

TPM Installation Instructions and User's Guide for TP9100

007-4382-002

CONTRIBUTORS

Written by Mick Taylor

Illustrated by Mick Taylor

Production by Karen Jacobson

Engineering contributions by Dave Lucas, Sammy Wilborn, Rob Novak.

COPYRIGHT

© 2001 Silicon Graphics, Inc. All rights reserved; provided portions may be copyright in third parties, as indicated elsewhere herein. No permission is granted to copy, distribute, or create derivative works from the contents of this electronic documentation in any manner, in whole or in part, without the prior written permission of Silicon Graphics, Inc.

LIMITED RIGHTS LEGEND

The electronic (software) version of this document was developed at private expense; if acquired under an agreement with the USA government or any contractor thereto, it is acquired as "commercial computer software" subject to the provisions of its applicable license agreement, as specified in (a) 48 CFR 12.212 of the FAR; or, if acquired for Department of Defense units, (b) 48 CFR 227-7202 of the DoD FAR Supplement; or sections succeeding thereto. Contractor/manufacturer is Silicon Graphics, Inc., 1600 Amphitheatre Pkwy 2E, Mountain View, CA 94043-1351.

TRADEMARKS AND ATTRIBUTIONS

Silicon Graphics is a registered trademark and SGI and the SGI logo are trademarks of Silicon Graphics, Inc.

Internet Explorer, Windows 2000, and Windows NT are trademarks of Microsoft Corporation, Netscape is a trademark of Netscape Communications, IRIX is a trademark of SGI, Linux is a trademark of Linus Torvalds and Red Hat is a trademark of Red Hat Inc., AcceleRAID and eXtremeRAID are trademarks of Mylex Corporation, Qlogic is a trademark of Qlogic corporation, IBM is a registered trademark of International Business Machines, Inc., and other names that are trademarks may be used in this document for the purpose of identifying the products or services of their respective owners.

Cover Design By Sarah Bolles, Sarah Bolles Design, and Dany Galgani, SGI Technical Publications.

Record of Revision

Version	Description
001	March 2001 Initial Rev
002	October 2001 Functionality updates

Contents

List of Figures	ix
List of Tables	xiii
About This Guide.	xv
Audience	xv
Structure of this Guide	xv
Product Support	xvi
Obtaining Publications	xvi
Reader Comments	xvi
1. Introduction	1
Product Overview	1
Features	2
Software Packaging	2
IRIX/Linux Software Packaging	2
Windows Software Packaging	3
2. Installing, Configuring, and Running TPM on IRIX, Linux, and Windows	5
Host Prerequisites	5
Installing the Software	7
Installing TPM on IRIX Platforms	7
Uninstalling TPM on IRIX Platforms	8
Installing TPM on Linux (RedHat 6.2 with SGI ProPack 1.3) Platforms	8
Uninstalling TPM on Linux (RedHat 6.2 with SGI ProPack 1.3) Platforms.	9
Installing TPM on Windows NT 4.0 and Windows 2000 Advance Server Platforms	10
Uninstalling TPM on Windows NT or Windows 2000 Advanced Server Platforms	15
Configuring Usernames and Passwords	15

Invoking TPM 16
Setting the Environment Variable for IRIX and Linux: 16
Starting SGI TP9100 Array Manager for IRIX and Linux 16
Starting SGI TP9100 Array Manager for Windows 17
Setting Up the Browser 17
Logon Screen 19
Additional Considerations 20
3. Using TPM 21
Main Menu 21
New Features 24
New Features for 7.75 Firmware Release on the FFX Controller 24
Specific New Functions for 7.75 FFX Firmware/TPM 1.2 25
New View/Modify RAID Controller Configuration Parameters for 7.75 FFX Firmware /TPM 1.2. 25
Configuration Functions 26
Select Physical Device (Configuration Function Menu) 27
Modify Disk Device Status (Configuration Function Menu) 30
Create/Destroy/Expand Logical Drives (Configuration Function Menu) 36
Initialize LUN(s) (Configuration Function Menu). 68
Initialize LUNs in Background (Configuration Function Menu). 73
Configure Host-to-LUN Mapping (Configuration Function Menu) 76
Installing Device Drivers for New LUNS (Configuration Function Menu) 83
Send Command to Host (Configuration Function Menu). 84
Administrative Functions 85
View/Modify RAID Controller Configuration. 86
Reset Controller(s). 95
Gracefully Bring a Controller Off Line 96
Gracefully Bring a 2nd Controller On Line 97
Perform Data Consistency Operations on a LUN 98
Enable/Disable Write Cache for LUN(s) 103
Modify/Purge Non-Volatile WWN Tables 104

Reporting Functions106
Set Default Screen Refresh Rate107
Display (Dual) Controller Status108
Topology Query109
Display Statistical Data by Physical Device110
Display Statistical Data by Logical (RAID) Disk (6.14 Firmware Only)110
Display SCSI/Fibre Device Information111
Display FULL Subsystem Configuration Information112
Display Physical Subsystem Information114
Display Logical Subsystem Information118
Display Environmental Subsystem Information120
Miscellaneous Functions122
Display Status of Background Jobs124
Flush Controller(s) Write Cache to Disk127
Save Current Controller Configuration128
Load Controller Configuration129
Flash New Firmware Onto Controller(s)130
Flash New Firmware Onto Supported Disk(s)131
Adjust Battery Settings134
Clear Configuration136
Set The Real Time Clock (7.01 and Later Firmware Only)137
View the Controller's Internal Event Log (7.01 and later Firmware Only)138
Identify a Disk (7.01 and later Firmware Only)140
Scan for New Enclosures and Disks143
Save Controller Debug Dump145
Stop Configurator Service Routine on Host146
Automatic Restart (Reboot) on Failure Parameter147
A. TPMWatch Event Monitor and Logger149
Running the TPM 1.0 Version of TPMWatch From Linux and IRIX150
Running the TPM 1.2 Version of TPMWatch From IRIX151
Running the TPM 1.0 Version of TPMWatch From Windows153
Master/Slave Controller Status Table155
Partner Status Table157
B. Error Codes161

List of Figures

Figure 2-1	TPM InstallShield Wizard	10
Figure 2-2	License Agreement Dialog Box	11
Figure 2-3	Choose Destination Dialog Box	12
Figure 2-4	Select Program Folder Dialog Box	13
Figure 2-5	Installation Complete Dialog Box.	14
Figure 2-6	Logon Screen	19
Figure 3-1	Main Menu (partial screen)	22
Figure 3-2	Select Physical Device Dialog Box	27
Figure 3-3	Modify Disk Device Status Dialog Box (for 6.14 Firmware)	31
Figure 3-4	Modify Disk Device Status Dialog Box (for 7.01 and Later Firmware) ³²	
Figure 3-5	Create/Destroy/Expand Logical Drives Dialog Box	36
Figure 3-6	New Configuration Screen	38
Figure 3-7	Drives Assigned Into Pack A (6.14 Firmware)	41
Figure 3-8	Drives Assigned Into Pack A (7.01 and Later Firmware)	42
Figure 3-9	All Drive Packs Created (6.14 Firmware)	43
Figure 3-10	All Drive Packs Created (7.01 and Later Firmware)	44
Figure 3-11	Select Drive Packs to Include in LUN	45
Figure 3-12	LUN is Being Defined (6.14 Firmware)	46
Figure 3-13	LUN is Being Defined (7.01 and Later Firmware)	47
Figure 3-14	LUN is Defined	48
Figure 3-15	All LUNs Defined (6.14 Firmware)	49
Figure 3-16	All LUNs Defined (7.01 and later Firmware)	50
Figure 3-17	Updating the RAID Controller Warning Screen (6.14, 7.01, and 7.03 Firmware)	51
Figure 3-18	Updating the RAID Controller Warning Screen (7.75 Firmware).	51
Figure 3-19	LUN Online Screen (6.14, 7.01, and 7.03 firmware only)	52
Figure 3-20	LUN Online Screen (7.75 firmware only)	52

Figure 3-21	Warning Screen (6.14, 7.01, and 7.03 Firmware) 53
Figure 3-22	Controller Resetting Information Screen 53
Figure 3-23	Controller Resetting Status Screen (1 of 2) 54
Figure 3-24	Controller Resetting Status Screen (2 of 2) 54
Figure 3-25	Create/Destroy/Expand Logical Drives Screen (6.14 Firmware) 56
Figure 3-26	Create/Destroy/Expand Logical Drives Screen (7.01 and Later Firmware) 57
Figure 3-27	Add LUN(s) Screen 58
Figure 3-28	Add LUN(s) Warning Screen (6.14, 7.01, and 7.03 Firmware) 59
Figure 3-29	Add LUN(s) Warning Screen (7.75 Firmware) 59
Figure 3-30	Add LUN(s) Confirmation Screen (6.14, 7.01, and 7.03 Firmware) 60
Figure 3-31	Add LUN(s) Confirmation Screen (7.75 Firmware) 60
Figure 3-32	Delete LUN Warning Screen 61
Figure 3-33	Create/Destroy/Expand Logical Drives Screen (6.14 Firmware) 64
Figure 3-34	Create/Destroy/Expand Logical Drives Screen (7.01 and Later Firmware) 65
Figure 3-35	Expand LUN Screen 66
Figure 3-36	Warning Message 67
Figure 3-37	Initialize LUN (6.14, 7.01, and 7.03 Firmware) 68
Figure 3-38	Initialize LUN (7.75 Firmware) 69
Figure 3-39	Initialization Confirmation. 70
Figure 3-40	Initialization Progress (6.14 Firmware) 71
Figure 3-41	Initialization Progress (7.01 and Later Firmware) 72
Figure 3-42	Initialize LUNs in Background (7.75 Firmware). 74
Figure 3-43	Initialize LUNs in Background (7.75 Firmware). 74
Figure 3-44	Configure Host to LUN Mapping (6.14 Firmware) 76
Figure 3-45	Storage Area Network 78
Figure 3-46	SAN Mapping 81
Figure 3-47	Install New Device Drivers Warning Screen 83
Figure 3-48	Send Command to Host 84
Figure 3-49	Administrative Functions Menu 85
Figure 3-50	View/Modify Controller Configuration (partial for 6.14 Firmware) 87
Figure 3-51	View/Modify Controller Configuration (partial for 6.14 Firmware) 88

Figure 3-52	View/Modify Controller Configuration (partial for 7.01 Firmware) .	89
Figure 3-53	View/Modify Controller Configuration (partial for 7.01 Firmware) .	90
Figure 3-54	View/Modify Controller Configuration (partial for 7.75 Firmware) .	91
Figure 3-55	View/Modify Controller Configuration (partial for 7.75 Firmware) .	92
Figure 3-56	View/Modify Controller Configuration (partial for 7.75 Firmware) .	93
Figure 3-57	Reset Controllers Warning Screen	95
Figure 3-58	Controller Booting Window	95
Figure 3-59	Controller Responding Window	96
Figure 3-60	Warning Screen.	96
Figure 3-61	Warning Screen.	97
Figure 3-62	Data Consistency Check/Restore.	98
Figure 3-63	Operation Started Dialog Box.	99
Figure 3-64	Background Initialization Status Screen (6.14 Firmware)	100
Figure 3-65	Background Initialization Status Screen (7.01 and 7.03 Firmware) .	101
Figure 3-66	Background Initialization Status Screen (7.75 Firmware)	102
Figure 3-67	Enable Write Cache For LUN(s)	103
Figure 3-68	World Wide Name Table for RAID Controller	104
Figure 3-69	Reporting Functions Menu (for 6.14 Firmware).	106
Figure 3-70	Reporting Functions Menu (for 7.01 and Later Firmware). . . .	106
Figure 3-71	Default Screen Refresh Rate	107
Figure 3-72	Controller Status Information.	108
Figure 3-73	Topology Query	109
Figure 3-74	Statistical Data for Logical Devices (for 6.14 Firmware)	110
Figure 3-75	Display SCSI/Fibre Device Information	111
Figure 3-76	Controller Information Values (partial screen)	112
Figure 3-77	Controller Tunable Parameters	112
Figure 3-78	GroupConfig Structure	113
Figure 3-79	Physical Subsystem Information (partial for 6.14 Firmware) . . .	114
Figure 3-80	Physical Subsystem Information (partial for 6.14 Firmware) . . .	115
Figure 3-81	Physical Subsystem Information (partial for 7.01 and later Firmware)	116
Figure 3-82	Physical Subsystem Information (partial for 7.01 and later Firmware)	117

Figure 3-83	Logical Subsystem Information (partial for 6.14 Firmware) . . .	118
Figure 3-84	Logical Subsystem Information (partial for 7.01 and Later Firmware)	119
Figure 3-85	Display Environmental Subsystem Information (6.14 Firmware) .	120
Figure 3-86	Display Environmental Subsystem Information (7.01 and Later Firmware)	121
Figure 3-87	Miscellaneous Functions Menu (for 6.14 Firmware)	122
Figure 3-88	Miscellaneous Functions Menu (for 7.01 and Later Firmware) .	122
Figure 3-89	Display Status of Background Jobs (6.14 Firmware)	124
Figure 3-90	Display Status of Background Jobs (7.01 and Later Firmware) .	125
Figure 3-91	Display Status of Background Jobs—Background Init in Progress (7.75 Firmware)	126
Figure 3-92	Cache Flush Warning	127
Figure 3-93	Cache Flush Successful	127
Figure 3-94	Save Current Controller Configuration	128
Figure 3-95	Load Controller Configuration	129
Figure 3-96	Flashing New Firmware	130
Figure 3-97	Flash New Firmware Onto Supported Disk(s) Dialog Box . . .	131
Figure 3-98	Flash New Firmware Dialog Box	132
Figure 3-99	Flash New Firmware Complete Box	133
Figure 3-100	Adjust Battery Settings (6.14, 7.01, and 7.03 Firmware).	134
Figure 3-101	Adjust Battery Settings (7.75 Firmware)	134
Figure 3-102	Clear Configuration.	136
Figure 3-103	Clear Configuration.	136
Figure 3-104	Set Real Time Clock.	137
Figure 3-105	Set Real Time Clock Confirmation Dialog Box.	137
Figure 3-106	View Controller’s Internal Event Log.	138
Figure 3-107	SES Drive Locations (Rackmount Configuration)	140
Figure 3-108	SES Drive Locations (Tower Configuration)	141
Figure 3-109	Identify a Disk	142
Figure 3-110	View Controller’s Internal Event Log (7.75 Firmware)	143
Figure 3-111	Dump Debug Configuration to Disk (7.75 Firmware)	145
Figure 3-112	TPM 1.0 Stop Service Screen	146
Figure 3-113	TPM 1.2 Shutdown Message	146

List of Tables

Table 1-1	Files in the Distribution (IRIX and Linux)	2
Table 1-2	Files in the Distribution (Windows)	3
Table 3-1	Select Physical Device Field Definitions	30
Table 3-2	Modify Disk Device Status Field Definitions	34
Table A-1	Record Layout154
Table A-2	Master/Slave Controller Status Table155
Table A-3	Partner Status Table157
Table B-1	Error Codes162

About This Guide

This guide provides information on preparing, installing, configuring, and using the Total Performance Manager (TPM) application for the TP9100 external RAID subsystem.

Note: At the time of production, an issue with the 7.70 controller firmware resulted in version 7.75 controller firmware code being released. All screen shots that were done with 7.70 controller firmware are still valid for 7.75 controller firmware.

Audience

This guide is intended for system administrators. Use this guide to:

- Gain a basic understanding of the TPM software.
- Learn how to install, configure, and run the TPM software in Linux, IRIX, and Windows environments.
- Learn about hardware and software requirements.
- Learn how to use the GUI to operate the TPM software.

Structure of this Guide

This guide contains the following chapters:

- Chapter 1, “Introduction” — Introduces the TPM software and provides information about features and packaging.
- Chapter 2, “Installing, Configuring, and Running TPM on IRIX, Linux, and Windows” — Describes the host prerequisites and how to install and configure the TPM software.
- Chapter 3, “Using TPM” — Describes how to use the various GUI menus to configure and control the RAID subsystem.
- Appendix A, “TPMWatch Event Monitor and Logger” — Describes how to configure and operate TPMWatch, a support program designed to poll RAID subsystems and report their health to an output file.
- Appendix B, “Error Codes” — Describes the error codes associated with the TPM software.

Product Support

SGI provides a comprehensive product support and maintenance program for its products. If you are in North America and would like assistance with your SGI-supported products, contact the Customer Support Center (CSC) at 1-800-800-4SGI (1-800-800-4744) or your authorized service provider. If you are outside North America, contact the SGI subsidiary or authorized distributor in your country.

Obtaining Publications

To obtain SGI documentation, go to the SGI Technical Publications Library at:

<http://techpubs.sgi.com>

Reader Comments

If you have comments about the technical accuracy, content, or organization of this document, please tell us. Be sure to include the title and document number of the manual with your comments. (Online, the document number is located in the front matter of the manual. In printed manuals, the document number can be found on the back cover.)

You can contact us in any of the following ways:

- Send e-mail to the following address:
techpubs@sgi.com
- Use the Feedback option on the Technical Publications Library World Wide Web page:
<http://techpubs.sgi.com>
- Contact your customer service representative and ask that an incident be filed in the SGI incident tracking system.
- Send mail to the following address:
Technical Publications
SGI
1600 Amphitheatre Pkwy., M/S 535
Mountain View, California 94043-1351
- Send a fax to the attention of "Technical Publications" at +1 650 932 0801.

We value your comments and will respond to them promptly.

Introduction

Product Overview

The Total Performance Manager (TPM) application is an HTML-based GUI RAID client-server application manager. It provides a user-friendly graphical user interface (GUI) for configuring and monitoring external RAID disk subsystems that use several families of RAID controller engines. TPM consists of two programs that run natively on an IRIX, Windows, or Linux system that is physically attached to a RAID storage subsystem. All communication is through the Fibre Channel (FC) controller, so that no RS-232 interface is required. The application incorporates an embedded Web server to provide the user interface to a user-supplied Web browser.

Note: TPM supports only external RAID controllers. It does not support internal RAID controllers, such as the Mylex AcceleRAID cards used in some SGI computing platforms.

Once the TPM service routine is running on a host server, it may be interfaced to any modern Web browser that supports HTML level 3 or higher. This includes Microsoft's Internet Explorer 4.x and above, and Netscape Navigator 3.x and above. The browser client does not have to be located on the computer running the TPM service routine. The browser also does not have to be running locally. If TPM is running on a computer with Internet access, the browser can communicate with TPM through a dial-up connection (provided there is no firewall blockage).

To communicate with the TPM service, enter the URL on which the TPM is listening. By default, TPM starts on port 2002. Therefore, if your host server is configured for IP address 192.168.1.99, set your web browser to the following URL:

`http://192.168.1.99:2002`

Features

TPM is designed to support the TP9100 external RAID subsystem. The TPM application provides the following functions:

- Configuration: Adds, deletes, and modifies LUNs, topologies, and device status.
- Administration: Configures controllers, administers LUNs, and takes controllers online and offline.
- Reporting: Presents statistical data and subsystem status.
- Maintenance Procedures, such as RAID controller and drive firmware downloads.

Software Packaging

The software is packaged differently for IRIX/Linux than for Windows. The following sections describe both.

IRIX/Linux Software Packaging

For IRIX and Linux, the TPM software is distributed on a CD-ROM that contains the files shown in Table 1-1. The files must be located in the `$DAM_HOME` directory tree. By convention, `$DAM_HOME` is set to `/opt/dam`, and this manual uses `$DAM_HOME` and `/opt/dam` interchangeably.

Table 1-1 Files in the Distribution (IRIX and Linux)

Files	Description
<code>cgi-bin/oemparts.txt</code>	Cross-reference file for spoofing make and model of subsystem.
<code>cgi-bin/oemtail.html</code>	HTML segment that appears at the bottom of most Web pages.
<code>tpm</code>	TPM service routine.
<code>dameventlog</code>	Log file (automatically created if not found).

Table 1-1 (continued) Files in the Distribution (IRIX and Linux)

Files	Description
damioscan	Script or executable that discovers SCSI/Fibre Channel devices and LUNs.
tpmwatch	Executable file that monitors subsystem health.
images/	Directory of image files displayed on various Web pages.
images/oemlogo.gif	The logo that is displayed on main screen.
images/wallpaper.gif	The background image displayed on most Web pages.
database/	Directory of files that store inquiry, log, and sense codes.
passwd.txt	The username and password file that is validated during logon.
/tmp	The directory where TPM places several small temporary files during program execution.

Windows Software Packaging

By convention, \$DAM_HOME is set to the current working directory (for example, "C:\Program Files\SGI\SGI TP1900 Array Manager"). The distribution files must be located in the current working directory.

Table 1-2 Files in the Distribution (Windows)

Files	Description
cgi-bin\oemparts.txt	Cross-reference file for spoofing make and model of subsystem.
cgi-bin\oemtail.html	HTML segment that appears at the bottom of most Web pages.
cygwin1.dll	A pseudo UNIX environment for Windows. This DLL provides a translation layer between the UNIX environment and the Windows environment.
database\	Directory files that store inquiry, log, and sense codes.

Table 1-2 (continued) Files in the Distribution (Windows)

Files	Description
<code>dmail.exe</code>	This executable provides the mailing function and is called from within the <code>tpmwatch</code> executable.
<code>images\oemlogo.gif</code>	The logo that is displayed on the main screen.
<code>images\wallpaper.gif</code>	The background image displayed on most Web pages.
<code>passwd.txt</code>	The username and password file that is validated during logon.
<code>See32.dll</code>	This is a support dll for the email program.
<code>sh.exe</code>	This executable is part of the CYGWIN environment and allows usage of shell commands using the 'C' system command
<code>tpm.exe</code>	TPM service routine.
<code>tpmwatch.exe</code>	Executable file that monitors subsystem health.

Installing, Configuring, and Running TPM on IRIX, Linux, and Windows

This chapter explains how to install and configure the TPM application on all operating systems. You may install TPM with other active users on the system if you wish. The entire process takes little time and no reboot is required.

Host Prerequisites

The host server where the TPM application is installed must have the following prerequisites:

- TCP/IP access. This access must exist between the host system and the machine with the Web browser. This can be over any medium, including Ethernet, token ring, ATM, or dial-up SLIP/PPP.
- A specific IP port or socket number. The software must communicate with a specific IP port or socket number. If there is a firewall or router, ensure the administrator does not restrict traffic over that socket (normally, the socket used is 2002).
- A compatible Web browser. HTML-compatible Web browsers with JavaScript support, such as Microsoft's Internet Explorer (IE) version 4.0 or 5.0, and Netscape's 4.x browsers have been tested. The browser can execute on any machine. The operating system of the client machine is not important.
- Your fibre channel host adapter and drivers must be properly configured.
- Operating System. TPM supports the following operating system versions:
 - IRIX 6.5.8 or later
 - Windows NT 4.0 with Service Pack 6 or later
 - Windows 2000 Advanced Server with Service Pack 1 or later.
 - Linux RedHat 6.2 with ProPack 1.3

Please note the following information on operating system support:

- TP9100 support for Linux has been frozen at Red Hat release 6.2 with Pro Pack 1.3, using TPM 1.0 with controller firmware 6.14, 7.01 and 7.03. Support for the TP9100 on newer releases of Linux will not be developed.
- TP9100 support for Windows has been frozen at Windows 2000 Advance Server and NT 4.0 with service pack 6, using TPM 1.0 and controller firmware 6.14, 7.01 and 7.03. Support for the TP9100 on newer releases of Windows will not be developed.
- Linux and Windows directories are provided on the TPM1.2 CD under the unsupported directory. The images provided have not been rigorously tested for reliable operation with all TP9100 configurations or features; therefore they are available under the following two conditions:
 - Operation is not guaranteed
 - Support will not be offered

Installing the Software

This section describes how to install and uninstall the TPM software on the IRIX, Linux, and Windows platforms.

Note: Prior to removing or upgrading the software, Tpm and TPMWatch must first be terminated.

Installing TPM on IRIX Platforms

Follow this procedure to install TPM on IRIX platforms:

1. Log in as root.
2. Insert the TPM CD-ROM into the CD-ROM drive.

Note: If the CD-ROM does not mount, refer to the appropriate IRIX Operating System Manual for instructions.

3. Remove the existing TPM 1.0 software from your system. Type the following command to launch the IRIX software versions tool to remove the `sgi_tpm` software image:

```
versions remove sgi_tpm
```
4. Type the following command to launch the IRIX installation tool (`inst`) to install the `sgi_tpm_1.2` software image:

```
inst -f /CDROM/irix/dist/sgi_tpm
```
5. To specify the package, type:

```
list
```
6. To install the software, type:

```
install
```
7. At the Install subsystem prompt, type:

```
sgi_tpm
```

8. To complete the installation, type:
`go`
9. Type the following command to exit the install program:
`quit`
The exit operation is automatically performed.
10. Type the following command to unmount the CD:
`umount /CDROM`
11. Remove the CD from the CD-ROM drive.

Uninstalling TPM on IRIX Platforms

1. Log in as root.
2. Type the following command to launch the IRIX `versions -remove`, to remove TPM software:
`versions remove sgi_tpm`

Installing TPM on Linux (RedHat 6.2 with SGI ProPack 1.3¹) Platforms

1. Log in as root.
2. Insert the TPM CD-ROM into the CD-ROM drive.
3. Type the following command to mount the CD:
`mount /dev/cdrom /mnt/cdrom`
4. Type the following command to launch Red Hat Package Manager (rpm), which installs the TPM software:
`rpm -iv /mnt/cdrom/unsupported/linux/rpm/sgi_tpm.rpm`

¹ At the time of this writing, TP9100 will not operate with SGI ProPack 1.4. SGI ProPack 1.3 works properly with drive firmware 6.14; however, drive downloads to update firmware do not work.

5. To unmount the CD, type the following command:

```
umount /mnt/cdrom
```
6. Remove the CD from the CD-ROM drive.
7. Make sure of the following:
 - a. All SCSI adapters are loaded.
 - b. The qllogicfc module is removed.
 - c. The qllogic2100 module is loaded.
 - d. The SCSI generic module is loaded.

Uninstalling TPM on Linux (RedHat 6.2 with SGI ProPack 1.3²) Platforms

1. Log in as root.
2. Type the following Red Hat Package Manager (rpm) command to uninstall the TPM software:

```
rpm -ev sgi_tpm
```

² At the time of this writing, TP9100 will not operate with SGI ProPack 1.4. SGI ProPack 1.3 works properly with drive firmware 6.14; however, drive downloads to update firmware do not work.

Installing TPM on Windows NT 4.0 and Windows 2000 Advance Server Platforms

1. Log in as administrator.
2. Insert the TPM software CD-ROM into the CD-ROM drive.
3. Select Start →Run and type the following in the Run dialog box:

`X:\unsupported\windows\setup.exe`

where X: corresponds to the drive letter of your CD-ROM drive, OR if you prefer, navigate to the \unsupported\windows folder on the CD-ROM and double click setup.exe.

The TPM InstallShield Wizard dialog box appears, as shown in Figure 2-1.

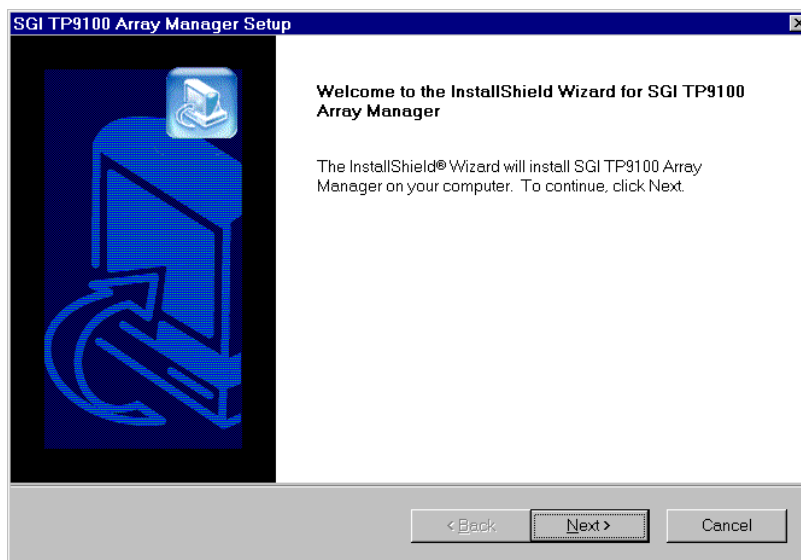


Figure 2-1 TPM InstallShield Wizard

4. Click Next.

The License Agreement dialog box appears, as shown in Figure 2-2.

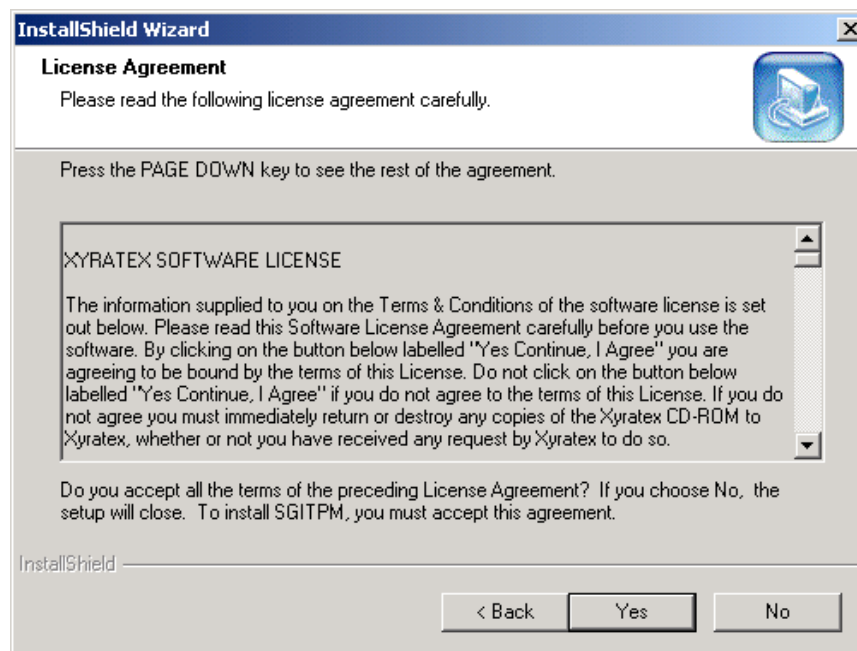


Figure 2-2 License Agreement Dialog Box

5. Click **Yes**.

The Choose Destination Location dialog box appears, as shown in Figure 2-3.

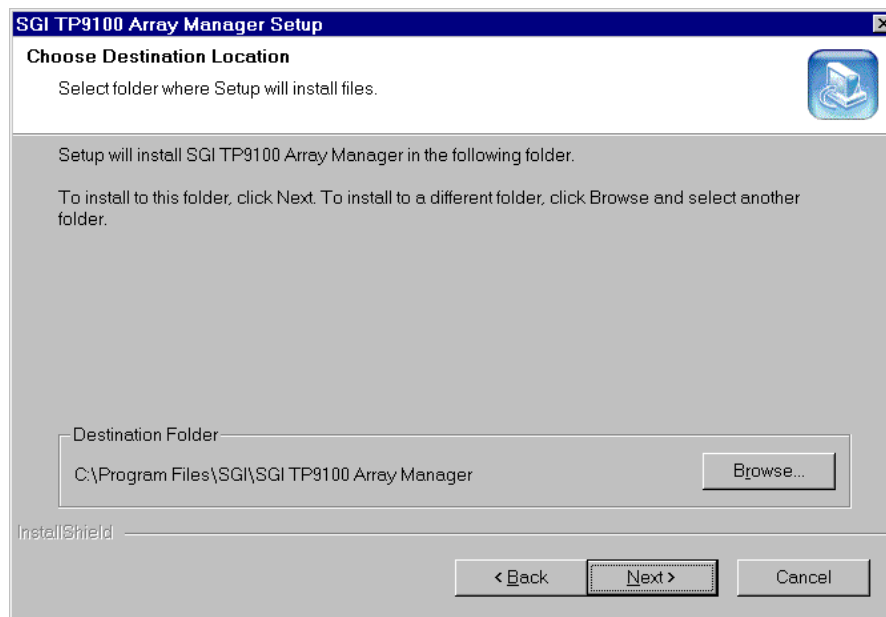


Figure 2-3 Choose Destination Dialog Box

6. You can accept the file installation default destination folder (C:\Program Files\SGI\SGI TP9100 Array Manager) or click **Browse...** to select a file installation different destination folder.
7. When you have selected a destination folder, click **Next**.
The Select Program Folder dialog box appears, as shown in Figure 2-4.

