraded mode, 112
 - failing drives and, 123
 - parity, 14
 - parity check/repair, 14, 142
 - recovery on, 113
 - reviving drives and, 114
 - when to use, 25
 - see also RAID level
- RAID 5
 - dead mode, 112
 - defined, 15
 - degraded mode, 112
 - failing drives and, 123
 - parity, 14
 - parity check/repair, 14, 142
 - recovery on, 113
 - reviving drives and, 114
 - when to use, 25
 - see also RAID level
- RAID configuration utility (`raidutil`), 161
- RAID level
 - changing, 76
 - RAID 0, 15
 - RAID 1, 15
 - RAID 3, 15
 - RAID 5, 15
 - setting, 62
 - when to use, 25
 - see also individual RAID level entries
- RAID Manager for DOS, 4
- RAID modules
 - adding, 43
 - boot device, 4
 - configuration
 - independent, 21
 - multi-host, 20
 - network, 24
 - overview, 58
 - single-host, 19
 - defined, 15
 - examples, 15
 - failure recovery, 110
 - finding, 43
 - identifying, 41
 - locating, 31, 44
 - module information display, 60
 - naming, 31, 43
 - profile information, 31, 47
 - saving, 56
 - removing, 43
 - selecting, 41
 - tasks, 31
- `raidcode.txt` file, 160
- `raidutil` command, 155, 161
- RDAC
 - controller failures and, 116
 - controller mode changes and, 154
 - defined, 15
 - failover conditions, 117
 - firmware download and, 134
 - independent configuration and, 22
 - LUN balancing and, 149, 151, 152
 - multi-host configuration and, 20
 - overview, 160
 - `rdaemon` utility, 161
 - `rdriver` utility, 161
 - redundant controllers and, 16
 - single-host configuration and, 19
 - `symmsdisk` utility, 162
- `rdac` text file, 160
- `rdacutil` command, 9, 73, 153, 154
- `rdaemon` utility, 161
- `rdriver` utility, 161
- Reconstructing drives, 122
- Reconstruction, 122, 123
 - defined, 16
 - monitoring, 115, 128
 - rate of, 16
- Reconstruction rate
 - changing, 33, 115, 128, 147
 - restrictions, 148
 - defined, 16
 - settings, 129, 148
- Recovery
 - component failures and, 25
 - failed controllers, 116
 - failed drives, 112
 - failed fans, 117

- failed power supplies, 117
 - independent controllers and, 22
 - instructions for, 120
 - manual recovery options, 121
 - overview, 110
 - reconstruction rate, 128
 - tasks, 33
- Recovery application, 118, 122
- description of, 29
 - overview, 110
- Recovery guru/health check, 25
- independent configuration and, 22
 - message log and, 93
 - operation, 111
 - show procedures, 120
 - starting, 118
 - viewing procedures, 110
 - when to use, 33, 111
- Redundant controller management utility
(`raidutil`), 161
- Redundant controllers
- active/active, 15, 17, 18, 153
 - active/passive, 15, 17, 153
 - changing mode of, 153
 - changing status of, 127
 - described, 16
 - failover support, 20, 22, 24
 - firmware requirements, 135
 - LUN ownership and, 18, 149
 - recommended configuration, 17
 - see also Controllers
 - see also RDAC
- Remove (module), 43
- Reservations, 20
- Resetting configuration, 89
- Reviving drives, 114, 122
- Reviving logical units, 124
- `rmevent` file, 160
- `rmparams` file, 162
- `rmscript`
- description, 160
 - file, 162
- S**
- SCSI bus
- independent controller configuration, 21, 23
 - module profile for, 50
 - multi-host configuration, 20
 - single-host configuration, 19
 - see also Host Bus
- SCSI ID
- drives, 13, 45
 - see also Host-side ID
- Segment size
- changing, 32, 80
 - defined, 17
 - selecting, 65, 71
- Segment size option, 71
- Serial link, 4
- Serial number
- controller, 49, 51
 - drives, 53
 - product, 51
- Show procedures, 120
- see also Recovery guru/health check
- Simple Network Management Protocol
- see SNMP
- Single drive failures, 113
- Single-host configuration, 19
- `rdac` support for, 19
 - see also RAID Modules, configuration
- SNMP, 94, 111
- defined, 17
- Software
- exiting, 37
 - installing, 4
 - online help for, 38
 - starting, 37
 - what's new, 34
- Starting the software, 37
- Status, 92, 111
- dead mode, 112
 - degraded mode, 112
 - drives, 53
 - inaccessible, 21
 - logical units, 55
 - message file, 31, 92, 111
 - backup file name, 101
 - changing settings, 99
 - clearing, 37, 101
 - copying, 98
 - default, 93

- maximum size limit, 37, 92, 101
- message types, 93
- opening, 98
- viewing, 94
- polling, 92, 99
- Status application, 103
 - description of, 28
- Status log, 31, 92
 - clearing, 37, 101
 - see also Message log file
- Stopping the software, 37
- storutil command, 161
- symping command, 161
- symsm file, 160
- symsmdisk utility, 162

T

- Tasks, 31
- Troubleshooting, 3
 - see also Installation and Support Guide

U

- Unassigned drive group, 10, 83
 - creating LUNs from, 62
 - viewing, 60
- Upgrading controller firmware, 133
- User guide
 - description of, 2

V

- Vendor
 - drives, 53
- Vendor ID
 - SCSI controllers, 51

W

- What's new in this version, 34
- Write cache mirroring, 155
 - defined, 9

- failures and, 117
- restrictions, 155
- setting, 55, 155
- when to use, 25
- see also Cache options

- Write caching, 155
 - defined, 9
 - failures and, 117
 - setting, 55, 155
 - see also Cache options