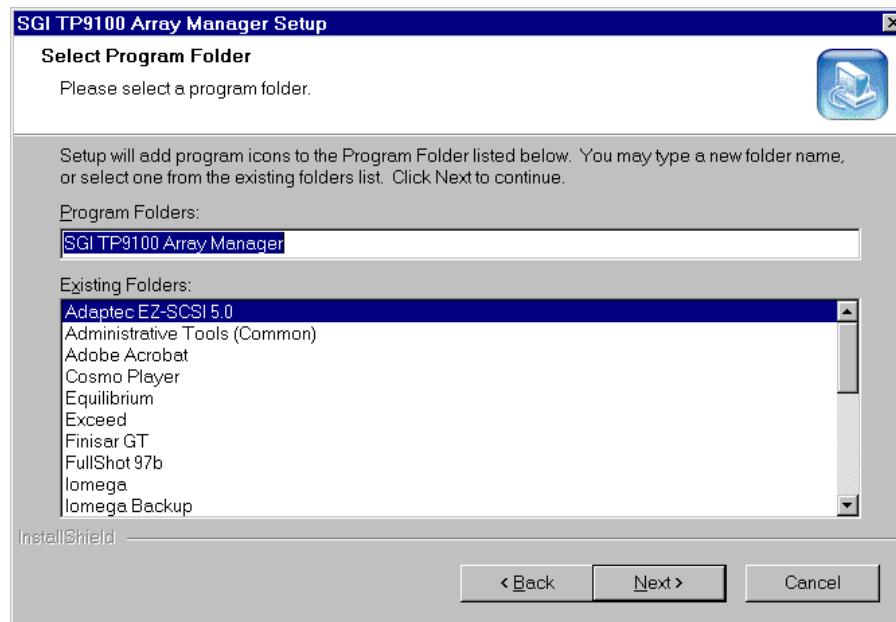


Figure 2-4 Select Program Folder Dialog Box

8. Select the program folder for storing program icons and click **Next**.
The Installation Complete dialog box appears, as shown in Figure 2-5.

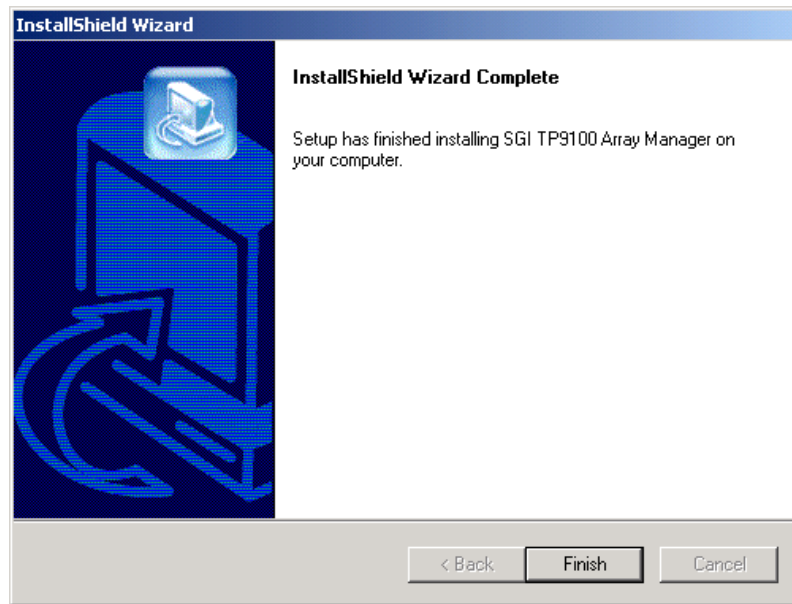


Figure 2-5 Installation Complete Dialog Box

9. Click **Finish** to complete the installation.

Uninstalling TPM on Windows NT or Windows 2000 Advanced Server Platforms

To perform the uninstall, follow the procedure listed here.

1. Go to Start -> Settings -> Control Panel and double-click **Add/Remove Programs**.
2. In the list that appears, select the SGI TP9100 Array Manager software entry and click **Change/Remove**.
3. Click **Yes in the Confirm File Deletion** dialog box.
A dialog box appears that indicates that the SGI TP9100 Array Manager software has been successfully uninstalled from your machine.
4. Click **OK**.
The TPM software is removed.
5. Click **Close** on the Add/Remove Programs Properties dialog box to close it.
6. Close the Control Panel window.

Configuring Usernames and Passwords

TPM's security mechanism is simple. Just edit the `passwd.txt` file, and make changes as necessary (only the system administrator should have access rights to this file). Each line serves as new username and password combination. The username and password are separated with a single colon (username:password). The colon character is not allowed in either the username or password, everything is case-sensitive, and the username, password, or both fields can be blank.

Invoking TPM

This section explains how to invoke TPM after installing it.

Setting the Environment Variable for IRIX and Linux:

To set the environment variable, follow these steps:

1. Log on as root (or have root access).
2. For csh, or tcsh SHELL, type:

```
setenv DAM_HOME /opt/dam
```

3. For sh, bash, or ksh SHELL, type:

```
DAM_HOME=/opt/dam
```

```
export DAM_HOME
```

Starting SGI TP9100 Array Manager for IRIX and Linux

Note: Only users with root permissions are allowed to run TPM.

To start TPM, enter the following command:

```
/opt/dam/tpm [-D][ -W port_number ]
```

where `port_number` is the IP socket/port number required to access the program. This not only hides TPM from standard Web surfers, but also prevents it from interfering with a system that happens to function as a Web server. Because normal Web traffic uses TCP/IP port number 80, the TPM server is invisible to client browsers unless the browsers are instructed to use a specific port.

By default, TPM interacts with port number 2002. If that port is busy, the program just initializes the service at the next available port number. It is also recommend that you run the service routine in the background to avoid tying up a terminal session. To do this, enter the following command:

```
# /opt/dam/tpm &
```


The flag `-D` can be added to display debug information on the terminal. Only do this if so instructed.

Note: When enabled, debug mode may cause the application to run more slowly.

Starting SGI TP9100 Array Manager for Windows

To start TPM, use one of the following methods:

- Start → Programs → SGI TP9100 Array Manager, or
- Double-click the SGI TP9100 Array Manager icon on the desktop.

To turn on debug mode, right-click the SGI TP9100 Array Manager icon on the desktop and click Properties.

On the resulting dialog box, add `-D` to the end of the entry in the Target text box (after the right-hand quote mark).

Note: When enabled, debug mode may cause the application to run more slowly.

Setting Up the Browser

Now that TPM has been invoked, you must set up a browser to run the TPM GUI.

Assume, for example, that the TCP/IP number of your host is 192.200.200.7, and it's name is `server1.sgi.com`. To interact with the TPM software, set the address (URL) of your browser to one of the following:

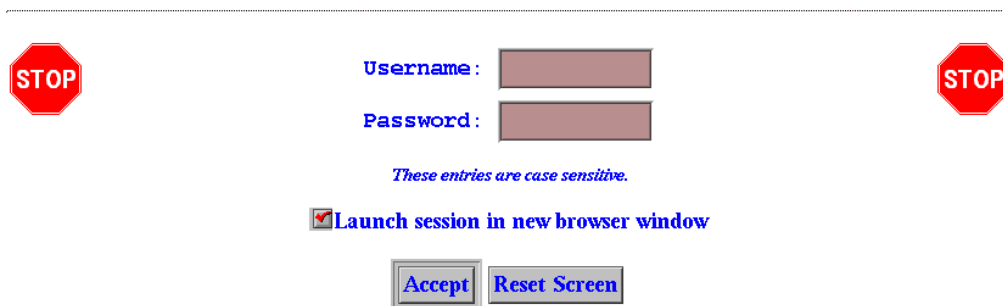
- `http://server1.sgi.com:2002`, or
- `http://192.200.200.7:2002`

If you have another server running the TPM application, and the IP name/number combination is 192.200.200.10, `server2.sgi.com`, and you started the program with `tpm -w 1234`, then you may access the software from either another PC, or an additional browser window, with:

- `http://server2.sgi.com:1234`, or
- `http://192.200.200.10:1234`

Logon Screen

After you invoke TPM, the logon screen appears in the browser window, as shown in Figure 2-6.



Username :

Password :

These entries are case sensitive.

Launch session in new browser window

Figure 2-6 Logon Screen

Follow these steps to log on:

1. Enter your username.
2. Enter your password.
3. Click **Accept**.

The username and password entered are validated against the file `$DAM_HOME/passwd.txt`. If the username and passwords match the file contents, the main window screen appears (see Figure 3-1 on page 22). Otherwise, the logon dialog box is redisplayed. TPM does not support multiple concurrent users. If another user (or specifically, another browser session, from any IP number) invokes TPM by setting their browser to the appropriate URL, TPM logs off the original user.

If you do not have networking installed, you can still access the GUI by using Netscape or Internet Explorer. Just set the browser to `http://localhost:2002`, or the appropriate port number.

Additional Considerations

Below is a list of additional considerations.

- Firewall administrators may block traffic on undefined port numbers such as 2002. Please talk with your security administrator to make sure you agree on what ports are acceptable for running TPM.
- Once the TPM job ends, it typically takes a minute or so for the port to automatically free up on your operating system. That means if you start another session of TPM before the port is free, you will see a message saying your default port is busy, and it will use the next available one.
- TPM is not designed to be a multi-user program. Only one Web browser at a time should attempt to interact with it.
- If you have an enterprise with multiple hosts and subsystems, you can open multiple windows as necessary with your browser to interact with an unlimited number of subsystems concurrently.
- Each record must be a fixed length (16 + 18 + 1) bytes long. The last byte is the new line character which will be added by your text editor. Do not edit this file on a PC. Be careful when FTP'ing, as well.

Using TPM

This chapter explains how to use TPM to configure and monitor your external RAID disk subsystems.

Note: The GUI screens vary according to the firmware (6.14, 7.01, 7.03, or 7.75) installed on the RAID controller. Where the screens are different, both are presented and explained.

Main Menu

After you successfully log on to TPM, the main menu is presented (see Figure 3-1).

Note: In many of the menus and screens shown in this document and presented by the software, additional information in the menu or screen is shown that may not appear in the body of the document. Be sure to read all information in each menu or screen before taking action on a particular menu.

Current Device: DAC960FFx-based RAID Controller at `/fw/scsi/20000080e5115f52/lun0/c3p1`

Configuration Functions:

- **Select physical device** - displays a table of all SCSI/Fibre devices and RAID subsystems. You must select the one which you wish to configure or inquire about (allow up to several minutes for discovery).
- **Modify disk device status** - bring disks on-line, off-line, or assign as global spares.
- **Create/Destroy/Expand Logical Drives** - allows you to set up new configurations, add logical drives to existing configurations, and expand the capacities of existing arrays.
- **Configure Host-to-LUN Mapping** - This allows logical drives to be made available or invisible to selected FC host adapters.
- **Initialize LUN(s)** - performs low-level format of a LUN. This is required before it can be used by your O/S.
- **Install device drivers for new LUNs** - This instructs your operating system to scan for new LUNs, and create device drivers as necessary.
- **Initialize LUNs in background** - This performs low-level format of all uninitialized LUNs in the background.
- **Send Command to host** - Use this option to send commands to your host. The commands must not be interactive, and all results will be displayed upon completion.

Administrative Functions:

- **View/Modify RAID controller configuration** - Although many changes may be made on-the-fly, some settings **must** be made before any RAID groups are defined.
- **Reset controller(s)** - This simultaneously cold resets all controllers in a subsystem. Multiple attempts are made for 90 seconds, in case the controllers have active I/Os.
- **Gracefully bring a controller off-line**.
- **Gracefully bring a 2nd controller on-line** - Do this after a failed controller has been removed, or you are upgrading from a simplex configuration to a dual-controller configuration.
- **Perform data consistency operations on a LUN** - You should perform a data consistency check regularly for all redundant LUNs.
- **Enable/Disable write cache for LUN(s)**.
- **Modify/Purge Non-Volatile WWN Tables** - Allows you to selectively add or remove WWN host entries for LUN Mapping.

Reporting Functions:

- **Set default screen refresh rate** - This lets you define the number of seconds between each screen refresh for status screens which automatically repaint.
- **Display (Dual) Controller Status** - This returns status information on dual controller status, and host addressing information on the connected controller.
- **Topology query** - Displays all host adapters on the SAN attached to the subsystem, and what controller/ports they are attached to.
- **Display statistical data by physical device** - This shows log page information for an individual disk drive.
- **Display SCSI/Fibre device information** - This issues a standard SCSI Inquiry, and reports all fields which describe the device.
- **Display FULL subsystem configuration information** - This is a complete hexadecimal dump of the controller's configuration data structures, and contains information which may be of interest to your supplier in the event of a problem.
- **Display physical subsystem information** - displays drive status, statistics, errors and physical locations for all disks in a subsystem, including expansion units.
- **Display logical subsystem information** - displays RAID groups status, statistics, errors and logical configuration for all RAID groups in a subsystem, including expansion units.
- **Display environmental subsystem information** - displays power, fans, temperature, battery backup, and other data relating to the chassis, including expansion units.

Figure 3-1 Main Menu (partial screen)

Figure 3-1 shows a portion of the main menu window that appears once you log on. The main menu windows are different for 6.14, 7.01, 7.03, and 7.75 firmware. The differences will be explained as each menu item is explained in subsequent sections of this chapter.

It is a good idea to disable the browser menu buttons when you run TPM because the BACK, RELOAD (Netscape Navigator), and REFRESH (Internet Explorer) buttons do not work with TPM. In fact, using these buttons may put the TPM application in an

undesired state. Instead, use the buttons and links that TPM presents at the bottom of each screen.

The main menu window has the following principal areas, with the associated explanations on the indicated pages:

- “Configuration Functions” on page 26
- “Administrative Functions” on page 85
- “Reporting Functions” on page 106
- “Miscellaneous Functions” on page 122

New Features

This section lists the new firmware and TPM features.

New Features for 7.75 Firmware Release on the FFX Controller

1. Reboot On Crash

This parameter controls the behavior of the Automatic Restart on Failure feature, and sets the following values:

- Maximum number of times a controller attempts automatic restart to recover from firmware detected errors
- Minimum time interval that the controller must operate before refreshing the number of restart attempts.

When a controller reaches the maximum restart attempts, the automatic restart feature becomes disabled until the value is refreshed. Any subsequent firmware detected errors require manual intervention to recover the controller. Possible values for the maximum automatic restart attempts are 0 to 15 attempts (the default setting is 3). Possible values for the time interval between refreshing the number of restart attempts are DISABLED (never reset automatically) to 7 days.

2. Debug Dump

Debug Dump is enabled by default. When enabled, this feature records controller state information when an abort occurs. After the abort has completed, the abort information can be retrieved and analyzed to help determine why the abort occurred. The information is generated while a controller abort is in progress and recorded to NVRAM and a disk drive. TPM 1.2 would then be used to retrieve the debug dump.

3. Hot Spare Polling

To increase data availability, the controller periodically reads and writes to online spare drives to make sure they are operational. The polling period is not user definable and set to once per 24 hours.

4. Instant RAID Availability (Background Initialization)

Background initialization makes the system drive instantly available for host read and write access. This parameter is enabled by default.

5. SAN Mapping Enhancement

SAN mapping tables store up to 64 host WWNs. With 7.75 firmware, the WWN table was modified to enable the user to remove or delete unused host WWNs. The user can determine which host WWNs are obsolete and need to be removed from the WWN table. As WWNs are removed and the WWN table updated, any WWNs following those deleted are moved up to fill the vacancies in the WWN table.

6. No Controller Reset Requirement on LUN Addition or Deletion

Controller reset is no longer required when deleting or creating new system drives or configurations.

7. No Controller Reset Requirement on Enclosure Addition

This feature allows users to add disk enclosures (one or more) to a configured system while the system continues to operate. After the enclosure or enclosures have been added to the system, this feature starts the SES monitoring process for the new enclosure. The user can then configure the additional disk capacity without restarting the system.

Specific New Functions for 7.75 FFX Firmware/TPM 1.2

- Configuration Functions: Initialize LUNs in Background.
- Administrative Functions: Modify/Purge Non-Volatile WWN Tables.
- Reporting Functions: No new functions.
- Miscellaneous Functions: Scan for New Enclosures and Disks, and Save Controller Debug Dump.

New View/Modify RAID Controller Configuration Parameters for 7.75 FFX Firmware /TPM 1.2

- Debug Dump
- Reboot on Failure (ROF) Rearm Interval
- ROF Reboot Count
- Enable Background Initialization

Configuration Functions

The Configuration Function menu has the following selections, with the associated explanations on the indicated pages:

- “Select Physical Device (Configuration Function Menu)” on page 27
- “Modify Disk Device Status (Configuration Function Menu)” on page 30
- “Create/Destroy/Expand Logical Drives (Configuration Function Menu)” on page 36
- “Initialize LUN(s) (Configuration Function Menu)” on page 68
- “Initialize LUNs in Background (Configuration Function Menu)” on page 73

Note: Initializing LUNs in the background is a new feature with the 7.75 firmware. With the 7.75 firmware, there is no controller reset on LUN creation and deletion.

- “Configure Host-to-LUN Mapping (Configuration Function Menu)” on page 76
- “Installing Device Drivers for New LUNS (Configuration Function Menu)” on page 83
- “Send Command to Host (Configuration Function Menu)” on page 84

Select Physical Device (Configuration Function Menu)

To view a table of all devices and RAID subsystems, click the **Select physical devices** link under **Configuration Functions** in the main menu. The dialog box shown in Figure 3-2 appears.

Select a single device below to inquire about or configure:

- Select a device by clicking the Select button below, then clicking on the **Select** button.
- The list of units below represents all SCSI, Fibre Channel, and SSA peripherals found on your system at this instant. Although you may only *configure* a RAID system, you are free to select any device and perform applicable statistical inquiries on it.
- If you want to select a RAID controller, choose any device that doesn't have the word **None** in the RAID Controller column, preferably at LUN0

Select	Physical device path	Channel	ID	LUN	Type	Vendor ID (SCSI Inquiry)	Product ID (SCSI Inquiry)	Controller(Slot#) Address	Partner Controller(Slot#) Address
<input checked="" type="checkbox"/>	/hw/scsi/20000080e5110413/lun0/c5				RAID	MYLEX	DACARMRB10498B6	DAC960FFx (0) 20-00-00-80-E5-11-04-13	DAC960FFx (1) Empty
<input type="checkbox"/>	/hw/scsi/20000080e5114c0d/lun0/c5				RAID	MYLEX	DACARMRB201366B5	DAC960FFx (0) 20-00-00-80-E5-11-4C-0D	DAC960FFx (1) 20-00-00-00-20-00-00-80

[\[Return to Main Page\]](#)
[\[Log Off\]](#)

Figure 3-2 Select Physical Device Dialog Box

By activating the desired checkbox under the Select column and clicking the **Select** button, you can select which RAID subsystem you wish to configure or monitor. You make your selection by choosing the physical device path associated with any ID/LUN combination displayed.

If you are running in a dual-controller configuration (in redundant mode), make sure to select controller 0 (C0). An example of this is DAC960FFx(0).

Note: All configuration and monitor operations must be through controller 0.

Operation

It is important for the administrator to know how physical device selection works, because it can serve as a good general debugging tool in the event that your computer does not “see” a specific LUN. The algorithm is the same, regardless of the operating system. In summary, TPM executes the following steps:

1. Creates a list of all SCSI and FC device drivers.
2. Issues a standard SCSI inquiry command to report the drive Vendor ID and Product ID fields. If the inquiry fails, TPM assumes the device driver is no good, and skips to the next driver in the list.
3. Issues the vendor-specific Inquiry command to determine if the device is a logical drive within a RAID subsystem. If so, it issues additional commands to report which controller and World Wide Name (WWN) is associated with that LUN.
4. Builds the record and reports what it has discovered if either the Vendor or Product fields are non-blank.

What to Do if a Device is Missing

If a device does not display, it is probably because the device driver is either missing or incorrect. A device may be missing, depending on what OS you have, and what you did to create drivers in the first place. Assuming your FC host adapter is properly installed and operational, and you have exclusive access to your host server, choose from the following:

- If running IRIX, issue the `scsiha -p bus#` command (see `scsiha(lm)`), followed by the `ioconfig -f /hw` command (see `ioconfig(lm)`).
- If running Linux, try removing and installing the device driver. If you have the Qlogic device driver, the command is `/sbin/rmmod qla2100; /sbin/insmod qla2100`. If that doesn't work, you must reboot.
- If running Windows, you must reboot.

Caution: If the Qlogic device driver is in use when you attempt to remove it (for example, an outstanding command is in progress), the host server might crash (the Linux OS will hang). Therefore, you should make sure that all applications that are using the Qlogic device driver have been terminated first before you remove and reinstall the device driver.

Note: Use the `qla2100` driver for SGI ProPack1.3. The **TP9100 cannot be used with ProPack 1.4.**

If all of the above remedies fail, a device might be masked because the controller is doing it intentionally. This occurs if the Affinity, LUN, or SAN mapping is used to make one or more LUNs invisible to a particular host adapter or controller. If you can, go to the main TPM menu (see Figure 3-1) and use the **Configure Host-to-LUN Mapping** selection under Configuration Functions to see if that is the problem. Otherwise, you may have to contact your SGI customer support representative to resolve the problem.

Field Definitions

The definition of each field in the **Select Physical Device** dialog box is given in Table 3-1.

Table 3-1 Select Physical Device Field Definitions

Field	Definition
Select	Selects which RAID subsystem you wish to configure or monitor.
Physical Device Path	In the case of Linux, the physical device path is the pass-through device driver name (<code>/dev/sgx</code>) for that particular device. Otherwise, it is the raw driver.
Controller, Channel, ID, LUN	These are additional fields reported by the OS that help identify the device driver.
Type	TPM reports all TP9100 RAID devices.
Vendor ID (SCSI Inquiry), Product ID (SCSI Inquiry)	The Vendor ID and Product ID parameters are returned by a standard SCSI Inquiry. You may wish to note some interesting information reported by LUNs. The first field typically starts with DAC. The second field describes the type and size of RAID LUN you have. For example, the selected LUN at <code>/hw/scsi/2000080e5110413/lun0/c5p1</code> (first row of Figure 3-2) points to a points to a 104,986MB RAID-0+1 system disk. The digits before the “ B ” indicate number of MB, and the character after the “ B ” indicates the type of RAID. RAID types of 0,1,3, and 5 indicate RAID-0, RAID-1, RAID-3, and RAID-5, respectively. A RAID type of 6 indicates RAID 0+1, and a RAID type of 7 indicates JBOD.
Controller (Slots) Address	This parameter shows the model of the controller, followed by the slot number in the subsystem in parentheses, followed by the unique MAC address for that controller, which assigned by the controller manufacturer.

Modify Disk Device Status (Configuration Function Menu)

To view or change the status of the drives, select **Modify Disk Device Status** under the **Configuration Function** menu. The dialog box shown in Figure 3-3 (for firmware 6.14) or Figure 3-4 (for firmware 7.01 and later) appears.

Modify Disk Device Status for RAID Subsystem at /hw/scsi/20000080e5110413/lun0/c5p1

1. You may change any number of drives that are not ON-LINE to **STANDBY**, **UNCONFIGURED**, **OFF-LINE**, or **ON-LINE**.
2. Allow 10 seconds for the drive state to be changed.
3. All other drives are defined as belonging to a RAID group. You must first delete the RAID group before you can change the state of these disks. If the RAID group is mounted, you must also dismount it, or your host may become confused.

Enclosure #0 (Rackmount view, rotate 90° clockwise for tower)

SGI ST173404FC [2706] S/N: 3CE02LFO 100Mhz, 2 Ports, FC-AL Physical: 70007 Useable: 69991 LoopID=125 (7Dh) Chan=0 ID=0 LUNs: 0	SGI ST173404FC [2705] S/N: 3CE02C6C 100Mhz, 2 Ports, FC-AL Physical: 70007 Useable: 69991 LoopID=124 (7Ch) Chan=1 ID=0 LUNs: 3	SGI ST173404FC [2705] S/N: 3CE02LS6 100Mhz, 2 Ports, FC-AL Physical: 70007 Useable: 69991 LoopID=123 (7Bh) Chan=0 ID=1 LUNs: 0	SGI ST173404FC [2705] S/N: 3CE02J9C 100Mhz, 2 Ports, FC-AL Physical: 70007 Useable: 69991 LoopID=122 (7Ah) Chan=1 ID=1 LUNs: 0
<input checked="" type="radio"/> ON-LINE <input type="radio"/> HOT SPARE <input type="radio"/> UNCONFIGURED	<input checked="" type="radio"/> ON-LINE <input type="radio"/> HOT SPARE <input type="radio"/> UNCONFIGURED	<input checked="" type="radio"/> ON-LINE <input type="radio"/> HOT SPARE <input type="radio"/> UNCONFIGURED	<input checked="" type="radio"/> ON-LINE <input type="radio"/> HOT SPARE <input type="radio"/> UNCONFIGURED
SGI ST173404FC [2705] S/N: 3CE02JP4 100Mhz, 2 Ports, FC-AL Physical: 70007 Useable: 69991 LoopID=121 (79h) Chan=0 ID=2 LUNs: 2	SGI ST173404FC [2705] S/N: 3CE02LRQ 100Mhz, 2 Ports, FC-AL Physical: 70007 Useable: 69991 LoopID=120 (78h) Chan=1 ID=2 LUNs: 1	SGI ST173404FC [2705] S/N: 3CE021HE 100Mhz, 2 Ports, FC-AL Physical: 70007 Useable: 69991 LoopID=119 (77h) Chan=0 ID=3 LUNs: 1	SGI ST173404FC [2705] S/N: 3CE02F1C 100Mhz, 2 Ports, FC-AL Physical: 70007 Useable: 69991 LoopID=118 (76h) Chan=1 ID=3 LUNs: 3
<input checked="" type="radio"/> ON-LINE <input type="radio"/> HOT SPARE <input type="radio"/> UNCONFIGURED	<input checked="" type="radio"/> ON-LINE <input type="radio"/> HOT SPARE <input type="radio"/> UNCONFIGURED	<input checked="" type="radio"/> ON-LINE <input type="radio"/> HOT SPARE <input type="radio"/> UNCONFIGURED	<input checked="" type="radio"/> ON-LINE <input type="radio"/> HOT SPARE <input type="radio"/> UNCONFIGURED

Legend:

VendorID ProductID [Firmware Revision]
 Clock Speed, # of Ports, InterfaceType
 Device Size in Megabytes
 Usable Size in Megabytes
 LoopID Decimal (HEX) Channel TargetID
 Lists of LUNs used by disk
 Current Status

1 MB = 1,048,576 Bytes, 1 GB = 1,073,741,824 Bytes, 1 Block = 512 Bytes

OK Cancel Reset Screen

Figure 3-3 Modify Disk Device Status Dialog Box (for 6.14 Firmware)

Modify Disk Device Status for RAID Subsystem at /hw/scsi/sc13d010

1. You may change any number of drives that are not ON-LINE to **STANDBY**, **UNCONFIGURED**, **OFF-LINE**, or **ON-LINE**.
2. Allow 10 seconds for the drive state to be changed.
3. All other drives are defined as belonging to a RAID group. You must first delete the RAID group before you can change the state of these disks. If the RAID group is mounted, you must also dismount it, or your host may become confused.

Enclosure #0 (Rackmount view, rotate 90° clockwise if in tower)

<p>SGI ST318304FC [2706] S/N: 3EL0098Q WUN: 2000002037653467 100Mhz, 2 Ports, FC-AL Drive Speed: 10016 RPM Physical: 17560 Useable: 17530 LoopID=125 (7Dh) Chan=0 ID=125 LUNs: None</p> <p><input checked="" type="radio"/> ON-LINE <input type="radio"/> HOT SPARE <input type="radio"/> UNCONFIGURED</p>	<p>SGI ST336704FC [2705] S/N: 3CD01S6L WUN: 20000020372A17B9 100Mhz, 2 Ports, FC-AL Drive Speed: 10016 RPM Physical: 35003 Useable: 34696 LoopID=124 (7Ch) Chan=1 ID=124 LUNs: None</p> <p><input checked="" type="radio"/> ON-LINE <input type="radio"/> HOT SPARE <input type="radio"/> UNCONFIGURED</p>	<p>SGI ST318304FC [2705] S/N: 3EL00W4T WUN: 2000002037659735 100Mhz, 2 Ports, FC-AL Drive Speed: 10016 RPM Physical: 17560 Useable: 17530 LoopID=123 (7Bh) Chan=0 ID=123 LUNs: None</p> <p><input checked="" type="radio"/> ON-LINE <input type="radio"/> HOT SPARE <input type="radio"/> UNCONFIGURED</p>	<p>SGI ST318304FC [2705] S/N: 3EL00W1C WUN: 20000020376598A1 100Mhz, 2 Ports, FC-AL Drive Speed: 10016 RPM Physical: 17560 Useable: 17530 LoopID=122 (7Ah) Chan=1 ID=122 LUNs: None</p> <p><input checked="" type="radio"/> ON-LINE <input type="radio"/> HOT SPARE <input type="radio"/> UNCONFIGURED</p>
<p>SGI ST318304FC [2705] S/N: 3EL00HNG WUN: 2000002037658ECD 100Mhz, 2 Ports, FC-AL Drive Speed: 10016 RPM Physical: 17560 Useable: 17530 LoopID=121 (79h) Chan=0 ID=121 LUNs: None</p> <p><input checked="" type="radio"/> ON-LINE <input type="radio"/> HOT SPARE <input type="radio"/> UNCONFIGURED</p>	<p>SGI ST318304FC [2705] S/N: 3EL00KDG WUN: 200000203765985D 100Mhz, 2 Ports, FC-AL Drive Speed: 10016 RPM Physical: 17560 Useable: 17530 LoopID=120 (78h) Chan=1 ID=120 LUNs: None</p> <p><input checked="" type="radio"/> ON-LINE <input type="radio"/> HOT SPARE <input type="radio"/> UNCONFIGURED</p>	<p>SGI ST318304FC [2705] S/N: 3EL00W88 WUN: 200000203765972F 100Mhz, 2 Ports, FC-AL Drive Speed: 10016 RPM Physical: 17560 Useable: 17530 LoopID=119 (77h) Chan=0 ID=119 LUNs: None</p> <p><input checked="" type="radio"/> ON-LINE <input type="radio"/> HOT SPARE <input type="radio"/> UNCONFIGURED</p>	<p>SGI ST318304FC [2705] S/N: 3EL00XDT WUN: 2000002037658EB6 100Mhz, 2 Ports, FC-AL Drive Speed: 10016 RPM Physical: 17560 Useable: 17530 LoopID=118 (76h) Chan=1 ID=118 LUNs: None</p> <p><input checked="" type="radio"/> ON-LINE <input type="radio"/> HOT SPARE <input type="radio"/> UNCONFIGURED</p>

Legends

VendorID ProductID [Firmware Revision]
World Wide Name
Clock Speed, # of Ports, InterfaceType
Device Size in Megabytes
Useable Size in Megabytes
LoopID Decimal (HEX) Channel TargetID
Lists of LUNs used by disk
Current Status

1 MB = 1,048,576 Bytes, 1 GB = 1,073,741,824 Bytes, 1 Block = 512 Bytes

OK Cancel Reset Screen

Figure 3-4 Modify Disk Device Status Dialog Box (for 7.01 and Later Firmware)

The Status dialog box for controllers with 7.01 and later firmware has the following additional information not displayed for the controllers with 6.14 firmware:

- Drive Speed (in RPM)
- WWN (worldwide name)

The **Modify Disk Device Status** dialog box allows you not only quickly view status of the drives, but also to define hot spares. Each drive in Figure 3-3 or Figure 3-4 maps to the same physical row and column of the disk drive chassis. If you are also using expansion enclosures, additional drive matrices are displayed for each chassis attached to the RAID enclosure.

Operation

Click on a button (or buttons) to change the drive state, then click **OK** at the bottom of the screen to activate the changes. You would typically use the buttons as follows:

- Click **ON-LINE** to put a drive into the online state.
- Click **HOT SPARE** to turn one or more drives into hot spares.
- Click **UNCONFIGURED** to change the state of a drive from online to dead (a dead drive acts as though it is not even plugged in).

Changes are effective immediately, and no reboot is required. It is safe to perform these changes at any time (providing you are not taking a mounted LUN offline by marking its drives as **UNCONFIGURED**, of course).

When you click **OK** after having made your selections, you are returned to the main menu. If you click **Reset Screen**, all radio buttons that you have changed are set back to their previous states.

Field Definitions

The definition of each field in the Modify Disk Device Status dialog box is given in Table 3-2.

Table 3-2 Modify Disk Device Status Field Definitions

Field	Definition
Make/Model/ [Firmware]	This field returns the SCSI vendor ID, product ID, and firmware revision for each drive. For proper operation, it is strongly recommended that each drive in a LUN have the same make, model, and firmware release.
Serial Number	The drive serial number
Worldwide Name (7.01 and later firmware only)	A 64-bit identifier assigned to a particular drive. It is used to distinguish one drive from another. The WWN may be used for network management purposes or whenever drive identification is needed.
Clock Speed	The clock speed in MHz of the device. This is not very interesting now, but when 200 Mhz and 400 MHz fibre channel drives become supported, it will become quite important.
Number of Ports	The number of ports used by each disk. The number should normally be 2 for FC and 1 for SCSI.
Interface	The drive controller interface. FC-AL = Fibre Channel Arbitrated Loop.
Drive Speed (7.01 and later firmware only)	The spindle speed in RPM of the hard disk drive
Physical	The physical number of blocks and MB on the disk. 1 MB = 1024 * 1024 bytes, and one block = 512 bytes.
Usable	The usable number of blocks and MB on the disk. These numbers are always less than the physical numbers because the RAID controller allocates a portion of disk space for it's Configuration On Disk (COD). Usable sizes vary, based on the version controller firmware and Configuration on Disk (COD) versions: SGI 6.11 and 6.14 firmware used COD version 1.0. SGI 7.01 and later firmware uses COD version 2.1.
LoopID	A unique hexadecimal number for a particular disk drive. The LoopID is basically the equivalent of a SCSI ID.
Channel/TargetID	The channel number and target ID for a drive. GAM uses this extensively to identify individual disks.
List of LUNS	A list of the logical drives that are using some or all of the space on a disk drive.
Status	The Status area is color-coded. In addition to HOT SPARE and ON-LINE, it is possible that the drive could be in another state, such as UNCONFIGURED.

Finally, TPM does not care to which state you change a drive, so use common sense. If the **Modify Disk Device Status** screen shows that a drive is used within a LUN, and you change the drive from ON-LINE to some other state, data loss could result.

Create/Destroy/Expand Logical Drives (Configuration Function Menu)

To set up new drive configurations or to add, delete, or expand drives in an existing configuration, select **Create/Destroy/Expand Logical Drives** under the **Configuration Function** menu. The dialog box shown in Figure 3-5 appears.

<p>This section lets you set up new configurations, or add/delete/expand logical drives in existing configurations. You will be walked through each choice, and the configurator will not change anything until specifically instructed to do so. Click here for a help screen which describes your RAID options in detail, and provides you with relative capacity, availability, and performance characteristics of each each type. Or chose one of the buttons below:</p>	
<p>New Configuration</p> <p>This creates a new disk configuration. <i>(All of your data will be destroyed in the process!)</i></p>	<input type="button" value="New Configuration"/>
<p>Add a LUN - Retains ALL Data</p> <p>This lets you use free disk space to create an additional logical RAID disk (LUN). Existing data will not be affected. You must have unconfigured disks available, and at least one LUN must already exist.</p>	<input type="button" value="Add LUN(s)"/>
<p>Delete last LUN - Retains ALL Data (Except that of the LUN which will be destroyed).</p> <p>This will let you delete the last LUN which was defined. (Limitations with the current firmware make it impossible to delete anything but the last LUN.) If you click on the Delete LUN, then the next screen will describe it's characteristics, and give you the chance to get out cancel the procedure. Note: The configurator will <i>not</i> attempt to determine if there is a mounted filesystem on the LUN.</p>	<input type="button" value="Delete LUN"/>
<p>Expand a LUN - Retains ALL Data</p> <p>This lets you add capacity to an existing LUN while the controller is online with the host(s). For example, a system using a 6-disk RAID5 set can add another disk to create a 7-disk drive set. This procedure is also referred to by the acronym, "ORE", for On-Line Raid Expansion..</p> <ul style="list-style-type: none"> • During the expansion, which includes re-striping data from the old (smaller) set to the larger set, the controller continues to service host I/O operations • MORE is supported in the simplex mode of operation only. One controller in a dual-active controller system must be disabled (failed-over). Attempting to do this operation in a dual-active environment will be rejected. • One to six drives can be added to a set at a time. The maximum number of physical disk drives is sixteen. • You can not have 16 LUNs defined. The expansion requires one free LUN in order to execute. • You can not perform an expansion if all 16 LUNs are defined. • There must be no SPANNED LUNs defined. • The disk drives being added must not be part of an array. • The capacity of each of the added disk drives must be greater than or equal to the size of the smallest disk drive in the set. • The added capacity parameters are kept in non-volatile memory. In the event of power loss to this host or to the subsystem, (or a controller failure) then the process will automatically resume when power is restored, or the controller is replaced. • In the event of a disk drive failure, the process continues to completion in CRITICAL mode. • ORE, initialize, Rebuild, and Consistency Check are mutually exclusive operations. Only one process may run at a time. • No configuration update commands will be accepted during the expansion. • Write-back cache is disabled during the expansion. • The LUN must be on-line 	<input type="button" value="Expand LUN"/>

Figure 3-5 Create/Destroy/Expand Logical Drives Dialog Box

The **Create/Destroy/Expand Logical Drives** dialog box contains the following buttons:

- **New Configuration:** allows you to create a new disk configuration. If LUNs have already been defined, they must be deleted.

- Add a LUN: allows you to use free disk space to create one or more additional LUN(s).
- Delete Last LUN: allows you to delete the last LUN that was defined.
- Expand a LUN: allows you to add capacity to an existing LUN.

Note: At the top of the screen is a link marked **Click here**. Clicking this link opens a window that provides detailed information on each type of RAID, along with performance characteristics and data reliability considerations. If one has not had factory training on the controllers, this information will probably be quite useful.

The following sections explain how to use these buttons in more detail.

New Configuration

When you click **New Configuration**, the screen appears as shown in Figure 3-6.

RAID Array Configuration for system at `/hw/scsi/sc13d010`

Drive Pack	Drive 0	Drive 1	Drive 2	Drive 3	Drive 4	Drive 5	Drive 6	Drive 7	Drive 8	Drive 9	Drive 10	Drive 11	Drive 12	Drive 13	Drive 14	Drive 15	Physical MB	Maximum Logical MB			
	Ch-ID	Ch-ID	Ch-ID	Ch-ID	Ch-ID	Ch-ID	Ch-ID	Ch-ID	Ch-ID	Ch-ID	Ch-ID	Ch-ID	Ch-ID	Ch-ID	Ch-ID	Ch-ID		RAID 0	RAID 1	RAID 3/5	RAID 0+1
A	Next Assign																				
Stripe Size (All LUNs) is TBD. Totals:																	0	0	0	0	0

RAID Subsystems:

Enclosure #0 (Rackmount view, rotate 90° clockwise if in tower)

17560MB 0-125	33003MB 1-124	17560MB 0-123	17560MB 1-122
17560MB 0-121	17560MB 1-120	17560MB 0-119	17560MB 1-118
17560MB 0-117	17560MB 1-116	17560MB 0-115	17560MB 1-114

- If any disks above are colored GREEN, then they may contain active LUNs. You will, however, be allowed to reconfigure them without requiring you to manually delete any active LUNs (Provided you selected the New Configuration option).
- In order to build LUNs, you first must organize the disk drives into packs of up to 16 drives.
- The **Next Assign** box will show you what pack and drive number that the next drive you **assign** will be assigned to.
- Pack numbers are only for the benefit of the configurator, and they are not saved within the controller.

[\[Return to Main Page\]](#) [\[Log Off\]](#)

Figure 3-6 New Configuration Screen

You may create a new LUN configuration if presently there are no configured LUNs in your array. Drive packs can be spanned, which provides the ability to configure multiple or parts of multiple drive packs as one system drive. This allows more than 8 or 16 (firmware version 7.01 or greater) physical disk drives to be configured as a system drive. A system drive can span up to 4 or 16 (firmware version 7.01 or greater) drive packs. The controller creates the spanned system drive during the array configuration process. Data is striped across the drive packs of the spanned system drive. See the section titled “Creating and Defining New LUNs” for more information.

Creating Drive Packs

The primary rules for creating packs are listed below. More details on how to do this are found in subsequent sections of this document.

- The maximum number of packs that can be combined into a LUN is four for controllers with 6.14 firmware.
- The maximum number of packs that can be combined into a LUN is 16 for controllers with 7.01 and later firmware.
- 1 to 8 disk drives can be combined into a pack for controllers with 6.14 firmware.
- 1 to 16 disk drives can be combined into a pack for controllers with 7.01 and later firmware.
- The number of drives in a pack determines the possible RAID levels.
- If spanning packs into a LUN, all packs must have the same number of disks.
- Any drive of any size may be used in a pack, but the amount of usable storage will be computed as the smallest disk times the number of drives in the pack.

Creating and Defining New LUNs

System drives are the logical devices (storage volumes) that are presented to the operating system. During the configuration process, after physical disk drive packs are defined, one or more system drives must be created from the drive packs. System drives have the following properties:

- More than one system drive can be defined on a single drive pack. Or, a system drive can span 4 (firmware below 7.01) or 16 (firmware greater than 7.01) packs.
- The minimum size of a system drive is 8 MB. The maximum size is 2 TB (see the note on next page)
- Up to 8 (firmware below 7.01) or 32 (firmware greater than 7.01) system drives can be created.
- Each system drive has a RAID level which is selectable (subject to the number of disk drives in the system drive's pack).
- Each system drive has its own write policy (write-back or write-through).
- Each system drive has its own affinity or LUN mapping.

Note: SGI does not support Spanning or MORE operations on the IRIX and Linux operating systems. See "Expand LUN" on page 62 for more information on MORE operations.

Note: The SGI IRIX operating system disk utility “fx” has limited support of up to 1 terabyte for a single LUN, or volume (.999 terabyte (2147483647 blocks [1099511627264 bytes]) required for fx volume header). Because the maximum supported TP9100 LUN capacity size is 2 terabytes, a bug has been submitted requesting that changes be made to Irix infrastructure (for example, fx, prtvtoc, and so on) in order to support a 2 terabyte LUN (1.999 terabytes (4294967295 blocks [2199023255040 bytes]) required for fx volume header). Please take this in account ahead of time and do not create and initialize system drives (LUNs) that exceed 1099511627264 bytes.

Follow this procedure to create and define new LUNs. More details on how to do this are found starting in Figure 3-7 on page 41 of this manual.

1. Activate the desired checkboxes in the **RAID Subsystems** area of the screen shown in Figure 3-6 and click **Select Marked** to group the desired drives into a pack.
The packs exist only for configuration purposes, and are used to group drives together for easier configuration.
2. To continue assigning drives into packs, repeat step 1 until all desired drives are assigned into packs.
3. After at least one pack has been created, you can use the **Select Previous Pack** button to reconfigure a previously configured pack.
4. After you have finished creating the desired drive packs, use the **Define LUNs** button to create a system disk (LUN) using up to eight of the drive packs (for 6.14 firmware), or up to 32 of the drive packs (for 7.01 and later firmware) you just created. You may use up to eight drive packs (32 drives maximum) with 6.14 controller firmware or up 32 drive packs maximum (60 drives maximum) with 7.01 and later firmware.
5. Select the drive pack(s) that you want to define as the new LUN and click **Configure New**.
6. Select the desired RAID level for the LUN you are creating and click **Apply**.
7. To continue building LUNs that incorporate drives in other packs, click **Select Pack(s)** and repeat the process of defining LUNs and their corresponding RAID levels and usable MB.
8. Update the RAID controller with the new LUN information.

The next sections explain how to execute these steps.

Group the First Set of Drives into a Pack

To group the first set of drives into a pack, follow these directions:

1. Activate the desired checkboxes in the RAID Subsystems area of the dialog box (see Figure 3-6) and click **Select Marked** to move drives from the map shown in the RAID Subsystems area of the screen into Pack A.

The top area of the screen contains a table that shows the drives assigned to drive pack A. **Next Assign** now appears in a new row to indicate that drives can now be selected and assigned to drive pack B through a similar process.

As you assign drives to the pack, the screen changes to the example shown in Figure 3-7 (6.14 firmware) or Figure 3-8 (7.01 and later firmware).

RAID Array Configuration for system at /hw/scsi/sc39d010

Drive Pack	Drive 0 Ch-ID	Drive 1 Ch-ID	Drive 2 Ch-ID	Drive 3 Ch-ID	Drive 4 Ch-ID	Drive 5 Ch-ID	Drive 6 Ch-ID	Drive 7 Ch-ID	Physical MB	Maximum Logical MB				
										RAID 0	RAID 1	RAID 3/5	RAID 0+1	
A	17501MB 0-0 <input type="checkbox"/> Undo	17501MB 1-0 <input type="checkbox"/> Undo	17501MB 2-0 <input type="checkbox"/> Undo	17501MB 3-0 <input type="checkbox"/> Undo					70007	69940	34970	52455	34970	
B	Next Assign													
Stripe Size (All LUNs) is 64KB. Totals:										70007	69940	34970	52455	34970

RAID Subsystems:

Enclosure #0 (Rackmount view, rotate 90° clockwise if in tower)

17501MB 0-0 <input type="checkbox"/>	17501MB 1-0 <input type="checkbox"/>	17501MB 0-1 <input type="checkbox"/>	17501MB 1-1 <input type="checkbox"/>
17501MB 0-2 <input type="checkbox"/>	17501MB 1-2 <input type="checkbox"/>	17501MB 0-3 <input type="checkbox"/>	17501MB 1-3 <input type="checkbox"/>
17501MB 0-5 <input type="checkbox"/>	17501MB 1-5 <input type="checkbox"/>	17501MB 0-4 <input type="checkbox"/>	17501MB 1-4 <input type="checkbox"/>

- If any disks above are colored GREEN, then they may contain active LUNs. You will, however, be allowed to reconfigure them without requiring you to manually delete any active LUNs (Provided you selected the New Configuration option).
- In order to build LUNs, you first must organize the disk drives into packs of up to 8 drives.
- The **Next Assign** box will show you what pack and drive number that the next drive you **assign** will be assigned to.
- Pack numbers are only for the benefit of the configurator, and they are not saved within the controller.

[\[Return to Main Page\]](#) [\[Log Off\]](#)

TP9100 Array Manager (TPM 1.0).

Figure 3-7 Drives Assigned Into Pack A (6.14 Firmware)

RAID Array Configuration for system at `/hw/scsi/sc3d010`

Drive Pack	Drive 0 Ch-ID	Drive 1 Ch-ID	Drive 2 Ch-ID	Drive 3 Ch-ID	Drive 4 Ch-ID	Drive 5 Ch-ID	Drive 6 Ch-ID	Drive 7 Ch-ID	Drive 8 Ch-ID	Drive 9 Ch-ID	Drive 10 Ch-ID	Drive 11 Ch-ID	Drive 12 Ch-ID	Drive 13 Ch-ID	Drive 14 Ch-ID	Drive 15 Ch-ID	Physical MB	Maximum Logical MB				
																		RAID 0	RAID 1	RAID 3/5	RAID 0+1	
A	17500MB 0-125 <input checked="" type="checkbox"/>	17500MB 1-124 <input checked="" type="checkbox"/>	17500MB 0-123 <input checked="" type="checkbox"/>	17500MB 1-122 <input checked="" type="checkbox"/>													70242	68664	34332	51498	34332	
B	Next Assign																					
Stripe Size (All LUNs) is TBD. Totals:																	70242	68664	34332	51498	34332	

RAID Subsystems:			
Enclosure #0 (Rackmount view, rotate 90° clockwise if in tower)			
17500MB 0-125 <input type="checkbox"/>	17500MB 1-124 <input type="checkbox"/>	17500MB 0-123 <input type="checkbox"/>	17500MB 1-122 <input type="checkbox"/>
17500MB 0-121 <input type="checkbox"/>	17500MB 1-120 <input type="checkbox"/>	17500MB 0-119 <input type="checkbox"/>	17500MB 1-118 <input type="checkbox"/>
17500MB 0-117 <input type="checkbox"/>	17500MB 1-116 <input type="checkbox"/>	17500MB 0-115 <input type="checkbox"/>	17500MB 1-114 <input type="checkbox"/>

- If any disks above are colored GREEN, then they may contain active LUNs. You will, however, be allowed to reconfigure them without requiring you to manually delete any active LUNs (Provided you selected the New Configuration option).
- In order to build LUNs, you first must organize the disk drives into packs of up to 16 drives.
- The **Next Assign** box will show you what pack and drive number that the next drive you **assign** will be assigned to.
- Pack numbers are only for the benefit of the configurator, and they are not saved within the controller.

[\[Return to Main Page\]](#) [\[Log Off\]](#)

Figure 3-8 Drives Assigned Into Pack A (7.01 and Later Firmware)

Figure 3-7 and Figure 3-8 show that four drives have been assigned into drive pack A, which could become a RAID set. Every time a set of drives is assigned, the table cell labeled **Next Assign** moves to the next row down, where a new drive pack can be created.

- To remove a drive from a pack, click the undo box (for 7.01 and later firmware) or click the **Undo** button (for 6.14 firmware), and the remaining drives in the pack shift to the left, while the removed drive reappears with a cleared checkbox in the RAID Subsystem table. Use the **Select Previous Pack** button to reconfigure the drive arrangement in an earlier drive pack.

Nothing is saved for several more screens, and you may cancel at any time by pressing the **Cancel** button. Use the **Clear All** button to deassign all drives from the drive packs.

Note: You do not have to allocate all the drives to packs (or LUNs). If you choose not to configure certain drives, they may be used at any time when you click **Add LUN(s)** or **Expand LUN** (see Figure 3-5), if the characteristics of the LUN qualify for expansion.

Group the Remaining Drives Into Packs

To create remaining drive packs, continue activating drive checkboxes in the RAID Subsystem area of the dialog box and using the **Select Next Pack** button as necessary until the drives are grouped as desired into separate drive packs.

Note: Refer to “New Configuration” on page 38 for rules on drive pack configuration.

When you are finished, the screen looks similar to the one pictured in Figure 3-9 (6.14 firmware) or Figure 3-10 (7.01 and later firmware). In both cases, the drives have been grouped into three separate drive packs.

RAID Array Configuration for system at `/hw/scsi/sc39d010`

Drive Pack	Drive 0	Drive 1	Drive 2	Drive 3	Drive 4	Drive 5	Drive 6	Drive 7	Physical MB	Maximum Logical MB				
	Ch-ID	Ch-ID	Ch-ID	Ch-ID	Ch-ID	Ch-ID	Ch-ID	Ch-ID		RAID 0	RAID 1	RAID 3/5	RAID 0+1	
A	17501MB 0-0 Undo	17501MB 1-0 Undo	17501MB 0-1 Undo	17501MB 1-1 Undo					70007	69940	34970	52455	34970	
B	17501MB 0-2 Undo	17501MB 1-2 Undo	17501MB 0-3 Undo	17501MB 1-3 Undo					70007	69940	34970	52455	34970	
C	17501MB 0-4 Undo	17501MB 1-4 Undo	17501MB 0-5 Undo	17501MB 1-5 Undo					70007	69940	34970	52455	34970	
D	Next Assign													
Stripe Size (All LUNs) is 64KB. Totals:										210021	209820	104910	157365	104910

RAID Subsystems:				
Enclosure #0 (Rackmount view, rotate 90° clockwise if in tower)				
17501MB 0-0	17501MB 1-0	17501MB 0-1	17501MB 1-1	
17501MB 0-2	17501MB 1-2	17501MB 0-3	17501MB 1-3	
17501MB 0-4	17501MB 1-4	17501MB 0-5	17501MB 1-5	

Clear All	Select Previous Pack	Define LUNs	Select Marked	Cancel
-----------	----------------------	-------------	---------------	--------

Figure 3-9 All Drive Packs Created (6.14 Firmware)

RAID Array Configuration for system at `/lun/scsi/sc3d010`

Drive Pack	Drive 0 Ch-ID	Drive 1 Ch-ID	Drive 2 Ch-ID	Drive 3 Ch-ID	Drive 4 Ch-ID	Drive 5 Ch-ID	Drive 6 Ch-ID	Drive 7 Ch-ID	Drive 8 Ch-ID	Drive 9 Ch-ID	Drive 10 Ch-ID	Drive 11 Ch-ID	Drive 12 Ch-ID	Drive 13 Ch-ID	Drive 14 Ch-ID	Drive 15 Ch-ID	Physical MB	Maximum Logical MB			
																		RAID 0	RAID 1	RAID 3/5	RAID 0+1
A	17560MB 0-125	17560MB 1-124	17560MB 0-122	17560MB 1-122													70242	68664	34332	51498	34332
B	17560MB 0-121	17560MB 1-120	17560MB 0-119	17560MB 1-119													70242	68664	34332	51498	34332
C	17560MB 0-117	17560MB 1-116	17560MB 0-115	17560MB 1-114													70242	68664	34332	51498	34332
D	Next Assign																				
Stripe Size [All LUNs] is TBD. Totals:																	210728	205992	102996	154494	102996

RAID Subsystems:			
Enclosure #0 (Rackmount view, rotate 90° clockwise if in tower)			
17560MB 0-125	17560MB 1-124	17560MB 0-122	17560MB 1-122
17560MB 0-121	17560MB 1-120	17560MB 0-119	17560MB 1-119
17560MB 0-117	17560MB 1-116	17560MB 0-115	17560MB 1-114

- If any disks above are colored GREEN, then they may contain active LUNs. You will, however, be allowed to reconfigure them without requiring you to manually delete any active LUNs (Provided you selected the New Configuration option).
- In order to build LUNs, you first must organize the disk drives into packs of up to 16 drives.
- The Next Assign box will show you what pack and drive number that the next drive you assign will be assigned to.
- Pack numbers are only for the benefit of the configurator, and they are not saved within the controller.

[\[Return to Main Page\]](#) [\[Log Off\]](#)

Figure 3-10 All Drive Packs Created (7.01 and Later Firmware)

Defining a LUN Using the Drive Packs

To define a LUN, follow these steps:

1. Click **Define LUNs**.

The screen shown in Figure 3-11 appears.

Logical Unit Configuration for system at `/hw/scsi/sc3d010`

Drive Pack																Raw MB	Usable MB	Allocated MB (% Used)	Physical Distribution				
Pack	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14				15	LUN #	Drive Ch-ID	Starting Block	# Blocks
A <input type="checkbox"/>	17560MB 0-125	17560MB 1-124	17560MB 0-123	17560MB 1-122													70242	68664	0 (0.0 %)				
B <input type="checkbox"/>	17560MB 0-121	17560MB 1-120	17560MB 0-119	17560MB 1-118													70242	68664	0 (0.0 %)				
C <input type="checkbox"/>	17560MB 0-117	17560MB 1-116	17560MB 0-115	17560MB 1-114													70242	68664	0 (0.0 %)				

Select the pack or pack(s) you wish to assign to the next group of LUNs you will create, and press on an action button. The rules for combining packs are:

- A pack may be split into multiple LUNs.
- Each pack must have the same number of disk drives (if you want to combine them).
- You need a minimum of 2 disks for RAID 0,1; 3 for RAID 3,5; and 3 for RAID 0+1. Additional constraints on allowable RAID structures may also be applied, depending on whether or not you are spanning packs. This configurator will prevent you from creating an invalid configuration, *once you select which pack(s) you want to use to define this next logical disk*.
- A maximum of 16 packs may be combined into a single LUN, but you can not exceed 4,294,967,295 blocks (approx 2.1 TB).

Logical Drive#	RAID Level	Pack(s) Used	Usable MB	Write Cache	Stripe Size
New	<i>Select pack(s) from above, then press an action button below.</i>				

Configure New

Cancel

Figure 3-11 Select Drive Packs to Include in LUN

2. Check the boxes at the left (A, B, or C) to select one or more packs to combine into a LUN.
3. To create a LUN from drive pack A, for example, activate the A box and click **Configure New** at the bottom of the screen.

The screen in Figure 3-13 (6.14 firmware) or Figure 3-14 (7.01 and later firmware) appears.

Note: For controllers with 6.14 firmware, you must go to the **View/Modify RAID Controller Configuration** dialog box (see Figure 3-50 on page 87) and select the stripe size before defining the LUN.

Logical Unit Configuration for system at `/hw/scsi/sc13d010`

Pack	Drive Pack															Raw MB	Usable MB	Allocated MB (% Used)	Physical Distribution				
	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14				15	LUN #	Drive Ch-ID	Starting Block	# Blocks
A	17560MB 0-125	17560MB 0-123	17560MB 1-115	17560MB 1-114													70242	70120	0 (0.0%)				
B	17560MB 0-121	17560MB 1-120	17560MB 1-118	35003MB 1-124													87685	70120	0 (0.0%)				
C	17560MB 1-122	17560MB 0-117	17560MB 1-116	17560MB 0-119													105364	105180	0 (0.0%)				

* A maximum of 16 packs may be combined into a single LUN, but do not exceed 96 drives, or 2 TB total MB.

Logical Drive#	RAID Level	Pack(s) Used	Usable MB	Write Cache
0	RAID 0 (stripe) Max MB=70120	TBD	ALL	<input checked="" type="checkbox"/>

If you enter ALL for the MB field, then the configurator will automatically allocate all remaining storage from the selected pack(s).

Figure 3-12 LUN is Being Defined (6.14 Firmware)

Logical Unit Configuration for system at `/hw/scsi/sc3d010`

Pack	Drive Pack															Raw MB	Usable MB	Allocated MB (% Used)	Physical Distribution					
	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14				15	LUN #	Drive Ch-ID	Starting Block	# Blocks	# MB
A	17560MB 0-125	17560MB 1-124	17560MB 0-123	17560MB 1-122													70242	68664	0 (0.0%)					
B	17560MB 0-121	17560MB 1-120	17560MB 0-119	17560MB 1-118													70242	68664	0 (0.0%)					
C	17560MB 0-117	17560MB 1-116	17560MB 0-115	17560MB 1-114													70242	68664	0 (0.0%)					

* A maximum of 16 packs may be combined into a single LUN, but you can not exceed 4,294,967,295 blocks (approx 2.1 TB).

Logical Drive#	RAID Level	Pack(s) Used	Usable MB	Write Cache	Stripe Size
0	RAID 0 (stripe) Max MB=68664	TBD	All	<input checked="" type="checkbox"/>	<input type="radio"/> 8 KB <input type="radio"/> 16 KB <input type="radio"/> 32 KB <input checked="" type="radio"/> 64 KB

Apply Cancel

If you enter **ALL** for the MB field, then the configurator will automatically allocate all remaining storage from the selected pack(s).

Figure 3-13 LUN is Being Defined (7.01 and Later Firmware)

These dialog boxes show that a Logical Drive #0 (LUN 0) is being created that allows you to select the RAID level and usable MB.

- Select the desired RAID level from the drop-down box shown.
- In this example, leave **All** in the **Usable MB** field, and click **Apply** to actually create the LUN.

The screen shown in Figure 3-14 appears.

Logical Unit Configuration for system at /hw/scsi/sc3d010

Pack	Drive Pack															Raw MB	Usable MB	Allocated MB (% Used)	Physical Distribution				
	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14				15	LUN #	Drive Ch-ID	Starting Block	# Blocks
A	17560MB 0-125	17560MB 1-124	17560MB 0-123	17560MB 1-122												70242	68664	68664 (100.0 %)	0 0 0 0	0-125 1-124 0-123 1-122	0 0 0 0	35155968 35155968 35155968 35155968	17166 17166 17166 17166
B	17560MB 0-121	17560MB 1-120	17560MB 0-119	17560MB 1-118												70242	68664	0 (0.0 %)					
C	17560MB 0-117	17560MB 1-116	17560MB 0-115	17560MB 1-114												70242	68664	0 (0.0 %)					

* A maximum of 16 packs may be combined into a single LUN, but you can not exceed 4,294,967,295 blocks (approx 2.1 TB).

Logical Drive#	RAID Level	Pack(s) Used	Usable MB	Write Cache	Stripe Size
0	RAID 0 (Stripe)	A	68664	Y	64 KB
1	RAID 0 (Stripe)		0	Y	64 KB

Select Pack(s) Next Screen Cancel

Figure 3-14 LUN is Defined

The right-hand side of the table shows how the LUN is physically arranged on each disk drive.

- To continue building LUNs that incorporate drives in other packs, click **Select Pack(s)** and repeat the process of defining LUNs and their corresponding RAID levels and usable MB.

Figure 3-15 (6.14 firmware) or Figure 3-16 (7.01 and later firmware) shows how the screen appears after using all the drive packs to define three separate LUNs.

Logical Unit Configuration for system at /hw/scsi/20000080e5110413/lun0/c5p1

Drive Pack								Raw MB	Usable MB	Allocated MB (% Used)	Physical Distribution					
Pack	Drive 0	Drive 1	Drive 2	Drive 3	Drive 4	Drive 5	Drive 6				Drive 7	LUN #	Drive Ch-ID	Starting Block	# Blocks	# MB
A	70007MB 0-0	70007MB 1-0	70007MB 0-1	70007MB 1-1					280028	279964	279964 (100.0%)	0 0 0	0-0 0-1 1-1	0 0 0	143341568 143341568 143341568	69991 69991 69991
B	70007MB 0-2	70007MB 1-2	70007MB 0-3	70007MB 1-3					280028	279964	279964 (100.0%)	1 1 1	0-2 1-2 0-3 1-3	0 0 0 0	143341568 143341568 143341568 143341568	69991 69991 69991 69991
C	70007MB 0-4	70007MB 1-4	70007MB 0-5	70007MB 1-5					280028	279964	279964 (100.0%)	2 2 2 2	0-4 1-4 0-5 1-5	0 0 0 0	143341568 143341568 143341568 143341568	69991 69991 69991 69991

* A maximum of 4 packs (total of 32 disks) may be combined into a single LUN, but you can not exceed 4,294,967,295 blocks (approx 2.1 TB).

Logical Drive#	RAID Level	Pack(s) Used	Usable MB	Write Cache
0	RAID 0 (Stripe)	A	279964	Y
1	RAID 0+1 (Mirrored Stripe)	B	139982	Y
2	RAID 5 (Parity Stripe)	C	209973	Y

Next Screen Cancel

Figure 3-15 All LUNs Defined (6.14 Firmware)

Logical Unit Configuration for system at `/hw/scsi/sc3d010`

Pack	Drive Pack															Raw MB	Usable MB	Allocated MB (% Used)	Physical Distribution				
	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14				15	LUN #	Drive Ch-ID	Starting Block	# Blocks
A	17560MB 0-125	17560MB 1-124	17560MB 0-123	17560MB 1-122												70242	68664	68664 (100.0 %)	0	0-125	0	35155968	17166
																			0	0-124	0	35155968	17166
																			0	0-123	0	35155968	17166
B	17560MB 0-121	17560MB 1-120	17560MB 0-119	17560MB 1-118												70242	68664	68664 (100.0 %)	1	0-121	0	35155968	17166
																			1	1-120	0	35155968	17166
																			1	0-119	0	35155968	17166
C	17560MB 0-117	17560MB 1-116	17560MB 0-115	17560MB 1-114												70242	68664	68664 (100.0 %)	2	0-117	0	35155968	17166
																			2	1-116	0	35155968	17166
																			2	0-115	0	35155968	17166

* A maximum of 16 packs may be combined into a single LUN, but you can not exceed 4,294,967,295 blocks (approx 2.1 TB).

Logical Drive#	RAID Level	Pack(s) Used	Usable MB	Write Cache	Stripe Size
0	RAID 0 (Stripe)	A	68664	Y	64 KB
1	RAID 0+1 (Mirrored Stripe)	B	34332	Y	64 KB
2	RAID 5 (Right Asymmetric)	C	51498	Y	64 KB

Next Screen Cancel

Figure 3-16 All LUNs Defined (7.01 and later Firmware)

At this point, all the LUNs have been defined. Next, the RAID controller must be updated with the new configuration.

Updating the RAID Controller

To update the RAID controller with the new LUN configuration, follow these steps:

1. Press the **Next Screen** button.

For 6.14, 7.01, and 7.03 firmware, the warning screen shown Figure 3-17 appears.




Figure 3-17 Updating the RAID Controller Warning Screen (6.14, 7.01, and 7.03 Firmware)

For 7.75 firmware, the warning screen shown Figure 3-18 appears.




Figure 3-18 Updating the RAID Controller Warning Screen (7.75 Firmware)

2. Click **YES** to save the LUN configuration; otherwise click **NO**.

If all goes well, the screen shown in Figure 3-19 appears (for 6.14, 7.01, and 7.03 firmware versions), indicating that the new LUN(s) are online. If you are running 7.75 firmware, the screen shown in Figure 3-20 appears.

3. If you are running 6.14, 7.01, or 7.03 firmware, go to step 5.
4. If you are running 7.75 firmware, go to step 9.

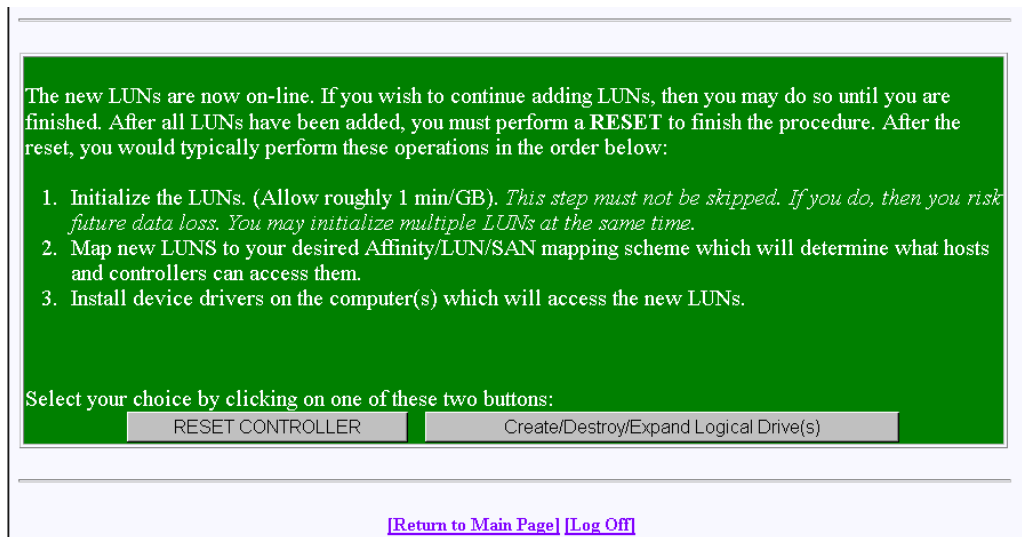


Figure 3-19 LUN Online Screen (6.14, 7.01, and 7.03 firmware only)

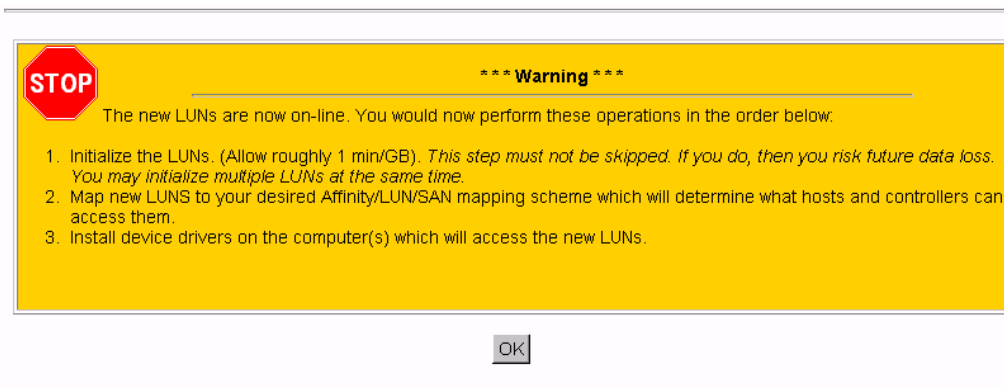


Figure 3-20 LUN Online Screen (7.75 firmware only)

For 6.14, 7.01, and 7.03 Firmware:

5. Click the **Create/Destroy/Expand Logical Drive(s)** button if you want to add more LUNs or reconfigure existing LUNs.
6. Click **Reset Controller** after all LUNs have been added and properly configured.

The screen shown in Figure 3-21 appears.




Figure 3-21 Warning Screen (6.14, 7.01, and 7.03 Firmware)

Note: If you are adding a new LUN to an existing configuration, you should unmount any LUNs before resetting. This is especially important if there is a possibility that the host will attempt to issue any I/Os to the LUNs before the controller(s) finish rebooting.

7. Click **YES**.

The window shown in Figure 3-22 appears.

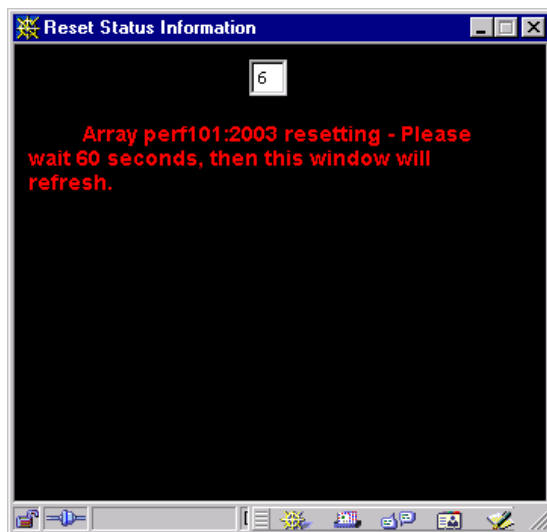


Figure 3-22 Controller Resetting Information Screen

A reset is required, and is not optional. The reason that TPM even asks you about initiating a reset is to give you the chance to make sure your host(s) are ready for a reset.

For 6.14, 7.01, and 7.03 firmware, the two Reset Status Information windows appear, as shown in Figure 3-23 and Figure 3-24.

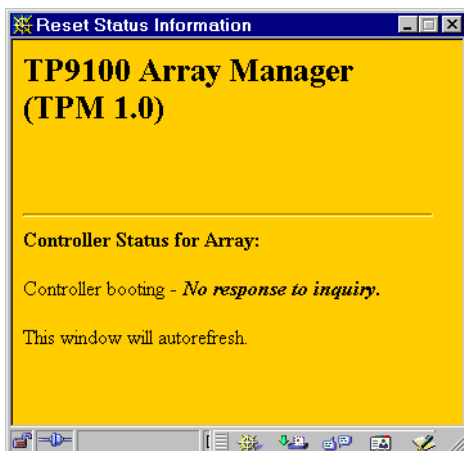


Figure 3-23 Controller Resetting Status Screen (1 of 2)

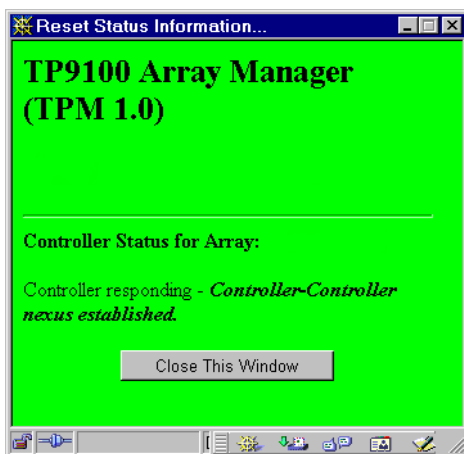


Figure 3-24 Controller Resetting Status Screen (2 of 2)

8. Click **Close This Window** on the dialog box shown in Figure 3-24.

After the controller(s) are reset (for 6.14, 7.01, and 7.03 firmware only), you must return to the main menu and do the following:

- a. Install the device driver for the new LUN (see “Installing Device Drivers for New LUNS (Configuration Function Menu)” on page 83).
- b. Initialize the LUNs (see “Initialize LUN(s) (Configuration Function Menu)” on page 68).

For 7.75 Firmware:

9. After you click OK on the screen shown in Figure 3-20, you are returned back to the main page. You must then follow these steps:
 - a. Install the device driver for the new LUN (see “Installing Device Drivers for New LUNS (Configuration Function Menu)” on page 83).
 - b. Initialize the LUNs (see “Initialize LUN(s) (Configuration Function Menu)” on page 68).

Add LUN(s)

Back at the main menu, click select **Create/Destroy/Expand Logical Drives** under the **Configuration Function** menu. The screen shown in Figure 3-25 appears for 6.14, 7.01, and 7.03 firmware. For 7.75 firmware, the screen shown in Figure 3-26 appears.

<p>This section lets you set up new configurations, or add/delete/expand logical drives in existing configurations. You will be walked through each choice, and the configurator will not change anything until specifically instructed to do so. Click here for a help screen which describes your RAID options in detail, and provides you with relative capacity, availability, and performance characteristics of each each type. Or chose one of the buttons below:</p>	
<p>New Configuration</p> <p>This creates a new disk configuration. The configurator will prevent you from running this if any LUNs are already defined. (You must perform the Delete LUN operation repeatedly until all LUNs have been removed.)</p>	<input type="button" value="New Configuration"/>
<p>Add a LUN - Retains ALL Data</p> <p>This lets you use free disk space to create an additional logical RAID disk (LUN). Existing data will not be affected. You must have unconfigured disks available, and at least one LUN must already exist.</p>	<input type="button" value="Add LUN(s)"/>
<p>Delete last LUN - Retains ALL Data (Except that of the LUN which will be destroyed).</p> <p>This will let you delete the last LUN which was defined. (Limitations with the current firmware make it impossible to delete anything but the last LUN.) If you click on the Delete LUN, then the next screen will describe it's characteristics, and give you the chance to get out cancel the procedure. Note: The configurator will not attempt to determine if there is a mounted filesystem on the LUN.</p>	<input type="button" value="Delete LUN"/>
<p>Expand a LUN - Retains ALL Data</p> <p>This lets you add capacity to an existing LUN while the controller is online with the host(s). For example, a system using a 6-disk RAID5 set can add another disk to create a 7-disk drive set. This procedure is also referred to by the acronym, "ORE", for On-Line Raid Expansion.</p> <ul style="list-style-type: none"> • During the expansion, which includes re-striping data from the old (smaller) set to the larger set, the controller continues to service host I/O operations • MORE is supported in the simplex mode of operation only. One controller in a dual-active controller system must be disabled (failed-over). Attempting to do this operation in a dual-active environment will be rejected. • One to six drives can be added to a set at a time. The maximum number of physical disk drives is eight. • You can not have 8 LUNs defined. The expansion requires one free LUN in order to execute. • You can not perform an expansion if all 8 LUNs are defined. • The disk drives being added must not be part of an array. • The capacity of each of the added disk drives must be greater than or equal to the size of the smallest disk drive in the set. • The added capacity parameters are kept in non-volatile memory. In the event of power loss to this host or to the subsystem, (or a controller failure) then the process will automatically resume when power is restored, or the controller is replaced. • In the event of a disk drive failure, the process continues to completion in CRITICAL mode. 	<input type="button" value="Expand LUN"/>

Figure 3-25 Create/Destroy/Expand Logical Drives Screen (6.14 Firmware)

This section lets you set up new configurations, or add/delete/expand logical drives in existing configurations. You will be walked through each choice, and the configurator will not change anything until specifically instructed to do so. Click [here](#) for a help screen which describes your RAID options in detail, and provides you with relative capacity, availability, and performance characteristics of each each type. Or chose one of the buttons below:

<p>New Configuration</p> <p>This creates a new disk configuration. <i>(All of your data will be destroyed in the process!)</i></p>	<p>New Configuration</p>
<p>Add a LUN - Retains ALL Data</p> <p>This lets you use free disk space to create an additional logical RAID disk (LUN). Existing data will not be affected. You must have unconfigured disks available, and at least one LUN must already exist.</p>	<p>Add LUN(s)</p>
<p>Delete last LUN - Retains ALL Data (Except that of the LUN which will be destroyed).</p> <p>This will let you delete the last LUN which was defined. (Limitations with the current firmware make it impossible to delete anything but the last LUN.) If you click on the Delete LUN, then the next screen will describe it's characteristics, and give you the chance to get out cancel the procedure. Note: The configurator will <i>not</i> attempt to determine if there is a mounted filesystem on the LUN.</p>	<p>Delete LUN</p>
<p>Expand a LUN - Retains ALL Data</p> <p>This lets you add capacity to an existing LUN while the controller is online with the host(s). For example, a system using a 6-disk RAID5 set can add another disk to create a 7-disk drive set. This procedure is also referred to by the acronym, "ORE", for On-Line Raid Expansion..</p> <ul style="list-style-type: none"> • During the expansion, which includes re-striping data from the old (smaller) set to the larger set, the controller continues to service host I/O operations • MORE is supported in the simplex mode of operation only. One controller in a dual-active controller system must be disabled (failed-over). Attempting to do this operation in a dual-active environment will be rejected. • One to six drives can be added to a set at a time. The maximum number of physical disk drives is sixteen. • You can not have 16 LUNs defined. The expansion requires one free LUN in order to execute. • You can not perform an expansion if all 16 LUNs are defined. • There must be no SPANNED LUNs defined. • The disk drives being added must not be part of an array. • The capacity of each of the added disk drives must be greater than or equal to the size of the smallest disk drive in the set. • The added capacity parameters are kept in non-volatile memory. In the event of power loss to this host or to the subsystem, (or a controller failure) then the process will automatically resume when power is restored, or the controller is replaced. • In the event of a disk drive failure, the process continues to completion in CRITICAL mode. • ORE, Initialize, Rebuild, and Consistency Check are mutually exclusive operations. Only one process may run at a time. • No configuration update commands will be accepted during the expansion. • Write-back cache is disabled during the expansion. • The LUN must be on-line 	<p>Expand LUN</p>

Figure 3-26 Create/Destroy/Expand Logical Drives Screen (7.01 and Later Firmware)

To add one or more LUNs, follow these directions:

1. Click **Add LUN(s)**.

A screen similar to the one in Figure 3-27 on page 58 appears.

RAID Array Configuration for system at `/hw/scsi/sc3d010`

Drive Pack	Drive 0 Ch-ID	Drive 1 Ch-ID	Drive 2 Ch-ID	Drive 3 Ch-ID	Drive 4 Ch-ID	Drive 5 Ch-ID	Drive 6 Ch-ID	Drive 7 Ch-ID	Drive 8 Ch-ID	Drive 9 Ch-ID	Drive 10 Ch-ID	Drive 11 Ch-ID	Drive 12 Ch-ID	Drive 13 Ch-ID	Drive 14 Ch-ID	Drive 15 Ch-ID	Physical MB	Maximum Logical MB				
																		RAID 0	RAID 1	RAID 3/5	RAID 0+1	
A	17560MB 0-125	17560MB 1-124	17560MB 0-123	17560MB 1-122													70242	68664	34332	51498	34332	
B	Next Assign																					
Stripe Size (All LUNs) is 64KB. Totals:																	70242	68664	34332	51498	34332	

RAID Subsystems:

Enclosure #0 (Rackmount view, rotate 90° clockwise if in tower)

17560MB 0-125	17560MB 1-124	17560MB 0-123	17560MB 1-122
17560MB 0-121 <input type="checkbox"/>	17560MB 1-120 <input type="checkbox"/>	17560MB 0-119 <input type="checkbox"/>	17560MB 1-118 <input type="checkbox"/>
17560MB 0-117 <input type="checkbox"/>	17560MB 1-116 <input type="checkbox"/>	17560MB 0-115 <input type="checkbox"/>	17560MB 1-114 <input type="checkbox"/>

- If any disks above are colored GREEN, then they may contain active LUNs. You will, however, be allowed to reconfigure them without requiring you to manually delete any active LUNs (Provided you selected the New Configuration option).
- In order to build LUNs, you first must organize the disk drives into packs of up to 16 drives.
- The **Next Assign** box will show you what pack and drive number that the next drive you **assign** will be assigned to.
- Pack numbers are only for the benefit of the configurator, and they are not saved within the controller.

Figure 3-27 Add LUN(s) Screen

To add a LUN to an existing configuration, at least one LUN must already exist. In addition, unconfigured disks must be available with enough disk space to create an additional LUN.

2. Follow essentially the same procedures as before listed under “New Configuration” on page 38.

The procedure that is followed and the screens that appear are similar to adding a new configuration to an unconfigured controller. The differences are as follows:

- All previously defined LUNs and packs (which are actually used in LUNs) are displayed when selecting drives for packs. The allocated drives, however, do not have an empty checkbox in them (only unused disks have an empty checkbox).
- You may not add a drive to an existing pack. To do this, you must use the **Expand LUN** button (see Figure 3-25) to perform an **online RAID expansion** procedure.

- Although you can fill out the screens, you must not attempt to save the configuration using the **Next Screen** button while users are online, because a controller reset is required.
3. When you have finished defining the new LUN, click **Next Screen**.
The warning screen shown in Figure 3-28 appears for 6.14, 7.01, and 7.03 firmware. For 7.75 firmware, the screen shown in appears.

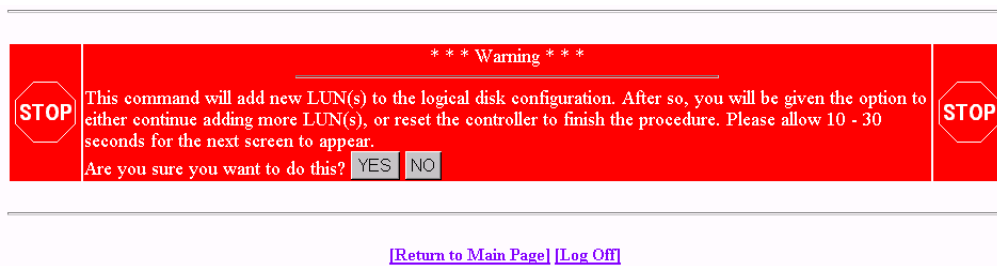


Figure 3-28 Add LUN(s) Warning Screen (6.14, 7.01, and 7.03 Firmware)


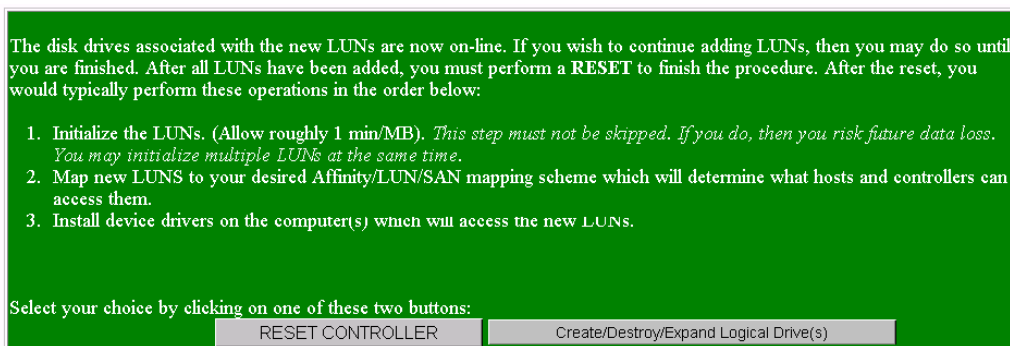


Figure 3-29 Add LUN(s) Warning Screen (7.75 Firmware)

4. Click **YES** to add the LUN, or **NO** to cancel and go back to the main menu.
5. If you are running 7.75 firmware, proceed to step 7; otherwise continue.
If you select **YES**, the confirmation screen in Figure 3-30 appears after a brief waiting period.



[\[Return to Main Page\]](#) [\[Log Off\]](#)

Figure 3-30 Add LUN(s) Confirmation Screen (6.14, 7.01, and 7.03 Firmware)

6. Click **RESET CONTROLLER** if no more LUNs are to be added, or **Create/Destroy/Expand Logical Drives** if you need to add more LUNs.

If you select **Create/Destroy/Expand Logical Drives**, you are returned to the dialog box of Figure 3-25 on page 56.

7. If you are running 7.75 firmware, the screen shown in Figure 3-31 appears.

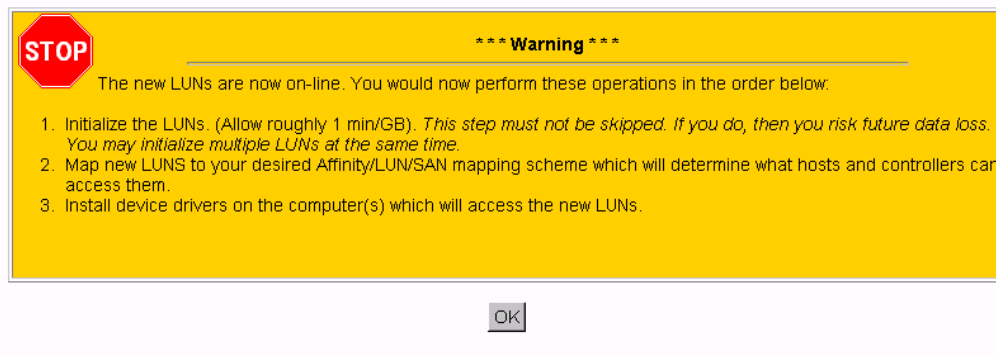


Figure 3-31 Add LUN(s) Confirmation Screen (7.75 Firmware)

8. After you click OK on the screen shown in Figure 3-31, you are returned back to the main page. You may then follow the steps shown in the screen of Figure 3-31.

Delete LUN

Use the **Delete LUN** button on the **Create/Destroy/Expand Logical Drives** screen (see Figure 3-25) to delete a LUN.

To delete a LUN, follow these directions:

1. Click **Delete LUN**.

The warning screen shown in Figure 3-32 on page 61 is displayed. This screen shows the size and characteristics of the last LUN that was created.

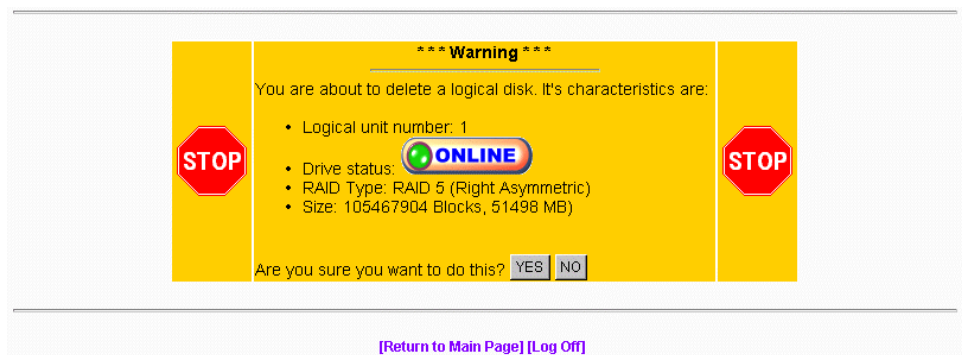


Figure 3-32 Delete LUN Warning Screen

2. Click **YES** to delete the LUN immediately.
3. Click **NO** to return to the main menu.

Don't delete a LUN with a mounted file system on it. This may lock up the server, or at the very least cripple it.

Warning: Make sure that you do not delete the LUN associated with the device driver you selected to talk to the RAID subsystem. If you do, TPM will not be able to communicate with the RAID controller until you choose another device driver from the Select Physical Device screen (see Figure 3-2 on page 27).

Expand LUN

RAID expansion allows capacity to be added to an existing RAID set while the controller is online. No resets are required, and I/Os to other LUNs will be serviced (a little slower). The following rules and conditions must be met to expand a LUN:

Note: SGI does not support LUN Expansion and MORE operations on the IRIX and Linux operating systems

1. The most important rule is that you must be using a file manager such as Veritas' Volume Manager. These file managers allow you to actually use an expanded LUN. If you wish to attempt an on-line expansion, and your environment is not properly configured, the request will fail, and you will receive an appropriate error message.
2. During the expansion, which includes re-striping data from the old (smaller) set to the larger set, the controller continues to service host I/O operation.
3. LUN expansion allows you to add capacity to an existing LUN while the controller is online with the host(s). For example, a system using a 6-disk RAID5 set can add another disk to create a 7-disk drive set. This procedure is also referred to by the acronym MORE.

MORE is supported in the simplex mode of operation only. One controller in a dual-active controller system must be disabled (failed over). Attempting to do this operation in a dual-active environment will be rejected.

4. One to six drives can be added to a set at a time. The maximum number of physical disk drives is eight with 6.14 controller firmware, or 16 with 7.01 and later firmware.
5. You can not have eight LUNs defined with 6.14 controller firmware or 32 LUNs defined with 7.01 and later controller firmware. The expansion requires one free LUN in order to execute.
6. You cannot perform an expansion if all eight LUNs are defined with 6.14 controller firmware or all 32 LUNs defined with 7.01 and later controller firmware.
7. The disk drives being added must be in STANDBY, and must not be part of an array.
8. The capacity of each of the added disk drives must be greater than or equal to the size of the smallest disk drive in the set.
9. The added capacity parameters are kept in non-volatile memory. In the event of power loss to this host or to the subsystem, (or a controller failure) the process automatically resumes when power is restored, or the controller is replaced.

10. In the event of a disk drive failure, the process continues to completion in CRITICAL mode.
11. MORE, Initialize, Rebuild, and Consistency Check are mutually exclusive operations. Only one process may run at a time.
12. No configuration update commands will be accepted during the expansion.
13. Write-back cache is disabled during the expansion.
14. The LUN must be online.

To expand a LUN, follow these steps:

1. Click **Expand LUN** in the **Create/Destroy/Expand Logical Drives** screen (see Figure 3-33 for 6.14 firmware, and Figure 3-34 for 7.01 and later firmware).

This section lets you set up new configurations, or add/delete/expand logical drives in existing configurations. You will be walked through each choice, and the configurator will not change anything until specifically instructed to do so. Click [here](#) for a help screen which describes your RAID options in detail, and provides you with relative capacity, availability, and performance characteristics of each type. Or chose one of the buttons below:

<p>New Configuration</p> <p>This creates a new disk configuration. The configurator will prevent you from running this if any LUNs are already defined. (You must perform the Delete LUN operation repeatedly until all LUNs have been removed.)</p>	<p>New Configuration</p>
<p>Add a LUN - Retains ALL Data</p> <p>This lets you use free disk space to create an additional logical RAID disk (LUN). Existing data will not be affected. You must have unconfigured disks available, and at least one LUN must already exist.</p>	<p>Add LUN(s)</p>
<p>Delete last LUN - Retains ALL Data (Except that of the LUN which will be destroyed).</p> <p>This will let you delete the last LUN which was defined. (Limitations with the current firmware make it impossible to delete anything but the last LUN.) If you click on the Delete LUN, then the next screen will describe it's characteristics, and give you the chance to get out cancel the procedure. Note: The configurator will not attempt to determine if there is a mounted filesystem on the LUN.</p>	<p>Delete LUN</p>
<p>Expand a LUN - Retains ALL Data</p> <p>This lets you add capacity to an existing LUN while the controller is online with the host(s). For example, a system using a 6-disk RAID5 set can add another disk to create a 7-disk drive set. This procedure is also referred to by the acronym, "ORE", for On-Line Raid Expansion.</p> <ul style="list-style-type: none"> • During the expansion, which includes re-striping data from the old (smaller) set to the larger set, the controller continues to service host I/O operations • MORE is supported in the simplex mode of operation only. One controller in a dual-active controller system must be disabled (failed-over). Attempting to do this operation in a dual-active environment will be rejected. • One to six drives can be added to a set at a time. The maximum number of physical disk drives is eight. • You can not have 8 LUNs defined. The expansion requires one free LUN in order to execute. • You can not perform an expansion if all 8 LUNs are defined. • The disk drives being added must not be part of an array. • The capacity of each of the added disk drives must be greater than or equal to the size of the smallest disk drive in the set. • The added capacity parameters are kept in non-volatile memory. In the event of power loss to this host or to the subsystem, (or a controller failure) then the process will automatically resume when power is restored, or the controller is replaced. • In the event of a disk drive failure, the process continues to completion in CRITICAL mode. 	<p>Expand LUN</p>

Figure 3-33 Create/Destroy/Expand Logical Drives Screen (6.14 Firmware)

This section lets you set up new configurations, or add/delete/expand logical drives in existing configurations. You will be walked through each choice, and the configurator will not change anything until specifically instructed to do so. Click [here](#) for a help screen which describes your RAID options in detail, and provides you with relative capacity, availability, and performance characteristics of each each type. Or chose one of the buttons below:

<p>New Configuration</p> <p>This creates a new disk configuration. <i>(All of your data will be destroyed in the process!)</i></p>	<input type="button" value="New Configuration"/>
<p>Add a LUN - Retains ALL Data</p> <p>This lets you use free disk space to create an additional logical RAID disk (LUN). Existing data will not be affected. You must have unconfigured disks available, and at least one LUN must already exist.</p>	<input type="button" value="Add LUN(s)"/>
<p>Delete last LUN - Retains ALL Data (Except that of the LUN which will be destroyed).</p> <p>This will let you delete the last LUN which was defined. (Limitations with the current firmware make it impossible to delete anything but the last LUN.) If you click on the Delete LUN, then the next screen will describe it's characteristics, and give you the chance to get out cancel the procedure. Note: The configurator will <i>not</i> attempt to determine if there is a mounted filesystem on the LUN.</p>	<input type="button" value="Delete LUN"/>
<p>Expand a LUN - Retains ALL Data</p> <p>This lets you add capacity to an existing LUN while the controller is online with the host(s). For example, a system using a 6-disk RAID5 set can add another disk to create a 7-disk drive set. This procedure is also referred to by the acronym, "ORE", for On-Line Raid Expansion..</p> <ul style="list-style-type: none"> • During the expansion, which includes re-striping data from the old (smaller) set to the larger set, the controller continues to service host I/O operations • MORE is supported in the simplex mode of operation only. One controller in a dual-active controller system must be disabled (failed-over). Attempting to do this operation in a dual-active environment will be rejected. • One to six drives can be added to a set at a time. The maximum number of physical disk drives is sixteen. • You can not have 16 LUNs defined. The expansion requires one free LUN in order to execute. • You can not perform an expansion if all 16 LUNs are defined. • There must be no SPANNED LUNs defined. • The disk drives being added must not be part of an array. • The capacity of each of the added disk drives must be greater than or equal to the size of the smallest disk drive in the set. • The added capacity parameters are kept in non-volatile memory. In the event of power loss to this host or to the subsystem, (or a controller failure) then the process will automatically resume when power is restored, or the controller is replaced. • In the event of a disk drive failure, the process continues to completion in CRITICAL mode. • ORE, Initialize, Rebuild, and Consistency Check are mutually exclusive operations. Only one process may run at a time. • No configuration update commands will be accepted during the expansion. • Write-back cache is disabled during the expansion. • The LUN must be on-line 	<input type="button" value="Expand LUN"/>

Figure 3-34 Create/Destroy/Expand Logical Drives Screen (7.01 and Later Firmware)

A screen similar to that shown in Figure 3-35 appears.

RAID Array Configuration for system at **/hw/scsl/sc3d010**

The LUNs listed below will support on-line expansion:

Drive Pack	Drive 0 Ch-ID	Drive 1 Ch-ID	Drive 2 Ch-ID	Drive 3 Ch-ID	Drive 4 Ch-ID	Drive 5 Ch-ID	Drive 6 Ch-ID	Drive 7 Ch-ID	Drive 8 Ch-ID	Drive 9 Ch-ID	Drive 10 Ch-ID	Drive 11 Ch-ID	Drive 12 Ch-ID	Drive 13 Ch-ID	Drive 14 Ch-ID	Drive 15 Ch-ID	Physical MB	LUN Number	RAID Type
A	17560MB 0-125	17560MB 1-124	17560MB 0-123	17560MB 1-122													70242	0	RAID 0 (Stripe)
B	17560MB 0-121	17560MB 1-120	17560MB 0-119	17560MB 1-118	17560MB 0-117	Next Assign											87803	1	RAID 0+1 (Mirrored Stripe)

RAID Subsystems:

Enclosure #0 (Rackmount view, rotate 90° clockwise if in tower)			
17560MB 0-125	17560MB 1-124	17560MB 0-123	17560MB 1-122
17560MB 0-121	17560MB 1-120	17560MB 0-119	17560MB 1-118
17560MB 0-117	17560MB 1-116 <input type="checkbox"/>	17560MB 0-115 <input type="checkbox"/>	17560MB 1-114 <input type="checkbox"/>

Select Another Pack EXPAND Selected LUN Select Marked Cancel

[\[Return to Main Page\]](#) [\[Log Off\]](#)

Figure 3-35 Expand LUN Screen

Note: In the screen shown in Figure 3-35, one disk drive has just been added to drive pack A, leaving three unused disk drives that can still be used to expand a LUN.

The drives shown in green are in an online state, which means they are configured into LUNs which happen to be online.

- When there are drives shown in the RAID Subsystems Map area of the screen with an empty checkbox visible, the drives are unused and can be assigned to expand a LUN. Click the **Select Another Pack** button until the cell labeled **Next Assign** is in the desired drive pack that is being expanded. Then activate the desired checkboxes and click **Select Marked** to add the drive(s) to the pack.

When you expand the LUN, the drive pack to which the drive has just been added will still be assigned to the LUN, but will have more capacity, thus expanding the LUN.

Note: If a LUN does not qualify for expansion, it will not be displayed.

- When you have finished adding drives to a particular pack, use the **Select Another Pack** button to select and expand another pack, or click **EXPAND Selected LUN** button to finish the LUN expansion process.

If you click **EXPAND Selected LUN**, the warning message shown in Figure 3-36 appears.

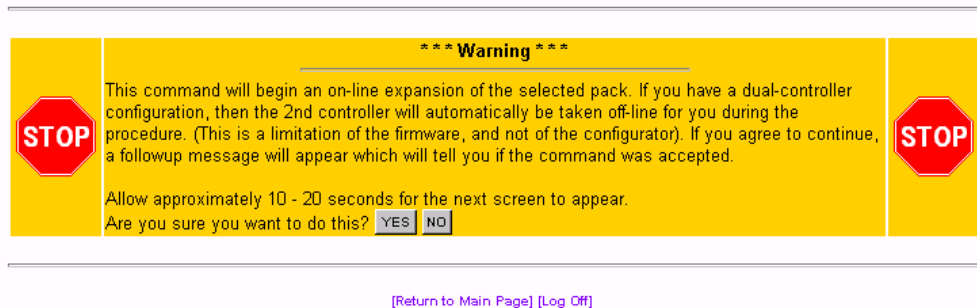


Figure 3-36 Warning Message

- Click **YES** to begin the LUN expansion, or **NO** to cancel and return to the main menu.

Initialize LUN(s) (Configuration Function Menu)




The LUNs must be initialized before the operating system can use them. The TPM application automatically selects uninitialized LUNs for you and displays their status (INITIALIZED, UNINITIALIZED, or INITIALIZING).

To initialize one or more LUNs, follow these steps:

1. Click **Initialize LUNs** on the main menu.

The screen shown in Figure 3-37 appears for 6.14, 7.01, and 7.03 firmware. The screen shown in Figure 3-38 appears for 7.75 firmware.

(Initialize LUN) Logical Disk Information for RAID Subsystem at */lwr/scsi/sc13d010*

Logical Drive #0 (70120 MB, 73,526,149,120 Bytes)	RAID 0 (Stripe) 
Affinity Map: SAN Mapping	Initialize? <input type="checkbox"/> UNINITIALIZED
Logical Drive #1 (35060 MB, 36,763,074,560 Bytes)	RAID 0+1 (Mirrored Stripe) 
Affinity Map: SAN Mapping	Initialize? <input type="checkbox"/> UNINITIALIZED
Logical Drive #2 (35060 MB, 36,763,074,560 Bytes)	RAID 5 (Right Asymmetric) 
Affinity Map: SAN Mapping	Initialize? <input type="checkbox"/> UNINITIALIZED

Enter "YES" in the fields corresponding to the LUN(s) you wish to initialize. You may select as many LUNs as you desire. Then press the **Initialize** button to immediately begin the process, or press the **Cancel** button. Initialization will run in the background, and the configurator will immediately taken you to the background job status screen.

[\[Return to Main Page\]](#)
[\[Log Off\]](#)

Figure 3-37 Initialize LUN (6.14, 7.01, and 7.03 Firmware)

(Initialize LUN) Logical Disk Information for RAID Subsystem at
/hw/scsi/2000080e5115f52/lun0/c3p1

Logical Drive #0 (34332 MB, 35,999,711,232 Bytes)	RAID 3 (Right Asymmetric)	
Affinity Map: SAN Mapping	Initialize? <input type="text" value="NO"/> (UNINITIALIZED)	Background Init? <input type="radio"/>
Logical Drive #1 (34332 MB, 35,999,711,232 Bytes)	RAID 5 (Right Asymmetric)	
Affinity Map: SAN Mapping	Initialize? <input type="text" value="NO"/> (UNINITIALIZED)	Background Init? <input type="radio"/>
Logical Drive #2 (34332 MB, 35,999,711,232 Bytes)	RAID 0+1 (Mirrored Stripe)	
Affinity Map: SAN Mapping	Initialize? <input type="text" value="NO"/> (UNINITIALIZED)	Background Init? <input type="radio"/>
Logical Drive #3 (17166 MB, 17,999,855,616 Bytes)	RAID 1 (Mirror)	
Affinity Map: SAN Mapping	Initialize? <input type="text" value="NO"/> (UNINITIALIZED)	Background Init? <input type="radio"/>

Enter "YES" in the fields corresponding to the LUN(s) you wish to initialize. You may select as many LUNs as you desire. Then press the **Initialize** button to immediately begin the process, or press the **Cancel** button. Initialization will run in the background, and the configurator will immediately taken you to the background job status screen.

Figure 3-38 Initialize LUN (7.75 Firmware)

Note: Do not initialize LUNs in the background from this screen. Instead, Click **Initialize LUNs in Background** on the main menu (see “Initialize LUNs in Background (Configuration Function Menu)” on page 73).

The “Background Init” button shown in this screen is a radio button, indicating that only one LUN can be selected at a time for background initialization. The “Initialize LUNs in Background” selection on the main menu must be used to background initialize multiple LUNs.

As shown in Figure 3-37, several LUNs are shown as **UNINITIALIZED**.

2. Type **YES** in the boxes corresponding to the devices to initialize, and click **Initialize** at the bottom of the screen.

The initialization confirmation screen shown in Figure 3-39 appears.

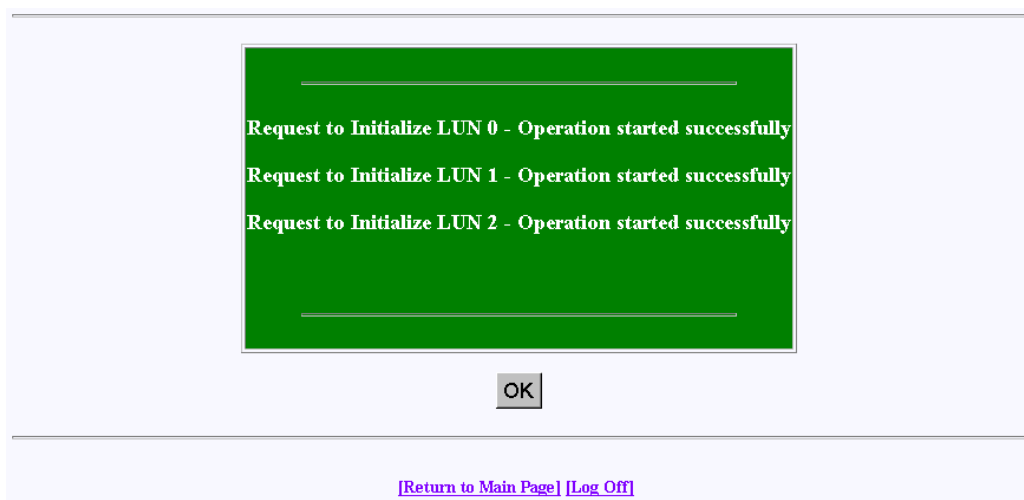


Figure 3-39 Initialization Confirmation

3. To continue with the initialization of the LUNs shown on the screen, click **OK**.

A status screen similar to the one shown in Figure 3-41 (for controller 6.14 firmware) or in Figure 3-41 (for controller 7.01 and later firmware) displays how the procedure is progressing, and updates every 10 seconds.

Background Job Status RAID Controller at /hw/scsi/sc3d1410


Rebuild	Not in progress
Consistency Check	Not in progress.
Initialization Status	LUN 0 In progress 0.1 % complete.
	LUN 1 In progress 0.2 % complete.
	LUN 2 In progress 0.1 % complete.
	LUN 3 In progress 0.0 % complete.
	LUN 4 Not in progress.
	LUN 5 Not in progress.
	LUN 6 Not in progress.
	LUN 7 Not in progress.

This screen will automatically refresh in approximately 10 seconds


[\[Return to Main Page\]](#) [\[Log Off\]](#)

Figure 3-40 Initialization Progress (6.14 Firmware)


Background Job Status RAID Controller at /lwr/scsi/sc13d010

Logical Drive #0 (70120 MB, 73,526,149,120 Bytes) RAID 0 (Stripe) 

Action	Notes
Initialization <input type="button" value="Cancel All"/>	In progress 0.3 % complete.

Logical Drive #1 (35060 MB, 36,763,074,560 Bytes) RAID 0+1 (Mirrored Stripe) 

Action	Notes
Rebuild	Not in progress
Consistency Check	Not in progress
Initialization <input type="button" value="Cancel All"/>	In progress 0.6 % complete.
Online RAID Expansion	Not in progress

Logical Drive #2 (35060 MB, 36,763,074,560 Bytes) RAID 5 (Right Asymmetric) 

Action	Notes
Rebuild	Not in progress
Consistency Check	Not in progress
Initialization <input type="button" value="Cancel All"/>	In progress 0.7 % complete.
Online RAID Expansion	Not in progress

This screen will automatically refresh in approximately 10 seconds

[\[Return to Main Page\]](#) [\[Log Off\]](#)

Figure 3-41 Initialization Progress (7.01 and Later Firmware)

The initialization procedure can be done at any time, and the controller(s) will service I/Os for other LUNs while this is in process. There is also a configurable parameter in the controller configuration that allows you to adjust how much controller CPU time to allow for background operations such as this.

Initialize LUNs in Background (Configuration Function Menu)

Background initialization makes uninitialized system drives consistent by setting the parity while allowing the host to have instantly available read and write access to the system drive. Background initialization is implemented only in version 7.75 and later firmware.

Note: Background initialization can only be performed on valid RAID levels (RAID levels 1, 3, 5 and 0 +1)

Background initialization is enabled by default. If the "Enable Background Init" controller parameter is changed and set to disabled, "Initialize LUNs in Background" will not be displayed under "Configuration Functions". The option to perform a background initialization is still provided under the "Initialize LUN(s)" function. However, only one LUN can be selected for background initialization at a time using this function. When performing "Background Init" on multiple LUNs, change the "Enable Background Init" parameter under the View/Modify RAID Controller Configuration. See Figure 3-56 on page 93.

To initialize LUNs in the background, follow these steps:

1. Click **Initialize LUNs in Background** on the main menu.

The screen shown in Figure 3-42 appears.



Figure 3-42 Initialize LUNs in Background (7.75 Firmware)

2. Click **INITIALIZE**.

The screen shown in appears.




Figure 3-43 Initialize LUNs in Background (7.75 Firmware)

3. Click **OK**.

Background initialization now begins and initializes any uninitialized system drives one at a time, and is paused by any of the following operations:

- Foreground Initialization
- Consistency Check and Restore
- Rebuild

If one of these operations is started while background initialization is executing, background initialization is paused until the interrupting operation is complete. Once paused, background initialization will only continue after a write operation to the LUN. The same is true of a controller reset.

Note: Performance is degraded during background initialization because every write requires access to all drives in the RAID group. Sites that require optimal performance should take this into account and initialize LUNs in the foreground. Published performance levels are not guaranteed when background initialization are in process.

Configure Host-to-LUN Mapping (Configuration Function Menu)

The **Configure Host -to-LUN** item on the **Configuration Functions** portion of the main menu (see Figure 3-1 on page 22) allows logical drives to be made available or invisible to selected Fibre Channel host adapters.

RAID controllers offer several drive mapping techniques and configuration modes for many different environments. These configuration modes define which hosts see the logical disks, and the effects of a hardware failure. When you select **Configure Host to LUN Mapping** from the main menu, the dialog box shown in Figure 3-44 (6.14 firmware) or in Figure 3-46 (7.01 and later firmware) appears.

<p>This section implements Host to LUN Mapping by allowing you to select a technique by which LUNs will be presented to computers. Carefully choose from below, and note that no changes will be made until you specifically tell the configurator to SAVE one. Also be sure not to have any mounted filesystems on LUNs if your new mapping makes them either disappear, or changes the LUN number that is reported to your host O/S.</p>	
<p>Affinity LUN Map</p> <p>This option lets you assign what controllers and ports see each of your logical drives. The LUNs which will be presented to your O/S will be automatically assigned.</p>	<input type="button" value="Affinity MAP"/>
<p>User-Supplied LUN Map</p> <p>This allows you to assign a specific SCSI/FibreChannel ID's for every logical drive that is configured into the array. You have the flexibility to assign them by specific controller/port numbers as well.</p>	<input type="button" value="User-Supplied LUN Map"/>
<p>SAN Map - (CURRENT)</p> <p><i>(Requires firmware version 6.0 or above)</i> Select this option to define which LUNs will be seen by specific Fibre-channel host adapters. This is extremely secure, and a must for multi-host Storage Area Networks.</p>	<input type="button" value="SAN MAP"/>
<p>[Return to Main Page] [Log Off]</p>	

Figure 3-44 Configure Host to LUN Mapping (6.14 Firmware)

The mapping strategy currently used by the controller is indicated with (CURRENT). In the example of Figure 3-44, the SAN Map is the current mapping strategy.

Note: SAN MAP is the only mapping strategy supported by the TP9100 RAID system.

Click **SAN MAP** to make any modifications to the mapping strategy. No changes are made until you complete the subsequent dialog boxes.

Heed these important warnings:

Warning: A reset is required if you change the mapping type (for example, from SAN to Affinity), or the topology type within SAN mapping. You can, however, set things up and postpone the reset until you have some downtime available.

Warning: If you make a change to the mapping, be sure to consider how the new mapping might affect device drivers on attached systems. Depending on what you are changing, you might make the controller invisible to TPM. You also might make one of your host device drivers now point to the wrong LUN, which could result in data loss if the LUN is in use. In other words, TPM will not stand in your way if you do something to configure the system in an undesirable manner.

Warning: Only the most experienced administrators should make changes in an on-line environment.

SAN Mapping

The Storage Area Network (SAN) Mapping feature, also known as Host-to-LUN Mapping feature, restricts host access to logical drives. Each drive is granted only to a single host or group of hosts, providing limited security control of data in an environment where multiple hosts are connected to the controller.

This feature is available only on controllers with a fibre channel host interface running on firmware version 6.0 or greater.

The SAN mapping feature is intended for use in configurations in which multiple host computers attach to one or more controllers. This is also referred to as a SAN configuration. The host computers are attached to the controller(s) through a fibre channel arbitrated loop, FC hub, or FC switch. An example of fibre channel arbitrated loop configuration is shown in Figure 3-45.

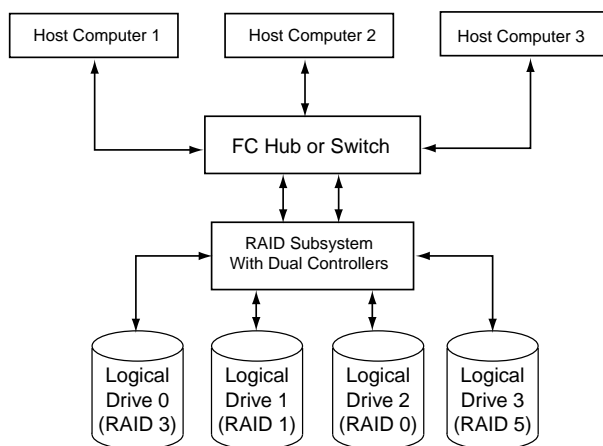


Figure 3-45 Storage Area Network

Without host to LUN mapping, each host computer (1 through 3) has complete access to all four system drives. When a host system boots, operating systems such as Windows NT attempt to mount all of these system drives. As a result, any data on the array can be accessed by the Windows NT system. In some cases, Windows NT automatically writes an identifying signature to these storage areas. This results in data corruption if another one of the host systems has stored data there.

If IRIX or Linux computers are involved, then an automatic mount might not occur, but you still have a security risk, and the possibility that a user with root privilege will mount one of these disks, or attempt to create a file system on one of these disks.

By utilizing SAN mapping, however, each logical drive can be configured to be visible to a single host computer only. If you are using a volume mapping tool such as Veritas' Volume Manager, or Tivoli's (previously Mercury's) SANergy product, then you can safely concurrently mount file systems on these LUNs to any number of these three hosts.

The controller uses the World Wide Name (WWN) to uniquely identify host computers that have logged in to the controller. A list of valid hosts and their corresponding WWNs, is provided to external configurators in order to configure the mapping.

After logical drives are configured, the controller maintains a table of WWNs for each one. This table defines the hosts that are granted access to each system drive and the controller port and the LUN number. The controller uses the table of WWNs to determine access to a specific system drive. If a host sends a new command to the controller, the controller validates the WWN, LUN, and controller port prior to servicing the command. If the WWN, LUN, and port information are valid for the system drive, the requested command is completed normally. If the WWN, LUN, and port combination are not valid for the system drive, the command is completed with SCSI Check Condition status, with the sense key set to Illegal Request (05h) and the sense code set to Logical Unit Not Supported (25h).

There are three exceptions to the response to commands when the WWN, LUN, and port combination are not valid:

1. If the request is an **Inquiry** command, the controller returns the Inquiry data with the peripheral qualifier set to indicate that the target is capable of supporting the specified device type on this LUN, but no device is currently connected to that LUN.
2. If the request is a Report LUNs command, and the addressed LUN is 0, the controller completes the command normally, reporting only the LUNs accessible by the host requesting the command.
3. If the request comes from TPM, however, the command is processed normally by the controller. This allows a controller that is not configured to be reconfigured to operate correctly with the attached hosts.

Supported SAN Mapping Topologies

The model of controller you use dictates what SAN Mapping topologies are supported. This manual covers all three possibilities (Inactive Port, MultiPort, and Multi-TID). Only one type of SAN topology can be active, and any changes to a topology requires a reboot for it to become effective. The topologies are:

- Inactive Port—in this topology, Controller0/Port0 and Controller1/Port1 are active. During failover, the inactive port on the partner takes over for the active port on the failed controller.
- MultiPort¹—in this topology, all ports are active. This topology does *not* provide transparent failover or failback and requires an alternate path driver to the host.
- Multi-TID—in this topology, all ports are active. This topology provides transparent failover and failback, but should not be used in conjunction with an alternate path driver.

This manual is not designed to be a tutorial on the strengths, weaknesses, and required external hardware configuration to use for each topology for the various operating systems. The documentation here merely covers how to configure each of them.

To display the SAN Logical Drive Map for controllers with 6.14 firmware, click the SAN Map button on the Configure Host to LUN Mapping screen (see Figure 3-44 on page 76). The screen shown in Figure 3-46 appears.

Note: For controllers with 7.01 and later firmware, the SAN Logical Drive Map appears when the **Configure Host -to-LUN** item on the **Configuration Functions** portion of the main menu (see Figure 3-1 on page 22) is selected.

¹ The SGI-recommended setting is MultiPort

SAN Logical Drive Map for RAID Subsystem at /hw/scsi/20000080e5115f52/1un0/c3p1

Host World Wide Name	SD #00	SD #01	SD #02
Drive to LUN Mapping	*LUN0	*LUN1	*LUN2
Allow ALL Hosts	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>
20-00-00-E0-8B-00-F3-C4	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>
10-00-00-60-69-20-15-72	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>

[\[Return to Main Page\]](#) [\[Log Off\]](#)

Figure 3-46 SAN Mapping

Regardless of the topology, the bottom portion of the screen shown in Figure 3-46 will be the same. Administrators must first choose the topology that best matches the SAN environment, then configure the mapping. To change the topology, click the **Topology** drop down box, then press the **Select** button. The screen is then changed to contain the appropriate values for the selected topology and controller/port combination.

Note: Your subsystem will probably have multiple controllers and ports. If this is the case, use the buttons in the **Portmap** area of the screen to configure the map for each of them (the topology type is the same for all ports).

After all the changes have been made, click **APPLY** at the bottom of the screen.

The warnings given earlier in this section are repeated here:

Warning: A reset is required if you change the mapping type (for example, from SAN to Affinity), or the topology type within SAN mapping. You can, however, set things up and postpone the reset until you have some downtime available.

Warning: If you make a change to the mapping, be sure to consider how the new mapping might affect device drivers on attached systems. Depending on what you are changing, you might make the controller invisible to TPM. You also might make the device drivers now point to the wrong LUN, which could result in data loss if the LUN is in use. In other words, TPM will not stand in your way if you do something to configure the system in an undesirable manner.

Warning: Only the most experienced administrators should make changes in an on-line environment.

Installing Device Drivers for New LUNS (Configuration Function Menu)

Selecting this item on the main menu (see “Main Menu” on page 21) instructs TPM to tell the host computer to scan for new devices and update device drivers. To initiate this process, click **Install device drivers for new LUNS** on the main menu. The warning screen shown in Figure 3-47 appears.

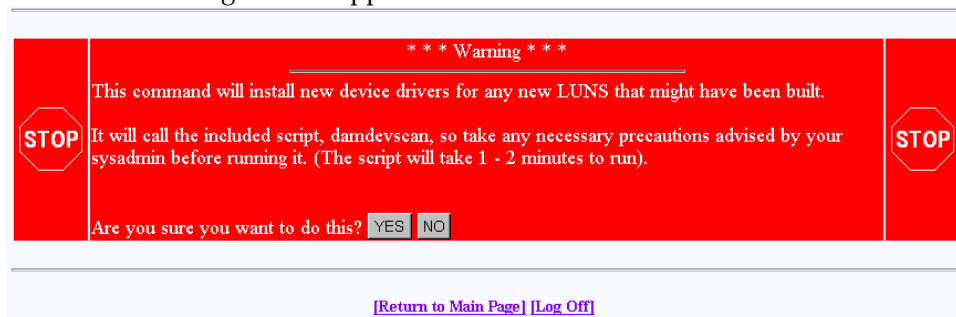


Figure 3-47 Install New Device Drivers Warning Screen

If you are running in a live environment with mounted LUNs, you need to be careful about executing this feature. For example, if you changed SAN mapping so logical drive #3 is mapped to LUN 0 instead of LUN 4, then you will make file systems disappear, and confuse your device drivers. This should be expected, of course. As stated earlier, TPM is not going to prevent you from doing something that would be detrimental to the system configuration.

Also, if you have cross-linked or improperly defined device drivers, running this function corrects the problem, but it also properly redefines drivers. This could also have an effect on mounted devices. Below are some operating system-specific comments:

- Linux: RedHat 6.2 does not support scanning for new devices. A reboot is recommended. However, if you manually remove your FC device driver, then re-install it, new devices are discovered. This saves you the trouble of rebooting, but the technique obviously can't be used if other LUNs are active and mounted. After unloading the FC driver module, perform a reload. You can select new LUNs, which calls the `$DAM_HOME/damdevscan` shell script.
- IRIX: TPM calls the `$DAM_HOME/damdevscan` shell script.

If the new LUNs haven't been discovered, a system reboot may be required in order to have the host recognize the new devices.

- Windows: A reboot is required after creating LUNs.

Send Command to Host (Configuration Function Menu)

Use this function to spawn a command line to the IRIX/Linux host that you are communicating with.

Note: The Windows OS does not support **Send Command to Host**. Therefore, it does not appear on the main menu.

When the program terminates, the resulting page is returned. In addition, any error messages going to **stderr** appear at the top of the result window. The output buffer is truncated to the first 95 KB worth of data.

This command is typically used to perform mounts, but TPM doesn't care what you do. Remember also that the command will be run as root, so please view the caveats above. To run the command, click **Send Command To Host** on the main menu. The screen shown in Figure 3-48 appears.

Send Command To Host

Enter any command below, and you must adhere to the following limitations:

- If the command runs in the background (trailing ampersand), then no return status will be shown on the next screen.
- You must not pipe stdout or stderr. This is the mechanism by which output will be sent to your browser.
- Do not run interactive commands, as there is no facility to run an interactive session using HTML. If this is what you require, then use a telnet session.
- Remember that the job will be run as root, and your environment will be the same as the service routine running in the background.

Command:

EXECUTE **Cancel**

[\[Return to Main Page\]](#) [\[Log Off\]](#)

Figure 3-48 Send Command to Host

Note: Use only non-interactive commands (those that do not require a user response).

Type the appropriate command, then click **EXECUTE**.

Administrative Functions

The Administrative Functions menu is located on the main screen and is shown in Figure 3-49.

Administrative Functions:

- [View/Modify RAID controller configuration](#) - Although many changes may be made on-the-fly, some settings **must** be made before any RAID groups are defined.
- [Reset controller\(s\)](#) - This simultaneously cold resets all controllers in a subsystem. Multiple attempts are made for 90 seconds, in case the controllers have active I/Os.
- [Gracefully bring a controller off-line](#).
- [Gracefully bring a 2nd controller on-line](#) - Do this after a failed controller has been removed, or you are upgrading from a simplex configuration to a dual-controller configuration.
- [Perform data consistency check/restore on a LUN](#) - You should perform a data consistency check regularly for all redundant LUNs.
- [Enable/Disable write cache for LUN\(s\)](#).

Figure 3-49 Administrative Functions Menu

This menu has the following selections, with the associated explanations on the indicated pages:

- “View/Modify RAID Controller Configuration” on page 86
- “Reset Controller(s)” on page 95
- “Gracefully Bring a Controller Off Line” on page 96
- “Gracefully Bring a 2nd Controller On Line” on page 97
- “Perform Data Consistency Operations on a LUN” on page 98
- “Enable/Disable Write Cache for LUN(s)” on page 103
- “Modify/Purge Non-Volatile WWN Tables” on page 104

Note: The Modify/Purge Non-Volatile WWN Tables feature is new with firmware version 7.75.

View/Modify RAID Controller Configuration

This menu selection lets you make changes to your controller. The most important thing to know is that changes labeled **On-the-fly** are immediate. Those marked **Reset** require a controller reset to become effective. The two choices labeled **New Config** can only be made when there are no defined LUNs, as they are data-destructive. When you click **View/Modify RAID controller configuration**, the dialog boxes appear as shown in Figure 3-50 and Figure 3-51 (for controllers with 6.14 firmware), as shown in Figure 3-52 and Figure 3-53 (for controllers with 7.01 firmware), and as shown in Figure 3-54 through Figure 3-56 (for controllers with 7.75 firmware). A partial dialog box is shown in each figure.

DAC960FFx Controller Configuration (Firmware Type: 5 Rev 6.14 Build #0) 128 MB Cache, Max #LUNs=8			
Current	Value	Category	Description
<input type="checkbox"/>	Auto Rebuild Management	On-the-fly	If enabled , it detects the replacement of a failed drive and performs an automatic rebuild once it has spun up, provided it is installed into a redundant array (RAID 1, RAID 3, RAID 5, RAID 6+L). If this feature is disabled , the administrator must issue the rebuild command manually through this configurator.
<input type="checkbox"/>	Operational Fault Management	Reset	Allows the controller to take autonomous actions when a failure occurs. This monitors and reports drive failures, background activity completion status, enclosure events, etc. This should remain enabled during normal controller operation. <i>(This is also known as SES, or SCSI Enclosure Services)</i>
16 KB	Stripe Size	New Config	This is the stripe size used by the controller to allocate cache lines in memory and data striping on disk. The value must not change once you have created LUNs as data loss will result. Also, certain firmware revisions won't even accept the command if any LUNs exist.
<input checked="" type="checkbox"/>	Read Ahead	Reset	The controller extends commands to the corresponding <i>stripe unit</i> size. The controller reads data from disk in chunks of one stripe-unit size. Given an 8KB stripe size, a 2KB read, for example, results in 8KB read being issued to the drive. The remaining 6KB of data stays in the cache.
<input type="checkbox"/>	Super Read Ahead	Reset	The controller extends the read-ahead algorithm by always reading an extra cache line on a read request, and reading a further cache line when a cache hit occurs on a pre-fetched cache line. This is primarily useful for applications with a high degree of sequential-ness.
<input type="checkbox"/>	Reassign limited to 1 sector	Reset	If enabled, reassigns will be restricted to only one block, the failing block. If this is disabled all reassigns will be for the entire current I/O, some possibly large number of blocks, not all of them failing. The single block reassign is further limited to recovered errors and medium errors.
<input type="checkbox"/>	True Verify	Reset	When enabled, and if the host enables verify on an I/O operation, and data is transferred, a true verify with data comparison is performed. When disabled, no data comparison is made.
<input type="checkbox"/>	Disk Write Through Verify	Reset	During error handling, this turns on Force Unit Access for reads and writes.
<input type="checkbox"/>	RAID-5 Algorithm Control	New Config	If disabled, the algorithm is right asymmetric, if enabled, left asymmetric. (It is doubtful you will ever see a performance difference, regardless of the setting).
5	Rebuild/Check Consistency Rate Default	On-the-fly	This value times 2 approximates the percentage of available rebuild cycles to be used when rebuilding a RAID group, or checking consistency. CPU utilization is always shared with data traffic. Range 0-50. Therefore, a value of 50 devotes the maximum allowable resources to a drive rebuild or expansion, allowing it to proceed at its fastest. A lower number provides more resources to service other I/Os.
64	Disk Queue Limit	On-the-fly	Sets the maximum allowed queue depth for tagged commands to each attached drive. This is further limited to the drive's own tag limit when that limit is reached. Do not change this value unless specifically directed to do so (The valid range is 1-230).
<input checked="" type="checkbox"/>	Queuing - Coalescing Optimization	On-the-fly	If enabled, this will join the data from adjacent I/Os into a single I/O to improve performance.
<input type="checkbox"/>	Enable On queue full give BUSY	Reset	Any time a command is received and the controller detects a queue full condition, it will normally return Queue Full status. If enabled, a queue full status will return a BUSY status, if disabled, it will return QUEUE FULL . This is intended to help hosts that are confused by QUEUE FULL .
<input checked="" type="checkbox"/>	Failover Node Name Retention	Reset	If enabled, each controller shares its node name with its partner controller and those names are used through all phases of failover and failback. If disabled, each controller still shares its node name with its partner controller, and those names are still used through all phases of a failover, BUT when a failback occurs the replacement controller uses its own node name. Not having this feature enabled will have serious ramifications if the controllers are connected to a host that uses node names to locate the LUNs.
<input checked="" type="checkbox"/>	SAF-TE Data for UPS Support	On-the-fly	If checked, then ups monitoring is disabled.
24	Max Number of Active Commands	Reset	Sets the maximum number of simultaneous commands. Range is (64-244).
<input checked="" type="checkbox"/>	Enable DEBUG Port	Reset	If checked, then debug output will go to the serial port. This is for Engineering and Diagnostic purposes only, and will result in a performance loss if enabled. If not checked, then the port will be in SLP (no jumper) or VT100 mode (jumper - Standard Mode).

Figure 3-50 View/Modify Controller Configuration (partial for 6.14 Firmware)

<input type="checkbox"/>	Enable Vendor Unique TUR	Reset	If checked, a Test Unit Ready command sent from the host to an off-line LUN will return a hard error status (4/00/00). If clear, then it will return a not ready status (2/04/03).
<input checked="" type="checkbox"/>	Disable Check Condition for Invalid LUN	Reset	If checked, the inquiry command will return data with the peripheral qualifier field set to 1 or 0x20 for the byte meaning peripheral not connected. If clear, the inquiry will fail with check condition of illegal request (sense=5/25/00).
<input type="checkbox"/>	Disable Pause when Not Ready	Reset	Normally, when controller is starting up, certain commands encounter a brief pause. If this field is checked, the pause is disabled, otherwise, it is enabled.
<input type="checkbox"/>	Disable BUSY status on fallback	Reset	During fallback, the survivor controller normally returns BUSY to new commands received from the host during the cache flush operation. If checked, requests are ignored. If unchecked, BUSY status is returned. This feature is intended to help hosts that are confused by a BUSY.
19200	Serial Port BAUD Rate	Reset	The baud rate of the serial port when in VT100 or Debug modes.
<input checked="" type="checkbox"/>	Force Simplex Mode	Reset	This is provided to allow duplex firmware to serve in a simplex environment, and makes duplex firmware skip some of the active-active operations. If you check this field in a duplex environment, then one of the controllers will go into reset.
<input type="checkbox"/>	Conservative Cache Mode	On-the-fly	This provides an extra degree of data safety when operating in failed over condition. This turns off write cache while the failed over condition persists. Checking this field enables conservative cache mode.
<input type="checkbox"/>	Simplex Disable Reset	Reset	For simplex-only. If checked, prevents a controller from asserting the reset signal to the partner controller.
2048 Bytes	Fibre Channel Frame Control	Reset	Provided to allow adjustment of the FC chip's frame size. Unless you are almost always doing very small block I/O, then you would want this set to 2048 .
512 Bytes	PCI Latency Control	Reset	Provided to allow adjustment of the FC chip's use of the internal bus. It controls the amount of data each FC processor chip can burst across the primary bus before relinquishing bus ownership to the next device. This takes effect only when all internal FC bus ports are active and arbitrating. Ordinarily you should leave this set to the factory default value of 512 . If you are in a high throughput environment, then you may see a slight performance advantage if you change the value to 2048 .
<input checked="" type="checkbox"/>	Enable Hard Loop IDs	Reset	This allows you to force the hard loop ID's for every controller and port. The fields below are used to assign them. Do not enable this feature unless you set appropriate values for your controllers. Changes will not go into effect until you issue a reset or recycle power to the array.
Ctrl0 Port0: 10 Ctrl1 Port0: 0	Hard Loop IDs for each controller and port	Reset	Allows option of using the same loop IDs all the time. Some fibre HBAs require non-default settings here. Contact your supplier for details. The valid range is 0-125 (0 - 0x7d). Enter a decimal number. Do not adjust these numbers unless instructed to do so by your vendor's technical support team.
<input checked="" type="checkbox"/>	Smart Large Host Transfers	On-the-fly	If checked, large transfers will coalesce into fewer I/Os, which means fewer disconnects on large transfers. This only takes effect for transfers larger than the stripe size. If enabled, you will have a slightly higher throughput, at a possible cost of some IOPs.
Automatic	Spin-Up Settings	On-the-fly	<input type="checkbox"/> # of Devices to spin up at a time. Ignored for "On Power" mode. <input type="checkbox"/> # of seconds to delay. For "Automatic", time between disk spin-up cycles. For "On Power", time before issuing start-up commands. For "On Command", time between disk spin-up cycles. <input type="checkbox"/> # of seconds between subsequent spin-ups. Only applicable to "On Power"

Notes:

- Settings characterized as **Reset** require that the controllers be reset before the new settings are invoked.
- Settings characterized as **On-the-fly** may be made at any time, and will become effective immediately.
- Settings characterized as **New Config** must be made during initial configuration, before any LUNs are created. **They will result in data loss of existing RAID groups.**

Figure 3-51 View/Modify Controller Configuration (partial for 6.14 Firmware)

DAC960FFx Controller Configuration (Firmware Type 7 Rev 7.01 Build #0) 128 MB Cache, Max #LUNs=32			
Current	Value	Category	Description
<input checked="" type="checkbox"/>	Auto Rebuild Management	On-the-fly	If enabled, it detects the replacement of a failed drive and performs an automatic rebuild once it has spun up, provided it is installed into a redundant array (RAID 1, RAID 2, RAID 5, RAID 6+1). If this feature is disabled, the administrator must issue the rebuild command manually through this configurator.
<input checked="" type="checkbox"/>	Operational Fault Management	Reset	Allows the controller to take autonomous actions when a failure occurs. This monitors and reports drive failures, background activity completion status, enclosure events, etc. This should remain enabled during normal controller operation. <i>(This is also known as SES, or SCSI Enclosure Services)</i>
<input checked="" type="checkbox"/>	Auto Failback	Reset	Allows the surviving controller to automatically sense and place an inserted replacement controller back in service.
<input checked="" type="checkbox"/>	Read Ahead	Reset	The controller extends commands to the corresponding <i>stripe unit</i> size. The controller reads data from disk in chunks of one stripe-unit size. Given an 8KB stripe size, a 2KB read, for example, results in 8KB read being issued to the drive. The remaining 6KB of data stays in the cache.
<input type="checkbox"/>	Super Read Ahead	Reset	The controller extends the read-ahead algorithm by always reading an extra cache line on a read request, and reading a further cache line when a cache hit occurs on a pre-fetched cache line. This is primarily useful for applications with a high degree of sequential-ness.
<input checked="" type="checkbox"/>	Reassign limited to 1 sector	Reset	If enabled, reassigns will be restricted to only one block, the failing block. If this is disabled all reassigns will be for the entire current I/O, some possibly large number of blocks, not all of them failing. The single block reassign is further limited to recovered errors and medium errors.
<input type="checkbox"/>	True Verify	Reset	When enabled, and if the host enables verify on an I/O operation, and data is transferred, a true verify with data comparison is performed. When disabled, no data comparison is made.
<input type="checkbox"/>	Disk Write Through Verify	Reset	During error handling, this turns on Force Unit Access for reads and writes.
<input type="text" value="5"/>	Rebuild/Check Consistency Rate Default	On-the-fly	This value times 2 approximates the percentage of available rebuild cycles to be used when rebuilding a RAID group, or checking consistency. CPU utilization is always shared with data traffic. Range 0-50. Therefore, a value of 50 devotes the maximum allowable resources to a drive rebuild or expansion, allowing it to proceed at its fastest. A lower number provides more resources to service other I/Os.
<input type="text" value="24"/>	Max Number of Active Commands	Reset	Sets the maximum number of simultaneous commands. Range is (64-244).
<input checked="" type="checkbox"/>	Queueing - Coalescing Optimization	On-the-fly	If enabled, this will join the data from adjacent I/Os into a single I/O to improve performance.
<input type="checkbox"/>	Enable On queue full give BUSY	Reset	Any time a command is received and the controller detects a queue full condition, it will normally return Queue Full status. If enabled, a queue full status will return a BUSY status, if disabled, it will return QUEUE FULL. This is intended to help hosts that are confused by QUEUE FULL.
<input type="checkbox"/>	Failover Node Name Retention	Reset	If enabled, each controller shares its node name with its partner controller and those names are used through all phases of failover and failback. If disabled, each controller still shares its node name with its partner controller, and those names are still used through all phases of a failover, BUT when a failback occurs the replacement controller uses its own node name. Not having this feature enabled will have serious ramifications if the controllers are connected to a host that uses node names to locate the LUNs.
<input type="checkbox"/>	SAF-TE Data for UPS Support	On-the-fly	If checked, then ups monitoring is disabled.
<input type="checkbox"/>	Disable Check Condition for Invalid LUN	Reset	If checked, the inquiry command will return data with the peripheral qualifier field set to 1 or 0x20 for the byte meaning peripheral not connected. If clear, the inquiry will fail with check condition of illegal request (sense=5/25/00).
<input checked="" type="checkbox"/>	Disable Pause when Not Ready	Reset	Normally, when controller is starting up, certain commands encounter a brief pause. If this field is checked, the pause is disabled, otherwise, it is enabled.
<input checked="" type="checkbox"/>	Disable BUSY status on failback	Reset	During failback, the survivor controller normally returns BUSY to new commands received from the host during the cache flush operation. If checked, requests are ignored. If unchecked, BUSY status is returned. This feature is intended to help hosts that are confused by a BUSY.

Figure 3-52 View/Modify Controller Configuration (partial for 7.01 Firmware)

<input checked="" type="checkbox"/>	Enable DEBUG Port	Reset	If checked, then debug output will go to the serial port. This is for Engineering and Diagnostic purposes only, and will result in a performance loss if enabled. If not checked, then the port will be in SLP (no jumper) or VT100 mode (jumper - Standard Mode).
<input type="checkbox"/>	Enable Vendor Unique TUR	Reset	If checked, a Test Unit Ready command sent from the host to an off-line LUN will return a hard error status (400/00). If clear, then it will return a not ready status (2/04/03).
19200	Serial Port BAUD Rate	Reset	The baud rate of the serial port when in VT100 or Debug modes.
<input type="checkbox"/>	Force Simplex Mode	Reset	This is provided to allow duplex firmware to serve in a simplex environment, and makes duplex firmware skip some of the active-active operations. If you check this field in a duplex environment, then one of the controllers will go into reset.
<input checked="" type="checkbox"/>	Conservative Cache Mode	On-the-fly	This provides an extra degree of data safety when operating in failed over condition. This turns off write cache while the failed over condition persists. Checking this field enables conservative cache mode.
<input type="checkbox"/>	Simplex Disable Reset	Reset	For simplex-only. If checked, prevents a controller from asserting the reset signal to the partner controller.
2048 Bytes	Fibre Channel Frame Control	Reset	Provided to allow adjustment of the FC chip's frame size. Unless you are almost always doing very small block I/O, then you would want this set to 2048 .
512 Bytes	PCI Latency Control	Reset	Provided to allow adjustment of the FC chip's use of the internal bus. It controls the amount of data each FC processor chip can burst across the primary bus before relinquishing bus ownership to the next device. This takes effect only when all internal FC bus ports are active and arbitrating. Ordinarily you should leave this set to the factory default value of 512 . If you are in a high throughput environment, then you <i>may</i> see a slight performance advantage if you change the value to 2048 .
<input checked="" type="checkbox"/>	Enable Hard Loop IDs	Reset	This allows you to force the hard loop ID's for every controller and port. The fields below are used to assign them. Do not enable this feature unless you set appropriate values for your controllers. Changes will not go into effect until you issue a reset or recycle power to the array.
Ctrl0 Port0 0 Ctrl1 Port0 1	Hard Loop IDs for each controller and port	Reset	Allows option of using the same loop IDs all the time. Some fibre HBAs require non-default settings here. Contact your supplier for details. The valid range is 0-125 (0 - 0x7d). Enter a decimal number. Do not adjust these numbers unless instructed to do so by your vendor's technical support team.
<input checked="" type="checkbox"/>	Smart Large Host Transfers	On-the-fly	If checked, large transfers will coalesce into fewer I/Os, which means fewer disconnects on large transfers. This only takes effect for transfers larger than the stripe size. If enabled, you will have a slightly higher throughput, at a possible cost of some IOPs.
Automatic	Spin-Up Settings	On-the-fly	<input type="text"/> # of Devices to spin up at a time. Ignored for "On Power" mode. <input type="text"/> # of seconds to delay. For "Automatic", time between disk spin-up cycles. For "On Power", time before issuing start-up commands. For "On Command", time between disk spin-up cycles. <input type="text"/> # of seconds between subsequent spin-ups. Only applicable to "On Power"

Notes:

- Settings characterized as **Reset** require that the controllers be reset before the new settings are invoked.
- Settings characterized as **On-the-fly** may be made at any time, and will become effective immediately.
- Settings characterized as **New Config** must be made during initial configuration, before any LUNs are created. **They will result in data loss of existing RAID groups.**

Apply Cancel Reset Screen

[\[Return to Main Page\]](#) [\[Log Off\]](#)

Figure 3-53 View/Modify Controller Configuration (partial for 7.01 Firmware)

DAC960FFx Controller Configuration (Firmware Type 7 Rev 7.70 Build #0) 128 MB Cache, Max #LUNs=32			
Current	Value	Category	Description
<input checked="" type="checkbox"/>	Auto Rebuild Management	On-the-fly	If enabled , it detects the replacement of a failed drive and performs an automatic rebuild once it has spun up, provided it is installed into a redundant array (RAID 1, RAID 3, RAID 5, RAID 0+1). If this feature is disabled , the administrator must issue the rebuild command manually through this configurator.
<input checked="" type="checkbox"/>	Operational Fault Management	Reset	Allows the controller to take autonomous actions when a failure occurs. This monitors and reports drive failures, background activity completion status, enclosure events, etc. This should remain enabled during normal controller operation. (This is also known as SES, or SCSI Enclosure Services)
<input checked="" type="checkbox"/>	Auto Failback	Reset	Allows the surviving controller to automatically sense and place an inserted replacement controller back in service.
<input checked="" type="checkbox"/>	Read Ahead	Reset	The controller extends commands to the corresponding <i>stripe unit</i> size. The controller reads data from disk in chunks of one stripe-unit size. Given an 8KB stripe size, a 2KB read, for example, results in 8KB read being issued to the drive. The remaining 6KB of data stays in the cache.
<input type="checkbox"/>	Super Read Ahead	Reset	The controller extends the read-ahead algorithm by always reading an extra cache line on a read request, and reading a further cache line when a cache hit occurs on a pre-fetched cache line. This is primarily useful for applications with a high degree of sequential access.
<input type="checkbox"/>	Reassign limited to 1 sector	Reset	If enabled, reassigns will be restricted to only one block, the failing block. If this is disabled all reassigns will be for the entire current I/O, some possibly large number of blocks, not all of them failing. The single block reassign is further limited to recovered errors and medium errors.
<input type="checkbox"/>	True Verify	Reset	When enabled, and if the host enables verify on an I/O operation, and data is transferred, a true verify with data comparison is performed. When disabled, no data comparison is made.
<input type="checkbox"/>	Disk Write Through Verify	Reset	During error handling, this turns on Force Unit Access for reads and writes.
<input type="text" value="50"/>	Rebuild/Check Consistency Rate Default	On-the-fly	This value times 2 approximates the percentage of available rebuild cycles to be used when rebuilding a RAID group, or checking consistency. CPU utilization is always shared with data traffic. Range 0-50. Therefore, a value of 50 devotes the maximum allowable resources to a drive rebuild or expansion, allowing it to proceed at its fastest. A lower number provides more resources to service other I/Os.
<input type="text" value="230"/>	Disk Queue Limit	Reset	Sets the maximum allowed queue depth for tagged commands to each attached drive. This is further limited to the drive's own tag limit when that limit is reached. Do not change this value unless specifically directed to do so. Range is (1-230).
<input checked="" type="checkbox"/>	Queuing - Coalescing Optimization	On-the-fly	If enabled, this will join the data from adjacent I/Os into a single I/O to improve performance.

Figure 3-54 View/Modify Controller Configuration (partial for 7.75 Firmware)

<input checked="" type="checkbox"/>	Enable On queue full give BUSY	Reset	Any time a command is received and the controller detects a queue full condition, it will normally return Queue Full status. If enabled, a queue full status will return a BUSY status, if disabled, it will return QUEUE FULL . This is intended to help hosts that are confused by QUEUE FULL.
<input checked="" type="checkbox"/>	Failover Node Name Retention	Reset	If enabled, each controller shares its node name with its partner controller and those names are used through all phases of failover and fallback. If disabled, each controller still shares its node name with its partner controller, and those names are still used through all phases of a failover, BUT when a fallback occurs the replacement controller uses its own node name. Not having this feature enabled will have serious ramifications if the controllers are connected to a host that uses node names to locate the LUNs.
<input type="checkbox"/>	SAF-TE Data for UPS Support	On-the-fly	If checked, then ups monitoring is disabled.
<input checked="" type="checkbox"/>	Disable Check Condition for Invalid LUN	Reset	If checked, the inquiry command will return data with the peripheral qualifier field set to 1 or 0x20 for the byte meaning peripheral not connected. If clear, the inquiry will fail with check condition of illegal request (sense=5/25/00).
<input type="checkbox"/>	Disable Pause when Not Ready	Reset	Normally, when controller is starting up, certain commands encounter a brief pause. If this field is checked, the pause is disabled, otherwise, it is enabled.
<input type="checkbox"/>	Disable BUSY status on fallback	Reset	During fallback, the survivor controller normally returns BUSY to new commands received from the host during the cache flush operation. If checked, requests are ignored. If unchecked, BUSY status is returned. This feature is intended to help hosts that are confused by a BUSY.
<input checked="" type="checkbox"/>	Enable DEBUG Port	Reset	If checked, then debug output will go to the serial port. This is for Engineering and Diagnostic purposes only, and will result in a performance loss if enabled. If not checked, then the port will be in SLP (no jumper) or VT100 mode (jumper - Standard Mode).
<input type="checkbox"/>	Enable Vendor Unique TUR	Reset	If checked, a Test Unit Ready command sent from the host to an off-line LUN will return a hard error status (4/00/00). If clear, then it will return a not ready status (2/04/03).
19200	Serial Port BAUD Rate	Reset	The baud rate of the serial port when in VT100 or Debug modes.
<input checked="" type="checkbox"/>	Force Simplex Mode	Reset	This is provided to allow duplex firmware to serve in a simplex environment, and makes duplex firmware skip some of the active-active operations. If you check this field in a duplex environment, then one of the controllers will go into reset.
<input type="checkbox"/>	Conservative Cache Mode	On-the-fly	This provides an extra degree of data safety when operating in failed over condition. This turns off write cache while the failed over condition persists. Checking this field enables conservative cache mode.
<input type="checkbox"/>	Simplex Disable Reset	Reset	For simplex-only. If checked, prevents a controller from asserting the reset signal to the partner controller.
2048 Bytes	Fibre Channel Frame Control	Reset	Provided to allow adjustment of the FC chip's frame size. Unless you are almost always doing very small block I/O, then you would want this set to 2048 .
512 Bytes	PCI Latency Control	Reset	Provided to allow adjustment of the FC chip's use of the internal bus. It controls the amount of data each FC processor chip can burst across the primary bus before relinquishing bus ownership to the next device. This takes effect only when all internal FC bus ports are active and arbitrating. Ordinarily you should leave this set to the factory default value of 512 . If you are in a high throughput environment, then you <i>may</i> see a slight performance advantage if you change the value to 2048 .

Figure 3-55 View/Modify Controller Configuration (partial for 7.75 Firmware)

<input checked="" type="checkbox"/>	Enable Hard Loop IDs	Reset	This allows you to force the hard loop ID's for every controller and port. The fields below are used to assign them. Do not enable this feature unless you set appropriate values for your controllers. Changes will not go into effect until you issue a reset or recycle power to the array.
Ctrl0 Port0 <input type="text" value="0"/> Ctrl1 Port0 <input type="text" value="3"/>	Hard Loop IDs for each controller and port	Reset	Allows option of using the same loop IDs all the time. Some fibre HBAs require non-default settings here. Contact your supplier for details. The valid range is 0-125 (0 - 0x7d). Enter a decimal number. Do not adjust these numbers unless instructed to do so by your vendor's technical support team.
<input checked="" type="checkbox"/>	Debug Dump	On-the-fly	If the controller fails, and this feature is enabled, then it will dump detailed diagnostic information into a reserved area of disk which can be sent to engineering for failure analysis. <i>This should always be enabled.</i>
<input type="text" value="DISABLED"/>	ROF Rearm Interval	On-the-fly	This reboot on failure rearm interval specifies how long the controller must stay up before the auto-reboot feature will be enabled.
<input type="text" value="0"/>	ROF Reboot Count	On-the-fly	# of times the controller will be rebooted when a controller failure occurs within the rearm interval above. If this count is exceeded in the rearm interval, no more reboots on failure will occur. To disable the ROF, set this field to 0 and the rearm interval above to something other than DISABLED .
<input checked="" type="checkbox"/>	Enable Background Init	On-the-fly	If enabled, you will then be allowed to perform background LUN initializations. This means that once you initiate a background init, then you will immediately be able to use the LUN for whatever I/O you wish to perform. The RAID engine will concurrently process application I/Os while initializing unused space. <i>Please do not change this value unless instructed to do so by your supplier.</i>
<input checked="" type="checkbox"/>	Smart Large Host Transfers	On-the-fly	If checked, large transfers will coalesce into fewer I/Os, which means fewer disconnects on large transfers. This only takes effect for transfers larger than the stripe size. If enabled, you will have a slightly higher throughput, at a possible cost of some IOPs.
<input type="text" value="Automatic"/>	Spin-Up Settings	On-the-fly	<input type="text" value="2"/> # of Devices to spin up at a time. Ignored for "On Power" mode. <input type="text" value="6"/> # of seconds to delay. For "Automatic", time between disk spin-up cycles. For "On Power", time before issuing start-up commands. For "On Command", time between disk spin-up cycles. <input type="text" value="0"/> # of seconds between subsequent spin-ups. Only applicable to "On Power"

Notes:

1. Settings characterized as **Reset** require that the controllers be reset before the new settings are invoked.
2. Settings characterized as **On-the-fly** may be made at any time, and will become effective immediately.
3. Settings characterized as **New Config** must be made during initial configuration, before any LUNs are created. **They will result in data loss of existing RAID groups.**

Apply Cancel Reset Screen

Figure 3-56 View/Modify Controller Configuration (partial for 7.75 Firmware)

The following selections appear on the 6.14 firmware dialog box that do not appear on the 7.01 and later firmware dialog box:

- **Stripe Size.** In controllers with 7.01 and later firmware, the stripe size is selected when the LUN is being defined (see Figure 3-13 on page 47). When controllers with 6.14 firmware are used, the stripe size must be defined using the **View/Modify Controller** dialog box before the LUN is defined.
- **Disk Queue Limit** (not viewable or changeable with controllers that have 7.01 and later firmware).

The Auto Failback selection appears for 7.01 and later firmware only. It is not viewable or changeable with controllers that have 6.14 firmware.

Make as many changes as you desire on the screen, then click the **Apply** button, which saves the new configuration on all controllers, as well as the COD area on your RAID subsystem's disk drives. The **Reset Screen** button changes the settings to the default values (the ones appearing when the screen was first selected).

Reset Controller(s)

When the Reset Controller(s) menu selection is made, the screen shown in Figure 3-57 appears.

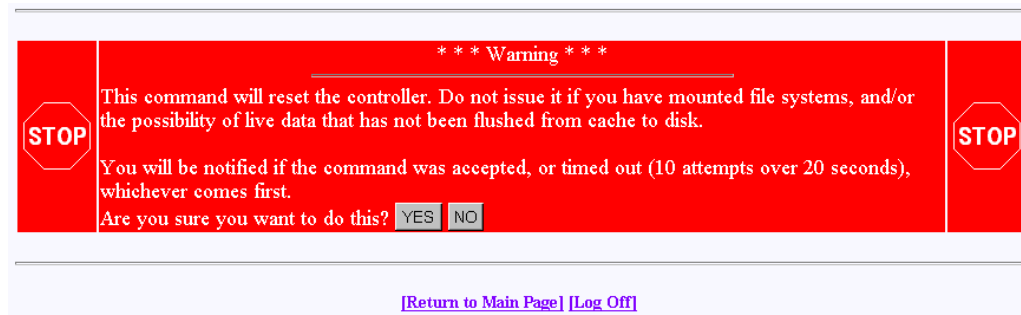


Figure 3-57 Reset Controllers Warning Screen

Click **YES** if you wish to reset your controller. If the system is a dual-controller configuration, clicking **YES** resets both controllers. Otherwise, click **NO** or the **[Return to Main Page]** link.

If you are not running Linux, you will see the screens shown in Figure 3-58 and Figure 3-59. These screens indicate when the controllers are back on-line. Click the close button (X) or press the **Close This Window** button after the controllers have reset.




Figure 3-58 Controller Booting Window




Figure 3-59 Controller Responding Window

As a protective measure, if the controllers are busy servicing I/Os, the Reset command will not be accepted.

Gracefully Bring a Controller Off Line

When the **Gracefully Bring a Controller Off Line** menu selection is made, the screen shown in Figure 3-60 appears.

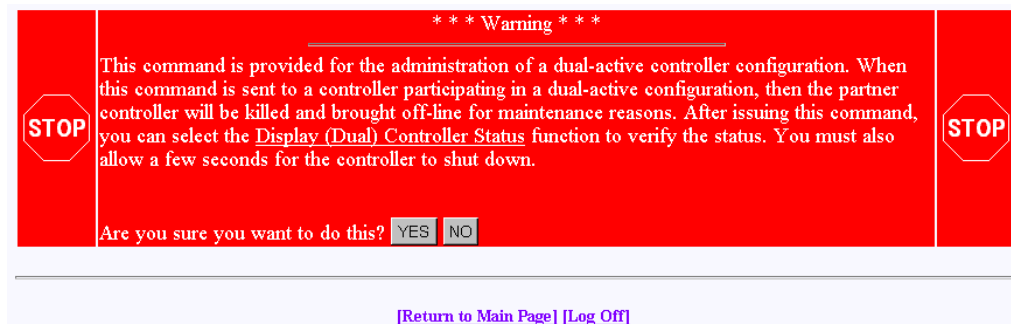


Figure 3-60 Warning Screen

Click **YES** button to initiate a controller failover. This is typically done for disaster recovery testing. You could also do the testing by physically removing a controller, but this lets you accomplish the same thing without touching the disk array.

Gracefully Bring a 2nd Controller On Line

When the **Gracefully Bring a 2nd Controller On Line** menu selection is made, the screen shown in Figure 3-61 appears.

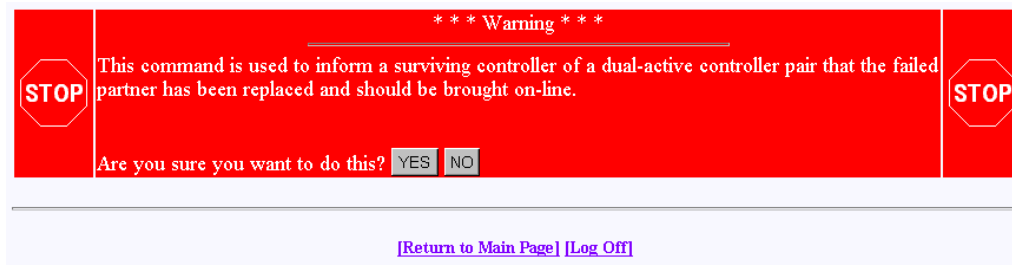


Figure 3-61 Warning Screen

When a failed controller is replaced, the system either automatically detects the replacement (if configured for automatic failback), or is informed of the replacement by issuing this command. The following steps outline the failback process executed by the surviving controller:

1. A replacement controller is detected.
2. The surviving controller releases its partner from reset.
3. Once the replacement controller completes initialization and is ready to resume I/O requests, the surviving controller quiesces both ports by responding with BUSY status to new I/O requests.
4. The surviving controller disables the failover port or secondary ID.
5. The surviving controller enables its primary ports.
6. The replacement controller enables its primary ports.
7. Both controllers disable conservative cache (if enabled) for write-back system drives and resume normal dual-active controller operation.




Click **YES** to force the failback.

Note: A replacement controller is held reset if a consistency check is in progress.

Perform Data Consistency Operations on a LUN

When the **Perform Data Consistency Check/Restore LUN** menu selection is made, the screen shown in Figure 3-62 appears.

Check and/or Restore Consistency of LUNs for RAID Subsystem at /hw/scsi/sc5a010

Logical Drive #0 (71618560 Blocks, 36,668,702,720 Bytes) Affinity Map: SAN Mapping	RAID 0+1 (Mirrored Stripe) 
Check Only? <input type="radio"/> Check and Restore? <input type="radio"/>	
Logical Drive #1 (107427840 Blocks, 55,003,054,080 Bytes) Affinity Map: SAN Mapping	RAID 3 (Parity Stripe) 
Check Only? <input type="radio"/> Check and Restore? <input type="radio"/>	
Logical Drive #2 (107427840 Blocks, 55,003,054,080 Bytes) Affinity Map: SAN Mapping	RAID 5 (Parity Stripe) 
Check Only? <input type="radio"/> Check and Restore? <input type="radio"/>	

Select the LUN you desire, then press the **Execute** button to immediately begin the process, or press the **Cancel** button to return to the main menu.

[\[Return to Main Page\]](#) [\[Log Off\]](#)

Figure 3-62 Data Consistency Check/Restore

Note: A check consistency will pause any background initialization process. Background initialization will not restart until a write is sent to the LUN against which background initialization was paused.

Use this screen to initiate a check and/or repair on a logical drive. To execute this process, the logical drive must be configured for high-availability RAID.

There are several reasons why the check or restore could be denied, and TPM reports the reasons if the request is rejected. The most common reasons are that there is no on-line spare disk to be used to repair the LUN, or more than one rebuild at a time is being attempted.

If you click **EXECUTE** to begin execute any of the check or restore actions, the dialog box of Figure 3-63 appears, indicating that the process has started.

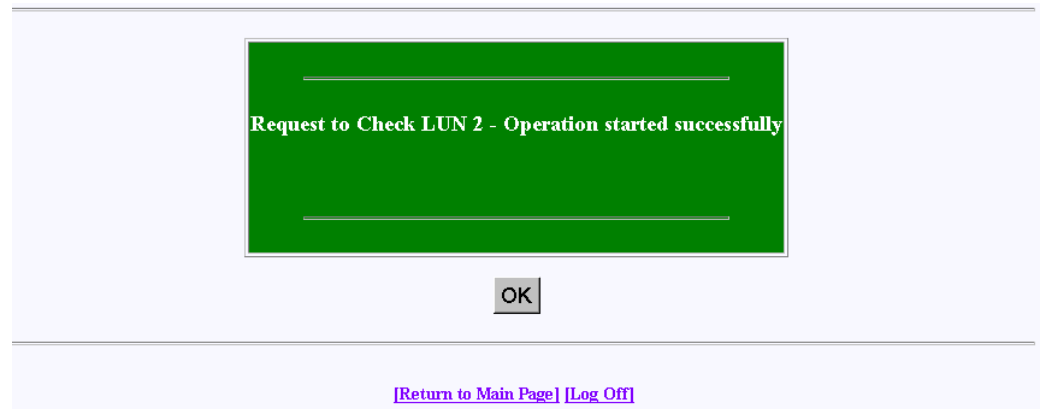


Figure 3-63 Operation Started Dialog Box

Click **OK** to proceed.

The window shown in Figure 3-64 (for 6.14 firmware), Figure 3-65 (for 7.01 and 7.03 firmware), and Figure 3-66 (for 7.75 firmware) shows the progress a few minutes after initiating a Check and Restore for Logical Drive #0, and 15 minutes after starting a LUN Initialization (format).


Background Job Status RAID Controller at /dev/rdisk/c1t0d0s0

Rebuild	In progress 0.7 % complete
Consistency Check	In progress 0.7 % complete.
Initialization Status	LUN 0 No initialization is in progress.
	LUN 1 In progress 78.8 % complete.
	LUN 2 No initialization is in progress.
	LUN 3 No initialization is in progress.
	LUN 4 No initialization is in progress.
	LUN 5 No initialization is in progress.
	LUN 6 No initialization is in progress.
	LUN 7 No initialization is in progress.


This screen will automatically refresh in approximately 10 seconds

Figure 3-64 Background Initialization Status Screen (6.14 Firmware)


Background Job Status RAID Controller at /hw/scsi/sc13d010

Logical Drive #0 (70120 MB, 73,526,149,120 Bytes) RAID 0 (Stripe) 

Action	Notes
Initialization	Initialization terminated by operator.

Logical Drive #1 (35060 MB, 36,763,074,560 Bytes) RAID 0+1 (Mirrored Stripe) 

Action	Notes
Rebuild	Not in progress
Consistency Check	Not in progress
Initialization	Initialization terminated by operator.
Online RAID Expansion	Not in progress

Logical Drive #2 (35060 MB, 36,763,074,560 Bytes) RAID 5 (Right Asymmetric) 

Action	Notes
Rebuild	Not in progress
Consistency Check <input type="button" value="Cancel All"/>	In progress 0.2 % complete.
Initialization	Initialization terminated by operator.
Online RAID Expansion	Not in progress

This screen will automatically refresh in approximately 10 seconds

[\[Return to Main Page\]](#) [\[Log Off\]](#)

Figure 3-65 Background Initialization Status Screen (7.01 and 7.03 Firmware)

**Background Job Status RAID Controller at
/hw/scsi/20000080e5115f52/Lun0/c3p1**





Logical Drive #0 (34332 MB, 35,999,711,232 Bytes)		RAID 3 (Right Asymmetric) 
Action	Notes	
Rebuild	Not in progress	
Consistency Check	Not in progress	
Initialization	Not in progress	
Background initialization <input type="button" value="Cancel All"/>	In progress 9.8 % complete (18 mins left).	
Online RAID Expansion	Not in progress	
Logical Drive #1 (34332 MB, 35,999,711,232 Bytes)		RAID 5 (Right Asymmetric) 
Action	Notes	
Rebuild	Not in progress	
Consistency Check	Not in progress	
Initialization	Not in progress	
Background initialization	Has not been initiated.	
Online RAID Expansion	Not in progress	
Logical Drive #2 (34332 MB, 35,999,711,232 Bytes)		RAID 0+1 (Mirrored Stripe) 
Action	Notes	
Rebuild	Not in progress	
Consistency Check	Not in progress	
Initialization	Not in progress	
Background initialization	Has not been initiated.	
Online RAID Expansion	Not in progress	
Logical Drive #3 (17166 MB, 17,999,855,616 Bytes)		RAID 1 (Mirror) 
Action	Notes	
Rebuild	Not in progress	
Consistency Check	Not in progress	
Initialization	Not in progress	
Background initialization	Has not been initiated.	
Online RAID Expansion	Not in progress	

Figure 3-66 Background Initialization Status Screen (7.75 Firmware)

Enable/Disable Write Cache for LUN(s)

When the **Enable/Disable Write Cache for LUN(s)** menu selection is made, the dialog box shown in Figure 3-67 appears.

(Enable/Disable Write Cache) Logical Disk Information for RAID Subsystem at /hw/scsi/sc13d010

Logical Drive #0 (70120 MB, 73,526,149,120 Bytes)	RAID 0 (Stripe)	
Affinity Map: SAN Mapping	Enable Write Cache <input checked="" type="checkbox"/>	
Logical Drive #1 (35060 MB, 36,763,074,560 Bytes)	RAID 0+1 (Mirrored Stripe)	
Affinity Map: SAN Mapping	Enable Write Cache <input checked="" type="checkbox"/>	
Logical Drive #2 (35060 MB, 36,763,074,560 Bytes)	RAID 5 (Right Asymmetric)	
Affinity Map: SAN Mapping	Enable Write Cache <input checked="" type="checkbox"/>	

Select as many LUNs as you desire, then press the **Apply** button to make the changes, or press the **Cancel** button.

[\[Return to Main Page\]](#) [\[Log Off\]](#)

Figure 3-67 Enable Write Cache For LUN(s)

Select the LUNs where you want the cache enabled or disabled and click **Apply** at the bottom of the screen.

Modify/Purge Non-Volatile WWN Tables

When the **Modify/Purge Non-Volatile WWN Tables** menu selection is made, the dialog box shown in Figure 3-68 appears.

World Wide Name Table for RAID controller at device:
/hw/scsi/2000080e5115f52/lun0/c3p1

- This is a list of the WWNs for each controller on your fibre channel loop, and what host ports they are connected to.
- If a RAID controller does not appear below, then the Fibre Channel controller can't see it.
- You should map each host adapter to both controller ports, if you have a dual active configuration, and wish to have failover support.
- World wide names in RED indicate that they are cached entries, and no longer connected to the fibre channel.
- If you do not wish to define an entry, then either set it to blanks, or leave it as 00-00-00-00-00-00-00.

World Wide Name	Controller 0	Controller 1
10-00-00-60-69-20-15-72	✓	✗
20-00-00-E0-8B-00-F3-C4	✓	✗
00-00-00-00-00-00-00-00		
00-00-00-00-00-00-00-00		
00-00-00-00-00-00-00-00		
00-00-00-00-00-00-00-00		
00-00-00-00-00-00-00-00		
00-00-00-00-00-00-00-00		
00-00-00-00-00-00-00-00		
00-00-00-00-00-00-00-00		
00-00-00-00-00-00-00-00		

SAVE

Cancel

Purge Old Entries

Reset Screen

[\[Return to Main Page\]](#) [\[Log Off\]](#)

Figure 3-68 World Wide Name Table for RAID Controller

The World Wide Name (WWN) table feature has been modified with firmware version 7.75 to include a delete WWN from WWN table option. The number of supported host WWNs is 64.

Note: All SAN mapping assignments created using firmware version 7.75 are lost if the firmware is downgraded to previous firmware version.

A controller currently maintains a host WWN table until the configuration is cleared. As WWNs are removed and the WWN table updated, any WWNs following those deleted are moved up to fill the vacancies in the WWN table. The SAN map uses the indices of the WWN table entries to specify hosts that have access to a specific system drive. When the WWN table entries are deleted, the indices change for any entries following those WWNs that were deleted.

Reporting Functions

The Reporting Functions menu is located on the main screen and is shown in Figure 3-69 for 6.14 controller firmware and in Figure 3-70 for 7.01 and later controller firmware.

Reporting Functions:

- [Set default screen refresh rate](#) - This lets you define the number of seconds between each screen refresh for status screens which automatically repaint.
- [Display \(Dual\) Controller Status](#) - This returns status information on dual controller status, and host addressing information on the connected controller.
- [Topology query](#) - Displays all host adapters on the SAN attached to the subsystem, and what controller/ports they are attached to.
- [Display statistical data by physical device](#) - This shows log page information for an individual disk drive.
- [Display statistical data by logical \(RAID\) disk](#) - Shows cumulative reads, writes, and cache hits.
- [Display SCSI/Fibre device information](#) - This issues a standard SCSI Inquiry, and reports all fields which describe the device.
- [Display FULL subsystem configuration information](#) - This is a complete hexadecimal dump of the controller's configuration data structures, and contains information which may be of interest to your supplier in the event of a problem.
- [Display physical subsystem information](#) - displays drive status, statistics, errors and physical locations for all disks in a subsystem, including expansion units.
- [Display logical subsystem information](#) - displays RAID groups status, statistics, errors and logical configuration for all RAID groups in a subsystem, including expansion units.
- [Display environmental subsystem information](#) - displays power, fans, temperature, battery backup, and other data relating to the chassis, including expansion units.

Figure 3-69 Reporting Functions Menu (for 6.14 Firmware)

Reporting Functions:

- [Set default screen refresh rate](#) - This lets you define the number of seconds between each screen refresh for status screens which automatically repaint.
- [Display \(Dual\) Controller Status](#) - This returns status information on dual controller status, and host addressing information on the connected controller.
- [Topology query](#) - Displays all host adapters on the SAN attached to the subsystem, and what controller/ports they are attached to.
- [Display statistical data by physical device](#) - This shows log page information for an individual disk drive.
- [Display SCSI/Fibre device information](#) - This issues a standard SCSI Inquiry, and reports all fields which describe the device.
- [Display FULL subsystem configuration information](#) - This is a complete hexadecimal dump of the controller's configuration data structures, and contains information which may be of interest to your supplier in the event of a problem.
- [Display physical subsystem information](#) - displays drive status, statistics, errors and physical locations for all disks in a subsystem, including expansion units.
- [Display logical subsystem information](#) - displays RAID groups status, statistics, errors and logical configuration for all RAID groups in a subsystem, including expansion units.
- [Display environmental subsystem information](#) - displays power, fans, temperature, battery backup, and other data relating to the chassis, including expansion units.

Figure 3-70 Reporting Functions Menu (for 7.01 and Later Firmware)

The Reporting Functions menu has the following selections, with the associated explanations on the indicated pages:

- “Set Default Screen Refresh Rate” on page 107
- “Display (Dual) Controller Status” on page 108
- “Topology Query” on page 109
- “Display Statistical Data by Physical Device” on page 110

- “Display Statistical Data by Logical (RAID) Disk (6.14 Firmware Only)” on page 110
- “Display SCSI/Fibre Device Information” on page 111
- “Display FULL Subsystem Configuration Information” on page 112
- “Display Physical Subsystem Information” on page 114
- “Display Logical Subsystem Information” on page 118
- “Display Environmental Subsystem Information” on page 120

Set Default Screen Refresh Rate

When the **Set Default Screen Refresh Rate** menu selection is made, the screen shown in Figure 3-71 appears.

Set Default Screen Refresh Rate

The field below allows you to specify the number of seconds between each screen refresh for functions that automatically update themselves. Once you make a change, it will be in effect until the configurator service routine running on your host is terminated. When you first start the program, the default rate is every 10 seconds. The valid range is 2 - 999999 seconds.

Current Refresh Rate (seconds)

SAVE **Cancel**

[\[Return to Main Page\]](#) [\[Log Off\]](#)

Figure 3-71 Default Screen Refresh Rate

Enter the desired screen refresh rate in seconds and click **SAVE**.

Display (Dual) Controller Status

When the **Display (Dual) Controller Status** menu selection is made, the screen shown in Figure 3-72 appears.

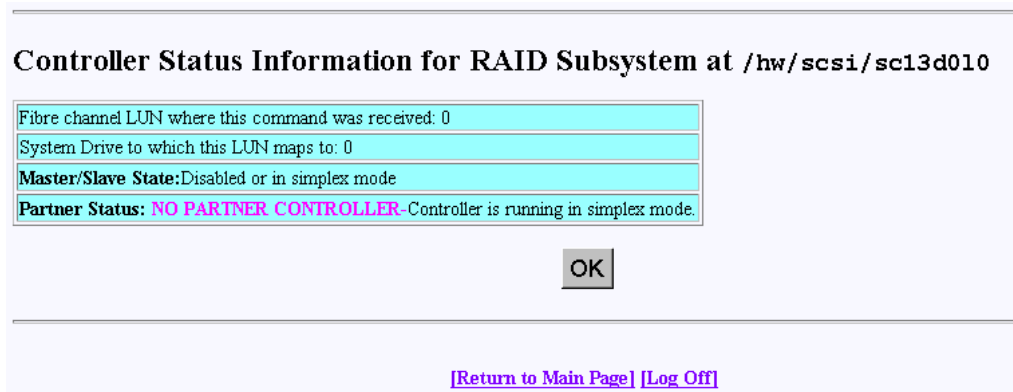


Figure 3-72 Controller Status Information

This Screen displays information about the controller(s) and whether or not they are working together. There are nearly 100 different error or warning messages that can be returned, and this could be quite useful in the event you have a controller failure.

You should also periodically check this screen during normal operations to make sure that all is well. In some cases you may have a controller failure which does NOT result in an audible or visual alarm.

Topology Query

When the **Topology Query** menu selection is made, the screen shown in Figure 3-73 appears.

World Wide Name Table for RAID controller at device: /hw/scsi/20000080e5110413/1un0/c5p1

- This is a list of the WWNs for each controller on your fibre channel loop, and what host ports they are connected to.
- If a RAID controller does not appear below, then the Fibre Channel controller can't see it.
- You should map each host adapter to both controller ports, if you have a dual active configuration, and wish to have failover support.
- World wide names in RED indicate that they are cached entries, and no longer connected to the fibre channel.

World Wide Name	Controller 0	Controller 1
20-00-00-E0-8B-02-E6-37		
20-00-08-00-69-04-85-D7		
10-00-00-60-69-20-12-86		
20-00-00-E0-8B-01-34-3B		
20-00-00-E0-8B-00-F3-C4		
20-00-08-00-69-04-87-D4		
20-00-08-00-69-04-87-E6		
20-00-08-00-69-00-00-07		
10-00-00-60-69-10-02-4E		
10-00-00-60-69-10-02-5D	<input checked="" type="checkbox"/>	
20-00-00-E0-8B-01-D9-BD		
20-00-00-E0-8B-00-3D-D8	<input checked="" type="checkbox"/>	
20-00-00-E0-8B-01-A2-11	<input checked="" type="checkbox"/>	
10-00-00-60-69-10-02-33		
10-00-00-60-69-10-1F-0E	<input checked="" type="checkbox"/>	
20-00-00-E0-8B-00-00-00	<input checked="" type="checkbox"/>	
20-00-00-E0-8B-01-37-39	<input checked="" type="checkbox"/>	

OK

[\[Return to Main Page\]](#) [\[Log Off\]](#)

Figure 3-73 Topology Query

This screen displays a list of Fibre Channel host adapters that are (or were) attached to the RAID controller. Use this screen to view limited topology information.

Note: This is a **read-only** display. No parameters can be changed.

Display Statistical Data by Physical Device

Note: This function is not supported by the TP9100 RAID system. It is only for JBOD environments.

Display Statistical Data by Logical (RAID) Disk (6.14 Firmware Only)

When the **Display Statistical Data by Logical RAID Disk** menu selection is made, the screen shown in Figure 3-74 appears.

Statistical Report by Logical Disk for Device at /hw/scsi/20000080e5110413/lun0/c5p1

Drive #	# of Read Commands	# of Blocks Read	# of Write Commands	# of Blocks Written	# of Read Cache Hits	MB Read /Sec	MB Written /Sec
Totals:	0	0	0	0	0	0.00	0.00

- These numbers are reset to zero every time the controller is powered up, and every time this screen is displayed.
- If no reads or writes are performed on a LUN, then a line for it will not be displayed.

This screen will automatically refresh in approximately 30 seconds

[\[Return to Main Page\]](#) [\[Log Off\]](#)

Figure 3-74 Statistical Data for Logical Devices (for 6.14 Firmware)

This screen displays cumulative reads and writes for all logical drives since the last polling period. If no I/Os occurred, nothing is displayed. In the screen above, drives number 0 through 3 were just initialized, and no other operations were done during that time.

All numbers reset every time the screen is brought up. Use this screen to view the actual number of I/O operations that get serviced by each LUN during the 10-second polling period.

Display SCSI/Fibre Device Information

When the **Display SCSI/Fibre Device Information** menu selection is made, the screen shown in Figure 3-75 appears.

SCSI Inquiry Dump for Device at /hw/scsi/sc13d010

<u>Device Characteristics</u>	<u>Value</u>
Vendor (Manufacturer)	MYLEX
Product Identifier	DACARMB 70120B0
Microcode Revision Level	7701
Unit Serial Number	
Device Type	Disk
ANSI SCSI Version	2

Device Capabilities:

- ISO Compliant** - Indicates if this device claims compliance to the International Organization for Standards (ISO) version of SCSI (ISO DIS 9316).
- ECMA Compliant** - Indicates if this device claims compliance to the European Computer Manufacturers Association (ECMA) version of SCSI (ECMA-111).

<input type="checkbox"/> 32-Bit Transfers	<input type="checkbox"/> 32-Bit Addressing	<input type="checkbox"/> Terminate Task Management
<input type="checkbox"/> 16-Bit Transfers	<input type="checkbox"/> 16-Bit Addressing	<input type="checkbox"/> Normal ACA (NormACA)
<input type="checkbox"/> Synchronous Data Transfers	<input type="checkbox"/> Relative Addressing Mode	<input type="checkbox"/> Linked Commands on LUN
<input type="checkbox"/> Transfer Disable Messages Supported	<input type="checkbox"/> Enclosure Services Available	<input type="checkbox"/> Dual-ported Device
<input type="checkbox"/> Handshake on Q Cable Supported	<input type="checkbox"/> Removable Media	<input checked="" type="checkbox"/> Tagged Command Queuing

```

0000h: 00 00 02 12 33 00 00 02 4D 59 4C 45 58 20 20 20 ...3...MYLEX
0010h: 44 41 43 41 52 4D 52 42 20 37 30 31 32 30 42 30 DACARMB 70120B0
0020h: 37 37 30 31 00 00 00 00 20 00 00 80 E5 11 4C 14 7701....L..
0030h: 00 00 00 00 00 00 00 00 - - - - - - - - .....
```

OK

[\[Return to Main Page\]](#) [\[Log Off\]](#)

Figure 3-75 Display SCSI/Fibre Device Information

The screen shows what the standard SCSI inquiry returns for a device or LUN. Use it for diagnostic reasons, or to help analyze any SCSI or Fibre Channel device attached to your computer.

Display FULL Subsystem Configuration Information

When the **Display FULL Subsystem Configuration Information** menu selection is made, the Controller Information Values (see Figure 3-76), Controller Tunable Parameters (see Figure 3-77), and GroupConfig Structure (see Figure 3-78) screens appear.

Hexadecimal Diagnostic Dump for Subsystem on device /hw/scsi/20000080e5115f52/lun0/c3p1:

Controller Information values: [HDACIOCTL_GETCONTROLLERINFO]																
0000h:	00	01	6B	00	64	00	01	00	01	00	00	00	00	00	00	..k.d.....
0010h:	52	41	49	44	42	52	49	43	4B	00	00	00	00	00	00	RAIDBRICK.....
0020h:	4D	59	4C	45	58	20	44	41	43	46	46	78	20	20	20	MYLEX DACFFx
0030h:	00	00	00	00	00	00	00	00	00	00	00	00	00	00	00
0040h:	07	46	00	00	00	00	00	00	00	00	00	00	00	00	00	.F.....
0050h:	00	00	00	00	00	00	00	00	7E	10	F8	7F	20	00	00~.....
0060h:	00	00	00	00	00	00	00	00	00	00	00	00	00	00	00
0070h:	00	00	00	00	00	00	00	00	00	00	00	00	00	00	00
0080h:	00	00	00	00	4D	59	4C	45	58	20	20	20	20	20	20	...MYLEX
0090h:	20	20	20	20	01	00	00	00	00	00	00	00	20	00	80
00A0h:	02	00	00	00	00	00	0D	0C	00	00	00	00	00	F4	00
00B0h:	02	00	02	00	7E	7E	00	00	00	00	00	00	00	00	00~.....
00C0h:	00	00	00	00	00	00	00	00	00	00	00	00	00	00	00
00D0h:	20	00	80	00	00	00	00	00	00	00	00	00	00	00	00
00E0h:	00	00	00	00	00	00	00	00	00	00	00	00	00	00	00

Figure 3-76 Controller Information Values (partial screen)

Controller Tunable Parameters: [HDACIOCTL_GETCONTROLLERPARAMETER]																	
0000h:	41	80	01	00	70	32	00	00	00	04	06	00	50	14	04	01	A...p2.....P...
0010h:	06	10	00	00	01	03	00	C8	00	00	03	00	4D	59	4C	45MYLE
0020h:	58	20	44	41	43	46	46	78	20	20	20	20	00	4D	59	4C	X DACFFx .MYL
0030h:	45	58	20	44	41	43	46	46	78	20	20	20	20	00	00	00	EX DACFFx ...
0040h:	00	00	12	01	00	00	00	00	00	00	00	00	00	00	00	00
0050h:	00	00	00	00	00	00	00	00	00	00	00	00	00	00	00	00
0060h:	00	00	00	00	00	00	00	00	00	00	00	00	00	00	00	00
0070h:	00	00	00	00	00	00	00	00	00	00	00	00	00	00	00	00

Figure 3-77 Controller Tunable Parameters


```

GroupConfig Structure: [HDACIOCTL_GETGROUPCONFINFO]
0000h: 0F 01 00 C0 00 70 18 02 00 00 7D 00 00 00 00 00 .....p....}.....
0010h: 0F 01 01 C0 00 70 18 02 00 01 7C 00 00 00 00 00 .....p....|.....
0020h: 0F 01 02 C0 00 70 18 02 00 00 7B 00 00 00 00 00 .....p....{.....
0030h: 0F 01 03 C0 00 70 18 02 00 01 7A 00 00 00 00 00 .....p....z.....
0040h: 0F 01 04 C0 00 70 18 02 00 00 79 00 00 00 00 00 .....p....y.....
0050h: 0F 01 05 C0 00 70 18 02 00 01 78 00 00 00 00 00 .....p....x.....
0060h: 0F 01 06 C0 00 70 18 02 00 00 77 00 00 00 00 00 .....p....w.....
0070h: 0F 01 07 C0 00 70 18 02 00 01 76 00 00 00 00 00 .....p....v.....
0080h: 0F 01 08 C0 00 70 18 02 00 00 75 00 00 00 00 00 .....p....u.....
0090h: 0F 01 09 C0 00 70 18 02 00 01 74 00 00 00 00 00 .....p....t.....
00A0h: 0F 01 0A C0 00 70 18 02 00 00 73 00 00 00 00 00 .....p....s.....
00B0h: 0F 01 0B C0 00 70 18 02 00 01 72 00 00 00 00 00 .....p....r.....
00C0h: 06 01 00 00 00 E0 30 04 04 08 08 10 00 70 18 02 .....0.....p..
00D0h: 00 01 00 C0 00 00 00 00 00 01 01 C0 00 00 00 00 .....
00E0h: 00 01 02 C0 00 00 00 00 00 01 03 C0 00 00 00 00 .....
00F0h: 03 01 01 00 00 50 49 06 04 08 08 10 00 70 18 02 .....PI.....p..
0100h: 00 01 04 C0 00 00 00 00 00 01 05 C0 00 00 00 00 .....
0110h: 00 01 06 C0 00 00 00 00 00 01 07 C0 00 00 00 00 .....

```

OK

Figure 3-78 GroupConfig Structure

These screens contain hexadecimal dump information from the controller configuration file. This would be of interest to the technical support group in the event of a problem, or would prove helpful to verify that several subsystem configurations have properly been cloned.

Display Physical Subsystem Information

When the **Display Physical Subsystem Information** menu selection is made, the screen shown in Figure 3-79 and Figure 3-80 appear (for controllers with 6.14 firmware). Figure 3-81 and Figure 3-82 appear for controllers with 7.01 and later firmware.

Physical Disk Information for RAID Subsystem at
/hw/scsi/20000080e5110413/lun0/c5p1
(See legend below)

Enclosure #0 (Rackmount view, rotate 90° clockwise for tower)			
S/GI ST173404FC [2706] S/N: 3CE02LFD 100Mhz, 2 Ports, FC-AL Physical: 70007 Useable: 69991 LoopID=125 (7Dh) Chan=0 ID=0 [ONLINE]	S/GI ST173404FC [2705] S/N: 3CE02C6C 100Mhz, 2 Ports, FC-AL Physical: 70007 Useable: 69991 LoopID=124 (7Ch) Chan=1 ID=0 [ONLINE]	S/GI ST173404FC [2705] S/N: 3CE02LS6 100Mhz, 2 Ports, FC-AL Physical: 70007 Useable: 69991 LoopID=123 (7Bh) Chan=0 ID=1 [ONLINE]	S/GI ST173404FC [2705] S/N: 3CE02J9C 100Mhz, 2 Ports, FC-AL Physical: 70007 Useable: 69991 LoopID=122 (7Ah) Chan=1 ID=1 [ONLINE]
Errors: 0 0 0 0	Errors: 0 0 0 0	Errors: 0 0 0 0	Errors: 0 0 0 0
Reads: 0	Reads: 0	Reads: 0	Reads: 0
Blocks Read: 0	Blocks Read: 0	Blocks Read: 0	Blocks Read: 0
Writes: 0	Writes: 0	Writes: 0	Writes: 0
Blocks Written: 0	Blocks Written: 0	Blocks Written: 0	Blocks Written: 0
S/GI ST173404FC [2705] S/N: 3CE02JP4 100Mhz, 2 Ports, FC-AL Physical: 70007 Useable: 69991 LoopID=121 (79h) Chan=0 ID=2 [ONLINE]	S/GI ST173404FC [2705] S/N: 3CE02LRQ 100Mhz, 2 Ports, FC-AL Physical: 70007 Useable: 69991 LoopID=120 (78h) Chan=1 ID=2 [ONLINE]	S/GI ST173404FC [2705] S/N: 3CE021HE 100Mhz, 2 Ports, FC-AL Physical: 70007 Useable: 69991 LoopID=119 (77h) Chan=0 ID=3 [ONLINE]	S/GI ST173404FC [2705] S/N: 3CE02F1C 100Mhz, 2 Ports, FC-AL Physical: 70007 Useable: 69991 LoopID=118 (76h) Chan=1 ID=3 [ONLINE]
Errors: 0 0 0 0	Errors: 0 0 0 0	Errors: 0 0 0 0	Errors: 0 0 0 0
Reads: 0	Reads: 0	Reads: 0	Reads: 0
Blocks Read: 0	Blocks Read: 0	Blocks Read: 0	Blocks Read: 0
Writes: 0	Writes: 0	Writes: 0	Writes: 0
Blocks Written: 0	Blocks Written: 0	Blocks Written: 0	Blocks Written: 0
S/GI ST173404FC [2705] S/N: 3CE02LNR 100Mhz, 2 Ports, FC-AL Physical: 70007 Useable: 69991 LoopID=117 (75h) Chan=0	S/GI ST173404FC [2705] S/N: 3CE02LRS 100Mhz, 2 Ports, FC-AL Physical: 70007 Useable: 69991 LoopID=116 (74h) Chan=1	S/GI ST173404FC [2705] S/N: 3CE02LNE 100Mhz, 2 Ports, FC-AL Physical: 70007 Useable: 69991 LoopID=115 (73h) Chan=0	S/GI ST173404FC [2705] S/N: 3CE02LDR 100Mhz, 2 Ports, FC-AL Physical: 70007 Useable: 69991 LoopID=114 (72h) Chan=1

Figure 3-79 Physical Subsystem Information (partial for 6.14 Firmware)

Errors:	0	0	0	0	Errors:	0	0	0	0	Errors:	0	0	0	0	Errors:	0	0	0	0
	0	0	0	OK		0	0	0	OK		0	0	0	OK		0	0	0	OK
Reads:	0				Reads:	0				Reads:	0				Reads:	0			
Blocks Read:	0				Blocks Read:	0				Blocks Read:	0				Blocks Read:	0			
Writes:	0				Writes:	0				Writes:	0				Writes:	0			
Blocks Written:	0				Blocks Written:	0				Blocks Written:	0				Blocks Written:	0			
SGI ST173404FC [2705] S/N: 3CE02LNR 100Mhz, 2 Ports, FC-AL Physical: 70007 Useable: 69991 LoopID=117 (75h) Chan=0 ID=4 [ONLINE]					SGI ST173404FC [2705] S/N: 3CE02LKS 100Mhz, 2 Ports, FC-AL Physical: 70007 Useable: 69991 LoopID=116 (74h) Chan=1 ID=4 [ONLINE]					SGI ST173404FC [2705] S/N: 3CE02LNE 100Mhz, 2 Ports, FC-AL Physical: 70007 Useable: 69991 LoopID=115 (73h) Chan=0 ID=5 [ONLINE]					SGI ST173404FC [2705] S/N: 3CE02LDR 100Mhz, 2 Ports, FC-AL Physical: 70007 Useable: 69991 LoopID=114 (72h) Chan=1 ID=5 [ONLINE]				
Errors:	0	0	0	0	Errors:	0	0	0	0	Errors:	0	0	0	0	Errors:	0	0	0	0
	0	0	0	OK		0	0	0	OK		0	0	0	OK		0	0	0	OK
Reads:	0				Reads:	0				Reads:	0				Reads:	0			
Blocks Read:	0				Blocks Read:	0				Blocks Read:	0				Blocks Read:	0			
Writes:	0				Writes:	0				Writes:	0				Writes:	0			
Blocks Written:	0				Blocks Written:	0				Blocks Written:	0				Blocks Written:	0			

Legend:

VendorID	ProductID	[Firmware Revision]		
Clock Speed	# of Ports	InterfaceType		
Device Size in Megabytes:				
Usable Size in Megabytes:				
LoopID	decimal (HEX)	Channel	TargetID	
Cumulative Errors	Parity	Soft	Hard	Miscellaneous
	Command Timeouts	Retries	Aborts	Predictive Faults
Read Operations				
Blocks Read				
Write Operations				
Blocks Written				

OK

[\[Return to Main Page\]](#) [\[Log Off\]](#)

Figure 3-80 Physical Subsystem Information (partial for 6.14 Firmware)

Physical Disk Information for RAID Subsystem at /hw/scsi/sc13d010

(See legend below)

Enclosure #0 (Rackmount view, rotate 90° clockwise if in tower)			
SGI ST318304FC [2705] S/N: 3E0098Q WWN: 2000002037653467 100Mhz, 2 Ports, FC-AL Drive Speed: 10016 RPM Physical: 17560 Useable: 17530 LoopID=125 (7Dh) Chan=0 ID=125 [OPTIMAL]	SGI ST336704FC [2705] S/N: 3CD0156L WWN: 20000020372A17B9 100Mhz, 2 Ports, FC-AL Drive Speed: 10016 RPM Physical: 35003 Useable: 34696 LoopID=124 (7Ch) Chan=1 ID=124 [OPTIMAL]	SGI ST318304FC [2705] S/N: 3E00W4T WWN: 2000002037659735 100Mhz, 2 Ports, FC-AL Drive Speed: 10016 RPM Physical: 17560 Useable: 17530 LoopID=123 (7Bh) Chan=0 ID=123 [OPTIMAL]	SGI ST318304FC [2705] S/N: 3E00W1C WWN: 20000020376598A1 100Mhz, 2 Ports, FC-AL Drive Speed: 10016 RPM Physical: 17560 Useable: 17530 LoopID=122 (7Ah) Chan=1 ID=122 [OPTIMAL]
Errors: 0 0 0 0 0 0 0 OK	Errors: 0 0 0 0 0 0 0 OK	Errors: 0 0 0 0 0 0 0 OK	Errors: 0 0 0 0 0 0 0 OK
Reads: 0	Reads: 0	Reads: 0	Reads: 0
Active Commands: 0	Active Commands: 0	Active Commands: 0	Active Commands: 0
Writes: 0	Writes: 0	Writes: 0	Writes: 0
Queued Commands: 0	Queued Commands: 0	Queued Commands: 0	Queued Commands: 0
SGI ST318304FC [2705] S/N: 3E00HNG WWN: 2000002037658ECD 100Mhz, 2 Ports, FC-AL Drive Speed: 10016 RPM Physical: 17560 Useable: 17530 LoopID=121 (79h) Chan=0 ID=121 [OPTIMAL]	SGI ST318304FC [2705] S/N: 3E00KDG WWN: 200000203765985D 100Mhz, 2 Ports, FC-AL Drive Speed: 10016 RPM Physical: 17560 Useable: 17530 LoopID=120 (78h) Chan=1 ID=120 [OPTIMAL]	SGI ST318304FC [2705] S/N: 3E00W88 WWN: 200000203765972F 100Mhz, 2 Ports, FC-AL Drive Speed: 10016 RPM Physical: 17560 Useable: 17530 LoopID=119 (77h) Chan=0 ID=119 [OPTIMAL]	SGI ST318304FC [2705] S/N: 3E00XDT WWN: 2000002037658EB6 100Mhz, 2 Ports, FC-AL Drive Speed: 10016 RPM Physical: 17560 Useable: 17530 LoopID=118 (76h) Chan=1 ID=118 [OPTIMAL]
Errors: 0 0 0 0 0 0 0 OK	Errors: 0 0 0 0 0 0 0 OK	Errors: 0 0 0 0 0 0 0 OK	Errors: 0 0 0 0 0 0 0 OK
Reads: 0	Reads: 0	Reads: 0	Reads: 0
Active Commands: 0	Active Commands: 0	Active Commands: 0	Active Commands: 0
Writes: 0	Writes: 0	Writes: 0	Writes: 0
Queued Commands: 0	Queued Commands: 0	Queued Commands: 0	Queued Commands: 0
SGI ST318304FC [2705] S/N: 3E00VAD	SGI ST318304FC [2705] S/N: 3E008JW	SGI ST318304FC [2705] S/N: 3E00769	SGI ST318304FC [2705] S/N: 3E00W3B

Figure 3-81 Physical Subsystem Information (partial for 7.01 and later Firmware)

SGI ST318304FC [2705] S/N: 3ELOOVAD WWN: 20000020376595D1 100Mhz, 2 Ports, FC-AL Drive Speed: 10016 RPM Physical: 17560 Useable: 17530 LoopID=117 (75h) Chan=0 ID=117 [OPTIMAL]	SGI ST318304FC [2705] S/N: 3ELO08JW WWN: 20000020376597A1 100Mhz, 2 Ports, FC-AL Drive Speed: 10016 RPM Physical: 17560 Useable: 17530 LoopID=116 (74h) Chan=1 ID=116 [OPTIMAL]	SGI ST318304FC [2705] S/N: 3ELO0769 WWN: 20000020376529C8 100Mhz, 2 Ports, FC-AL Drive Speed: 10016 RPM Physical: 17560 Useable: 17530 LoopID=115 (73h) Chan=0 ID=115 [OPTIMAL]	SGI ST318304FC [2705] S/N: 3ELO0W3B WWN: 2000002037659890 100Mhz, 2 Ports, FC-AL Drive Speed: 10016 RPM Physical: 17560 Useable: 17530 LoopID=114 (72h) Chan=1 ID=114 [OFF-LINE]
Errors: 0 0 0 0 0 0 0 OK	Errors: 0 0 0 0 0 0 0 OK	Errors: 0 0 0 0 0 0 0 OK	Errors: 0 0 0 0 0 0 0 OK
Reads: 0	Reads: 0	Reads: 0	Reads: 0
Active Commands: 0	Active Commands: 0	Active Commands: 0	Active Commands: 0
Writes: 0	Writes: 0	Writes: 0	Writes: 0
Queued Commands: 0	Queued Commands: 0	Queued Commands: 0	Queued Commands: 0

Legend:

VendorID ProductID [Firmware Revision]				
Clock Speed, # of Ports, InterfaceType				
Device Size in Megabytes:				
Usable Size in Megabytes:				
LoopID decimal (HEX) Channel TargetID				
Cumulative Errors	Parity	Soft	Hard	Miscellaneous
	Command Timeouts	Retries	Aborts	Predictive Faults
Read Operations				
# of Active Commands				
Write Operations				
# of Queued Commands				

OK

[\[Return to Main Page\]](#) [\[Log Off\]](#)

Figure 3-82 Physical Subsystem Information (partial for 7.01 and later Firmware)


The 7.01 and later controller firmware returns WWN and Drive Speed information, which are not returned by controllers with 6.14 firmware.


These screens shows statistical data, drive status, and errors for all drives in all subsystems. It also correctly displays the enclosure number, row, and column of each disk drive. All numbers are cumulative, starting from zero when the subsystem is powered on. Under normal operation you should rarely see any errors.

Display Logical Subsystem Information

When the **Display Logical Subsystem Information** menu selection is made, the windows shown in Figure 3-83 (for controllers with 6.14 firmware) and Figure 3-84 (for controllers with 7.01 and later firmware) appear.

Logical Disk Information for RAID Subsystem at /hw/scsi/20000080e5110413/lun0/c5p1

Logical Drive #0 (104986 MB 110,086,324,224 Usable Bytes)		RAID 0+1 (Mirrored Stripe) [Write Cache Enabled, Stripe=16KB]							
Affinity Map: SAN Mapping									
Drive#	Span	Physical Enclosure	Physical Row	Physical Column	Channel	ID	Starting Block #	Span Size In blocks	
0	0	0	0	0	0	0	0	143341568	
1	0	0	0	2	0	1	0	143341568	
2	0	0	0	3	1	1	0	143341568	

Logical Drive #1 (139982 MB 146,781,765,632 Usable Bytes)		RAID 0 (Stripe) [Write Cache Enabled, Stripe=16KB]							
Affinity Map: SAN Mapping									
Drive#	Span	Physical Enclosure	Physical Row	Physical Column	Channel	ID	Starting Block #	Span Size In blocks	
0	0	0	1	1	1	2	0	143341568	
1	0	0	1	2	0	3	0	143341568	


Logical Drive #2 (139982 MB 146,781,765,632 Usable Bytes)		RAID 5 (Parity Stripe) [Write Cache Enabled, Stripe=16KB]							
Affinity Map: SAN Mapping									
Drive#	Span	Physical Enclosure	Physical Row	Physical Column	Channel	ID	Starting Block #	Span Size In blocks	
0	0	0	1	0	0	2	0	143341568	
1	0	0	2	1	1	4	0	143341568	
2	0	0	2	2	0	5	0	143341568	

Figure 3-83 Logical Subsystem Information (partial for 6.14 Firmware)

Logical Disk Information for RAID Subsystem at /hw/scsi/sc13d010									
Logical Drive #0 (70120 MB 73,526,149,120 Usable Bytes)				RAID 0 (Stripe) [Write Cache Enabled, Stripe=64KB]					
Affinity Map: SAN Mapping									
Background Tasks: No background jobs active									
Drive#	Span	Physical Enclosure	Physical Row	Physical Column	Channel	ID	Starting Block #	Span Size In blocks	
0	0	0	0	0	0	125	0	35901440	
1	0	0	0	1	1	124	0	35901440	
2	0	0	0	2	0	123	0	35901440	
3	0	0	0	3	1	122	0	35901440	
Logical Drive #1 (35060 MB 36,763,074,560 Usable Bytes)				RAID 0+1 (Mirrored Stripe) [Write Cache Enabled, Stripe=64KB]					
Affinity Map: SAN Mapping									
Background Tasks: No background jobs active									
Drive#	Span	Physical Enclosure	Physical Row	Physical Column	Channel	ID	Starting Block #	Span Size In blocks	
0	0	0	1	0	0	121	0	35901440	
1	0	0	1	1	1	120	0	35901440	
2	0	0	1	2	0	119	0	35901440	
3	0	0	1	3	1	118	0	35901440	

Figure 3-84 Logical Subsystem Information (partial for 7.01 and Later Firmware)

The 7.01 and later controller firmware returns a Background Task status line, which is not returned by controllers with 6.14 firmware.

The screen shows all configured LUNs, their status, mapping information, and how they are laid out. If one of the drives were removed, you would see the **ONLINE** indicator change to **CRITICAL**. If you were to view the screen shown in “Display Physical Subsystem Information” on page 114 under this condition, you would see the disk being rebuilt only if an action was taken (for example, if a disk had failed and was replaced). In that case, a rebuild operation should be in progress. If this wasn’t a test, the information in that screen would show that the drive is either off-line or the slot is empty, depending on how damaged the disk drive is.

Display Environmental Subsystem Information

When the **Display Environmental Subsystem Information** menu selection is made, the screen shown in Figure 3-85 (for controllers with 6.14 firmware) and Figure 3-86 (for controllers with 7.01 and later firmware) appears.

Battery Backup Status		Value		
Current power in hours (minutes)		30.1 (1805)		
Maximum power in hours (minutes)		30.1 (1805)		
Power threshold in hours (minutes)		24.1 (1444)		
Charge level (per cent)		100		
Hardware Version		1		
Battery Type		NiCAD		
Status		<ul style="list-style-type: none"> • No reconditioning cycle since power on. • Reconditioning cycle needed. 		

FAN #	Status	Speed	Enclosure ID (Switch Setting)
0 (RHS from rear)	Operational	Low	0
1 (LHS from rear)	Operational	Low	0
2	Not Present	N/A	0

Power Supply #	Status	Enclosure ID (Switch Setting)
0 (RHS from rear)	Operational	0
1 (LHS from rear)	Operational	0
2	Not Present	0

Temperature Sensor #	Status	OverTemp Warning	Current Temp Celsius/(F)	Enclosure ID (Switch Setting)
0	Operational	Normal	30 (86)	0

Alarm #	Status	Value	Enclosure ID (Switch Setting)
0	Operational	Normal	0

<i>No UPS data is available -- Not connected to one?!</i>						
---	--	--	--	--	--	--

Enclosure #	Status	Service	Primary Path	Secondary Path	Slots	Identifier Info
0	Operational	SES	Normal	Normal	12	<ul style="list-style-type: none"> • WWN: 50-05-0C-C0-00-00-16-7F • Enclosure ID: SGI • Product ID: TP9100 • Revision # B1

This screen will automatically refresh in approximately 10 seconds

Figure 3-85 Display Environmental Subsystem Information (6.14 Firmware)

Battery Backup Status		Value
Current power in hours (minutes)		31.1 (1869)
Maximum power in hours (minutes)		31.1 (1869)
Power threshold in hours (minutes)		24.1 (1444)
Charge level (per cent)		100
Hardware Version		1
Battery Type		NiCAD
Status		OK

Enclosure #	Status	Service	Primary Path	Secondary Path	Slots	Identifier Info
0	Operational	SES	Normal	Normal	12	<ul style="list-style-type: none"> • WWN: 50-05-0C-C0-00-00-16-7F • Enclosure ID: SG1 • Product ID: TP9100 • Revision # B1

FAN #	Status	Speed	Enclosure ID
0 (RHS from rear)	Operational	Low	50-05-0C-C0-00-00-16-7F
1 (LHS from rear)	Operational	Low	50-05-0C-C0-00-00-16-7F
2	Not Present	N/A	50-05-0C-C0-00-00-16-7F

Power Supply #	Status	Enclosure ID
0 (RHS from rear)	Operational	50-05-0C-C0-00-00-16-7F
1 (LHS from rear)	Operational	50-05-0C-C0-00-00-16-7F
2	Not Present	50-05-0C-C0-00-00-16-7F

Temperature Sensor #	Status	OverTemp Warning	Current Temp Celsius/(F)	Enclosure ID
0	Operational	Normal	30 (86)	50-05-0C-C0-00-00-16-7F

Alarm #	Status	Value	Enclosure ID
0	Operational	Normal	50-05-0C-C0-00-00-16-7F

No UPS data is available -- Not connected to one?!

This screen will automatically refresh in approximately 30 seconds

[\[Return to Main Page\]](#) [\[Log Off\]](#)

Figure 3-86 Display Environmental Subsystem Information (7.01 and Later Firmware)

The two windows for the different controller firmware versions contain the same information, but arranged differently. The windows display status information results of polling the enclosure (and all expansion enclosures) every 10 seconds.

Note: If expansion enclosures are attached, all of the same information would be reported for them as well, only with a different **Enclosure ID**.

Miscellaneous Functions

The Miscellaneous Functions menu is located on the main screen and is shown in Figure 3-87.

Miscellaneous Functions:

- **Display status of background jobs** - Shows status of all rebuilds, consistency checks, and initialization (formatting) jobs.
- **Flush controller(s) write cache to disk**.
- **Save current controller configuration** - Use this in combination with Load to clone a configuration.
- **Load controller configuration** - Use this in combination with Save to clone a configuration.
- **Flash new firmware onto controller(s)**
- **Flash new firmware onto supported disk(s)**
- **Adjust battery settings** - Allows setting thresholds, and forcing reconditioning or charging of BBU battery.
- **Clear configuration** - This clears (erases) all configuration data structures, and in the process, destroys all data.
- **Stop configurator service routine on host** - this kills the service job running on your host computer. If you select this option then nobody will be able to access the service routine until the job is manually restarted.

Figure 3-87 Miscellaneous Functions Menu (for 6.14 Firmware)

Miscellaneous Functions:

- **Display status of background jobs** - Shows status of all rebuilds, consistency checks, and initialization (formatting) jobs.
- **Flush controller(s) write cache to disk**.
- **Save current controller configuration** - Use this in combination with Load to clone a configuration.
- **Load controller configuration** - Use this in combination with Save to clone a configuration.
- **Flash new firmware onto controller(s)**
- **Flash new firmware onto supported disk(s)**
- **Adjust battery settings** - Allows setting thresholds, and forcing reconditioning or charging of BBU battery.
- **Clear configuration** - This clears (erases) all configuration data structures, and in the process, destroys all data.
- **Set the real-time Clock** - This sets the real time clock imbedded in the RAID controller to the time of your host system.
- **View the controller's internal event log** - This reports diagnostic messages saved in the internal RAID controller's event log.
- **Identify a Disk** - Select this function to identify a disk by causing the lights to slowly blink for 10 seconds.
- **Scan for new enclosures and disks** - You must perform this task after attaching new enclosures before they can be monitored.
- **Save controller debug dump** - This diagnostic function saves controller debug information to a data file. Please perform this action if instructed to do so by your supplier.
- **Stop configurator service routine on host** - this kills the service job running on your host computer. If you select this option then nobody will be able to access the service routine until the job is manually restarted.

Figure 3-88 Miscellaneous Functions Menu (for 7.01 and Later Firmware)

The Reporting Functions menu has the following selections, with the associated explanations on the indicated pages:

- “Display Status of Background Jobs” on page 124
- “Flush Controller(s) Write Cache to Disk” on page 127
- “Save Current Controller Configuration” on page 128

- “Load Controller Configuration” on page 129
- “Flash New Firmware Onto Controller(s)” on page 130
- “Flash New Firmware Onto Supported Disk(s)” on page 131
- “Adjust Battery Settings” on page 134
- “Clear Configuration” on page 136
- “Set The Real Time Clock (7.01 and Later Firmware Only)” on page 137
- “View the Controller’s Internal Event Log (7.01 and later Firmware Only)” on page 138
- “Identify a Disk (7.01 and later Firmware Only)” on page 140
- “Scan for New Enclosures and Disks” on page 143

Note: The Scan for New Enclosures and Disks feature is new with 7.75 firmware

- “Save Controller Debug Dump” on page 145

Note: The Save Controller Debug Dump feature is new with 7.75 firmware

- “Stop Configurator Service Routine on Host” on page 146
- “Automatic Restart (Reboot) on Failure Parameter” on page 147


Display Status of Background Jobs

When the **Display Status of Background Jobs** menu selection is made, the windows appear as shown in Figure 3-89 (for 6.14 firmware) and Figure 3-89 (for 7.01 and later firmware) appear.


Background Job Status RAID Controller at /hw/scsi/20000080e5110413/lun0/c5p1	
Check or Rebuild (Manual)	Not in progress
Initialization Status	LUN 0 Not in progress.
	LUN 1 Not in progress.
	LUN 2 Not in progress.
	LUN 3 Not in progress.
This screen will automatically refresh in approximately 30 seconds	
[Return to Main Page] [Log Off]	

Figure 3-89 Display Status of Background Jobs (6.14 Firmware)


Background Job Status RAID Controller at /hw/scsi/sc13d010

Logical Drive #0 (70120 MB, 73,526,149,120 Bytes) RAID 0 (Stripe) 

Action	Notes
Initialization	Initialization terminated by operator.

Logical Drive #1 (35060 MB, 36,763,074,560 Bytes) RAID 0+1 (Mirrored Stripe) 

Action	Notes
Rebuild	Not in progress
Consistency Check	consistency check terminated by operator.
Initialization	Initialization terminated by operator.
Online RAID Expansion	Not in progress

Logical Drive #2 (35060 MB, 36,763,074,560 Bytes) RAID 5 (Right Asymmetric) 

Action	Notes
Rebuild	Not in progress
Consistency Check	consistency check terminated by operator.
Initialization	Initialization terminated by operator.
Online RAID Expansion	Not in progress


This screen will automatically refresh in approximately 10 seconds

[\[Return to Main Page\]](#) [\[Log Off\]](#)


Figure 3-90 Display Status of Background Jobs (7.01 and Later Firmware)

Note: If a background init is in progress (7.75 firmware), the screen shown in Figure 3-91 appears.


Background Job Status RAID Controller at
/hw/scsi/20000080e5115f52/1un0/c3p1

Logical Drive #0 (34332 MB, 35,999,711,232 Bytes) RAID 3 (Right Asymmetric) 


Action	Notes
Rebuild	Not in progress
Consistency Check	Not in progress
Initialization	Not in progress
Background initialization <input type="button" value="Cancel All"/>	In progress 9.8 % complete (18 mins left).
Online RAID Expansion	Not in progress

Logical Drive #1 (34332 MB, 35,999,711,232 Bytes) RAID 5 (Right Asymmetric) 

Action	Notes
Rebuild	Not in progress
Consistency Check	Not in progress
Initialization	Not in progress
Background initialization	Has not been initiated.
Online RAID Expansion	Not in progress

Logical Drive #2 (34332 MB, 35,999,711,232 Bytes) RAID 0+1 (Mirrored Stripe) 

Action	Notes
Rebuild	Not in progress
Consistency Check	Not in progress
Initialization	Not in progress
Background initialization	Has not been initiated.
Online RAID Expansion	Not in progress

Logical Drive #3 (17166 MB, 17,999,855,616 Bytes) RAID 1 (Mirror) 

Action	Notes
Rebuild	Not in progress
Consistency Check	Not in progress
Initialization	Not in progress
Background initialization	Has not been initiated.
Online RAID Expansion	Not in progress

Figure 3-91 Display Status of Background Jobs—Background Init in Progress (7.75 Firmware)

These windows show you how initialization, rebuilds, or checks are progressing. The browser title bar (not shown in Figure 3-90) displays the date and time of the last poll. Press the **[Return to Main Page]** link to exit.

Flush Controller(s) Write Cache to Disk

When the **Flush Controller(s) Write Cache to Disk** menu selection is made, the screen shown in Figure 3-92 appears.

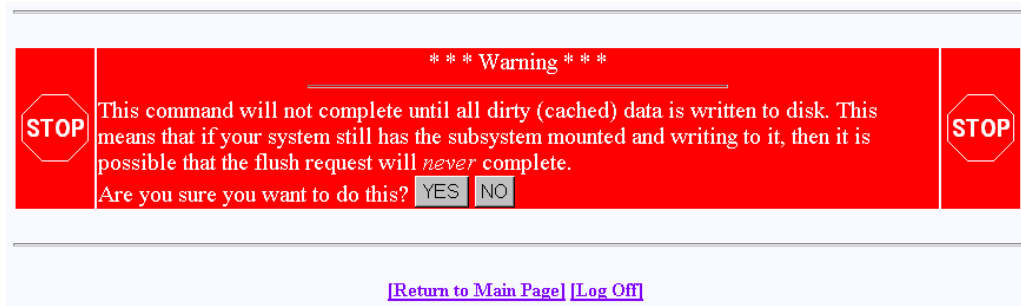


Figure 3-92 Cache Flush Warning

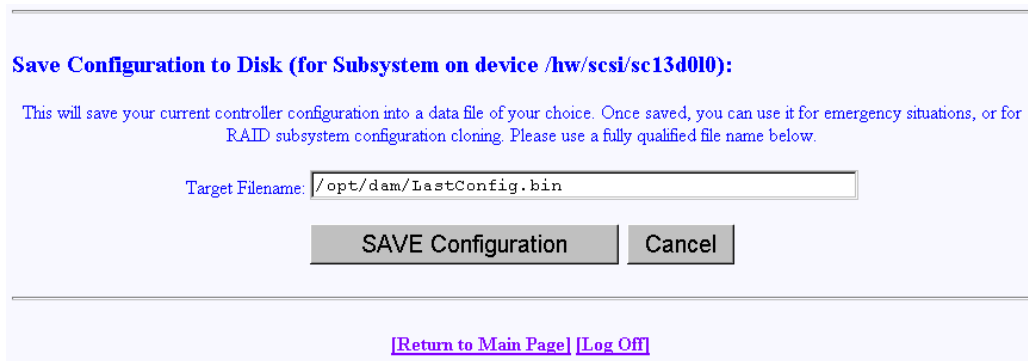
Click **YES** to force a cache flush. You would ordinarily perform this after all LUNs are unmounted, and before a power down. If you were to do a cache flush on mounted file systems that are servicing write requests, there would be a small risk that this request would never complete. However, it would be pointless to flush the cache in this situation, as the cache would instantly be dirty after the controller receives the next write command. If the cache flush is successful, the window shown in Figure 3-93 appears.



Figure 3-93 Cache Flush Successful

Save Current Controller Configuration

When the **Save Current Controller Configuration** menu selection is made, the screen shown in Figure 3-94 appears.



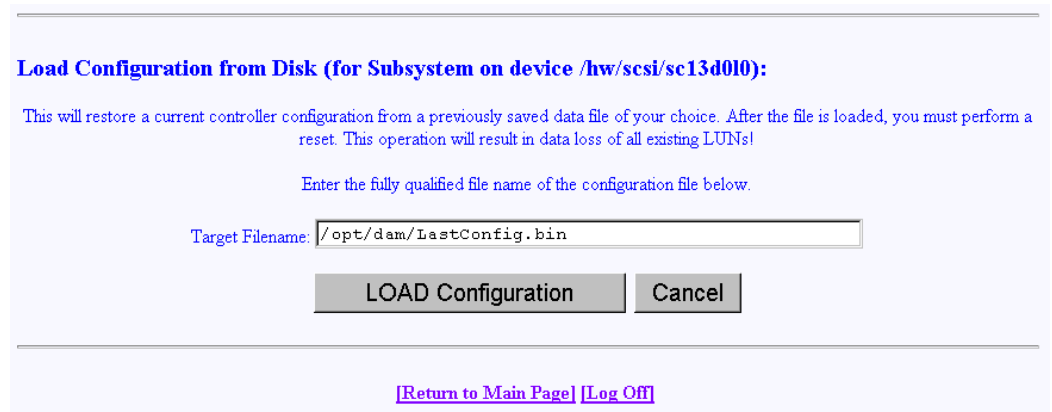
The screenshot shows a web interface for saving configuration. At the top, it says "Save Configuration to Disk (for Subsystem on device /hw/scsi/sc13d010):". Below this is a paragraph of text: "This will save your current controller configuration into a data file of your choice. Once saved, you can use it for emergency situations, or for RAID subsystem configuration cloning. Please use a fully qualified file name below." There is a text input field labeled "Target Filename:" containing the text "/opt/dam/LastConfig.bin". Below the input field are two buttons: "SAVE Configuration" and "Cancel". At the bottom of the interface, there are two links: "[Return to Main Page]" and "[Log Off]".

Figure 3-94 Save Current Controller Configuration

To save the current controller configuration to a file, type the name of the file in the area provided on the screen and click **SAVE Configuration**. The file is saved locally on the server that launched TPM, not the client machine running the Web browser. The file may be used later to restore or clone a RAID configuration.

Load Controller Configuration

When the **Load Controller Configuration** menu selection is made, the screen shown in Figure 3-95 appears.



The screenshot shows a web interface for loading a controller configuration. At the top, the title is "Load Configuration from Disk (for Subsystem on device /hw/scsi/sc13d010):". Below this, a warning message states: "This will restore a current controller configuration from a previously saved data file of your choice. After the file is loaded, you must perform a reset. This operation will result in data loss of all existing LUNs!". A prompt asks the user to "Enter the fully qualified file name of the configuration file below.". A text input field labeled "Target Filename:" contains the text "/opt/dam/LastConfig.bin". Below the input field are two buttons: "LOAD Configuration" and "Cancel". At the bottom of the page, there are two links: "[Return to Main Page]" and "[Log Off]".

Figure 3-95 Load Controller Configuration

This screen allows you to load the controller configuration file into another RAID array. You can use this feature to clone a RAID configuration. To clone a configuration this way, the two RAID subsystems must be exactly the same, including the disk drives.

The file is or loaded from the sever that launched TPM, not the client machine running the Web browser.

Flash New Firmware Onto Controller(s)

When the **Flash New Firmware Onto Controller(s)** menu selection is made, the screen shown in Figure 3-96 appears.

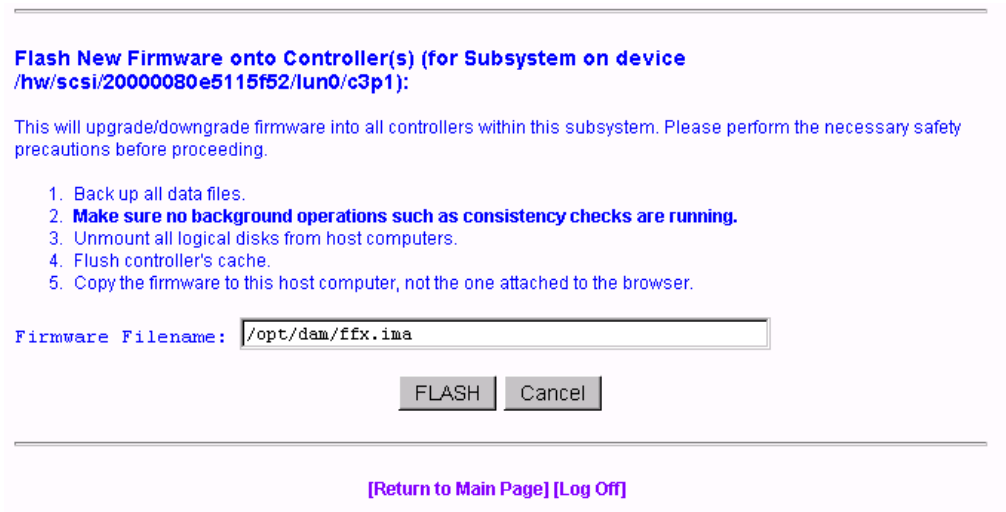


Figure 3-96 Flashing New Firmware

Warning: If you attempt to flash new firmware to the disk drive or controller while background jobs are running (the drive is transferring data), drive operation could become unpredictable and the drive may even become inoperable. In this event, drive data recovery may have to be performed by a professional data recovery lab.

After you enter the filename of the image, click **FLASH**. TPM then verifies the image size and type for your subsystem. If verification fails, a report is made.

Flashing the firmware takes a few seconds, following which TPM immediately jumps to the **Reset Controller** screen. Allow the reset to proceed.

Be sure you are aware of all the information that you need before upgrading (or downgrading) to certain firmware revisions.

Flash New Firmware Onto Supported Disk(s)

When the **Flash New Firmware Onto Supported Disk(s)** menu selection is made, the screen shown in Figure 3-97 appears (provided that the Operational Fault Management and Auto Rebuild Management functions on the View/Modify RAID Controller Configuration dialog box are not disabled—see Figure 3-52 on page 89). If the functions are already disabled when you click the **Flash New Firmware Onto Supported Disk(s)** menu selection, the dialog box shown in Figure 3-98 appears.

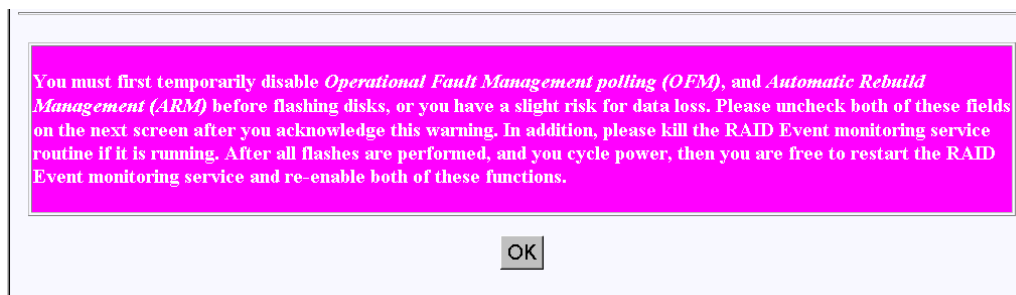


Figure 3-97 Flash New Firmware Onto Supported Disk(s) Dialog Box

Warning: The TPMWatch application must be terminated prior to updating disk drive firmware. Failure to do so may cause one or more disk drives to become inoperable.

Warning: If you attempt to flash new firmware to the disk drive or controller while background jobs are running (the drive is transferring data), drive operation could become unpredictable and the drive may even become inoperable. In this event, drive data recovery may have to be performed by a professional data recovery lab.

When you click **OK**, the View/Modify RAID Controller Configuration dialog box appears (see Figure 3-52 on page 89). Make sure you disable the Operational Fault Management and Auto Rebuild Management functions on this dialog box, then click **Apply**. You are returned to the main TPM menu.

Go to the Administrative Functions menu and click **Reset Controllers** (see Figure 3-57 on page 95). After the reset is complete, click the **Flash New Firmware Onto Supported Disk(s)** menu selection.

The dialog box shown in Figure 3-98 appears.

Disk Drive Flashing for RAID Subsystem at /hw/scsi/sc12d010

Firmware Filename:

Select	Channel	ID	Make	Model	FW Rev	Enclosure	Row	Col	# of Bloks
<input type="checkbox"/>	0	73	SGI	ST318304FC	2706	0	2	2	35964301
<input type="checkbox"/>	0	75	SGI	ST318304FC	2706	0	2	0	35964301
<input type="checkbox"/>	0	77	SGI	ST318304FC	2705	0	1	2	35964301
<input type="checkbox"/>	0	79	SGI	ST318304FC	2705	0	1	0	35964301
<input type="checkbox"/>	0	7B	SGI	ST318304FC	2705	0	0	2	35964301
<input type="checkbox"/>	0	7D	SGI	ST318304FC	2706	0	0	0	35964301
<input type="checkbox"/>	1	72	SGI	ST318304FC	2705	0	2	3	35964301
<input type="checkbox"/>	1	74	SGI	ST318304FC	2705	0	2	1	35964301
<input type="checkbox"/>	1	76	SGI	ST318304FC	2705	0	1	3	35964301
<input type="checkbox"/>	1	78	SGI	ST318304FC	2705	0	1	1	35964301
<input type="checkbox"/>	1	7A	SGI	ST318304FC	2705	0	0	3	35964301
<input type="checkbox"/>	1	7C	SGI	ST318304FC	2705	0	0	1	35964301

Comments:

1. Select a device, then press the FLASH button to initiate the firmware flash.
2. This function is designed for CE's only, and should never be run unless specifically instructed to do so.
3. Each drive will typically flash within 30 - 90 seconds, depending on the model.
4. Do not run any I/Os through the subsystem while a flash is being performed.
5. Be sure to flush disk cache before performing this operation.
6. **Make sure no background operations such as consistency checks are running.**
7. You must recycle power to the array after drive(s) are flashed.

- This feature works on supported IBM and Seagate disk drives only.
- Please dismount LUNs on this array from your hosts before continuing.

Once you select the FLASH button, then there will not be a confirming message, and flashing will begin immediately.

Figure 3-98 Flash New Firmware Dialog Box

Read all the instructions on the dialog box, select one or more devices to flash, then click **FLASH** to flash the firmware. When the process is complete, the screen shown in Figure 3-99 appears.

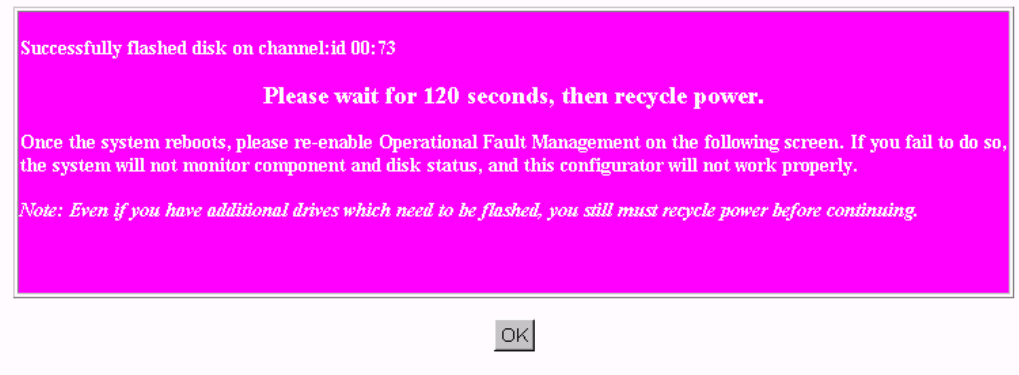


Figure 3-99 Flash New Firmware Complete Box

Wait 120 seconds, as the screen instructs, then cycle the power. After power up and reboot are complete, enable the Operational Fault Management and Auto Rebuild Management functions on the View /Modify RAID Controller Configuration dialog box.

Adjust Battery Settings

When the **Adjust Battery Settings** menu selection is made, the screen shown in Figure 3-100 appears for 6.14, 7.01, and 7.03 firmware. The screen shown in appears for 7.75 firmware.

Configure/Maintain Battery Settings (for Subsystem on device /hw/scsi/20000080e5110413/lun0/c5p1):

Make changes as necessary, then select an action button. Changes will be immediate, and may be made while background operations are occurring. The current battery status is:

(Full details can be seen on the enclosure status screen.)

Battery threshold (minutes)

[\[Return to Main Page\]](#) [\[Log Off\]](#)

Figure 3-100 Adjust Battery Settings (6.14, 7.01, and 7.03 Firmware)

Configure/Maintain Battery Settings (for Subsystem on device /hw/scsi/20000080e5115f52/lun0/c3p1):

Make changes as necessary, then select an action button. Changes will be immediate, and may be made while background operations are occurring. The current battery status is:

- Reconditioning cycle needed.

(Full details can be seen on the enclosure status screen.)

Battery threshold (minutes)

[\[Return to Main Page\]](#) [\[Log Off\]](#)

Figure 3-101 Adjust Battery Settings (7.75 Firmware)

When the remaining battery power (in minutes) falls below the Battery Threshold (minutes) value entered, a low battery power alarm is triggered and is displayed in the Battery Backup Status portion of the Display Environmental Subsystem Information window (see Figure 3-85 on page 120 and Figure 3-86 on page 121).

The buttons at the bottom of the screen operate as follows:

- Save Threshold: saves the battery threshold value entered in the text box.

Note: The specified threshold value entered must not exceed the battery's maximum power value in minutes.

- Recondition Battery: conditions the battery so that it can achieve maximum life. Conditioning involves fully discharging the battery, then recharging it.

Note: While the battery is being reconditioned, the system cache operates in the conservative cache mode (write-through) for the duration of the reconditioning process.

- Charge Battery: initiates a battery charge cycle.
- Shutdown Battery: this option shuts down charging to the battery. The Battery Backup Unit (BBU) maintains memory content in the presence of AC power failures, AC power glitches, and short power outages. This protection is important when the write-back cache is enabled, and data is waiting to be flushed to the disk drive. Therefore SGI does not recommend or support use of the Shutdown Battery option provided in TPM 1.2 with 7.75 firmware. If a shutdown of the battery is performed, memory retention is not guaranteed if power loss occurs. Any data in the controller's on-board cache memory will be lost.

Clear Configuration

When the **Clear Configuration** menu selection is made, the screen shown in Figure 3-102 appears.




Figure 3-102 Clear Configuration

Warning: Read the information in the screen before you take any action.

To erase the configuration structures, click **CLEAR**.

The confirmation message shown in Figure 3-103 appears.

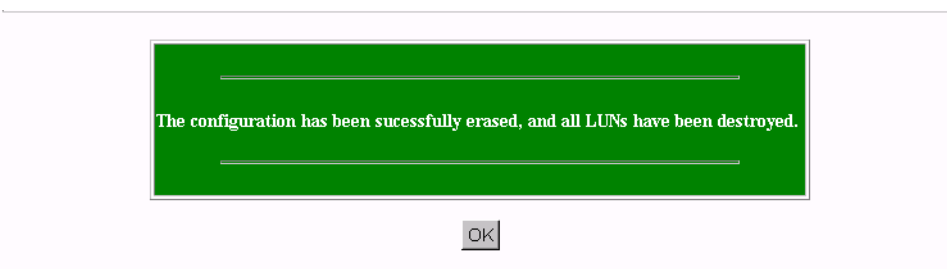


Figure 3-103 Clear Configuration

Set The Real Time Clock (7.01 and Later Firmware Only)

When the **Set Real Time Clock** menu selection is made, the dialog box shown in Figure 3-104 appears.




Figure 3-104 Set Real Time Clock

Click **SET CLOCK** to synchronize the controller clock with the host computer.

The confirmation screen shown in Figure 3-105 appears.




Figure 3-105 Set Real Time Clock Confirmation Dialog Box.

View the Controller's Internal Event Log (7.01 and later Firmware Only)

When the **View the Controller's Internal Event Log** menu selection is made, the dialog box shown in Figure 3-106 appears.

Event history since unit was powered on (*Only unreported events are shown*) :

Seq Num	Time Date	Type Date	Ch	ID	LUN	Code	Parm	Sense Data	Meaning/Action Required
25	09:41:53 01/22/2001	Info	0	115	0	14	0		A hard disk has been removed. Cause: User removed an unconfigured physical device. An unconfigured physical device failed. A controller was removed. A controller powered off. <i>Action Required: Replace the device if needed.</i>
26	09:42:03 01/22/2001	Info	0	115	0	13	0		A new hard disk has been found. Cause: A physical device has been powered on. A new physical device has been added. Controller was powered on. Controller was added. System has rebooted. <i>Action Required: None</i>

Figure 3-106 View Controller's Internal Event Log

Each time you bring up this dialog box, it displays the events that occurred since the last time the dialog box was brought up.

The buttons at the bottom of the dialog box provide these functions:

- **Save All:** saves all events to a raw data text file named `eventhistory.log`.
- **Save New:** saves the new events to a raw data text file named `eventhistory.log`.

Note: A **Save New** operation overwrites the current `eventhistory.log` file. If you do not want to overwrite it, it must be renamed before you perform the operation.

- **Append All:** appends all events to the raw data text file named `eventhistory.log` file.

- Append New: appends new events to the raw data text file named `eventhistory.log` file.
- View All Events: displays all events, including old and new events.

Note: If you are running `gamevent` on the system, this button must be used to see all log events, because `gamevent` is continually emptying the event log.

Identify a Disk (7.01 and later Firmware Only)

For the **Identify a Disk** menu selection to work properly, Operational Fault Management (OFM) must be enabled. This is accomplished by enabling the Operational Fault Management function in the View/Modify RAID Controller Configuration dialog box (see Figure 3-52 on page 89). Also, at least one of the SES disk drives must be present and operational (see Figure 3-107 and Figure 3-108).




Figure 3-107 SES Drive Locations (Rackmount Configuration)




Figure 3-108 SES Drive Locations (Tower Configuration)

When the **Identify a Disk** menu selection is made, the dialog box shown in Figure 3-109 appears.

Physical Devices in RAID Subsystem at `/hwr/scsi/20000080e5114243/1un0/c35p1`

Select	Channel	ID	Make	Model	Enclosure	Row	Col
<input type="radio"/>	0	73	IBM	DNEF-309170	0	2	2
<input type="radio"/>	0	75	IBM	DNEF-309170	0	2	0
<input type="radio"/>	0	77	IBM	DNEF-309170	0	1	2
<input type="radio"/>	0	79	IBM	DNEF-309170	0	1	0
<input type="radio"/>	0	7B	IBM	DNEF-309170	0	0	2
<input checked="" type="radio"/>	0	7D	IBM	DNEF-309170	0	0	0
<input type="radio"/>	1	72	IBM	DNEF-309170	0	2	3
<input type="radio"/>	1	74	IBM	DNEF-309170	0	2	1
<input type="radio"/>	1	76	IBM	DNEF-309170	0	1	3
<input type="radio"/>	1	78	IBM	DNEF-309170	0	1	1
<input type="radio"/>	1	7A	IBM	DNEF-309170	0	0	3
<input type="radio"/>	1	7C	IBM	DNEF-309170	0	0	1
Total Drives: 12							

Select a single device to identify, then press the PING button to initiate the identification, which will cause the drive light to blink. The disk will stop blinking either after 10 seconds, or when you select another device to identify.

Figure 3-109 Identify a Disk

Click PING to identify a selected drive.

Scan for New Enclosures and Disks

When the **Scan for New Enclosures and Disks** menu selection is made (7.75 firmware), the dialog box shown in Figure 3-110 appears.

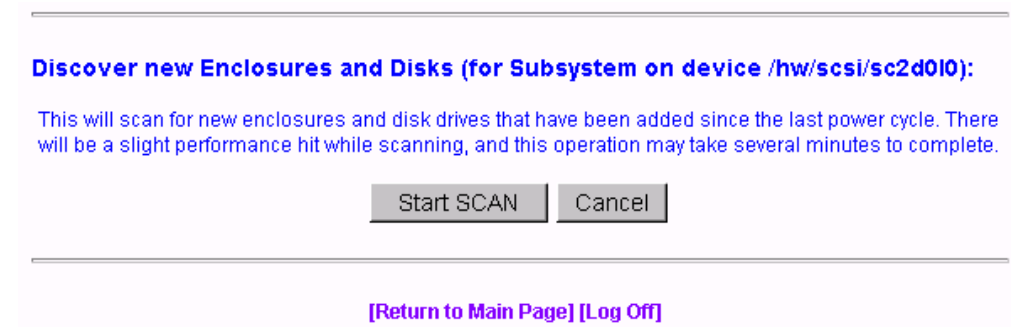


Figure 3-110 View Controller's Internal Event Log (7.75 Firmware)

This feature allows users to add one or more disk enclosures to a configured system while the system continues to operate. After the enclosure or enclosures have been added to the system, clicking **Start SCAN** on the screen shown in Figure 3-110 starts the SES monitoring process for the new enclosure. The user can then configure the additional disk capacity without restarting the system.

Additional enclosures are added to the configured system using the following procedure.

1. Check for ID conflicts. Each enclosure and disk drive must have a unique ID.
2. Resolve any ID conflicts.
3. Connect drive channels from the existing system to the new enclosure or enclosures.
4. Supply power to the new enclosure or enclosures. This causes a Loop Initialization Primitive (LIP) on the drive channel to notify the controller that new disk drives have been added to the fibre loop.
5. Wait for the controller to supply power to the disk drives. All disk drives must have completed the spin-up process before proceeding.
6. Issue the scan for additional enclosures. This may be an option incorporated in the configuration utility, or issued as a direct SCSI command.

7. After the SES process completes polling the loop, the new enclosures and disk drives are ready for configuration.
8. If the new enclosure(s) are not detected, remove and insert a disk drive from the new enclosure. Removing and inserting a disk drive generates an LIP, and the new disk drives will join the existing fibre loop.
9. Issue the scan for additional enclosures a second time.

Save Controller Debug Dump

When the **Save Controller Debug Dump** menu selection is made (7.75 firmware), the dialog box shown in Figure 3-110 appears.

Dump Debug Configuration to Disk (for Subsystem on device /hw/scsi/sc2d010):

This feature saves information about errors that might cause a controller abort. If a controller aborts due to a firmware detected error, information relating to the nature of the error is saved in the controller NVRAM. This allows that dump to be saved to a data file, which can be sent to your supplier for analysis.

Please note that the file is in binary format, so make sure that you transfer it to your email system in binary mode.

Dump Filename:

[\[Return to Main Page\]](#) [\[Log Off\]](#)

Figure 3-111 Dump Debug Configuration to Disk (7.75 Firmware)

This feature records controller state information when an abort occurs. After the abort has completed, you can click **SAVE Debug Dump** to retrieve and analyze the abort information to help determine why the abort occurred. The information is generated while a controller abort is in progress and recorded to NVRAM and a disk drive, if enabled. The aborting controller blocks requests from the partner controller so that the dump can be generated. The Debug Dump is written to a Reserved Disk Area (RDA) on one selected disk drive. The contents of the dump are not user configurable.

The default operation is that when the dump is complete, status information is written to the Debug Dump header.

The Debug Dump data entry consists of an abort code, an error code, and an event flag, and is overwritten by each successive abort. The abort code specifies where in the firmware image the abort occurred. The abort code is written to NVRAM when the controller is aborting and generates an event. The abort code takes up two bytes of data.

Debug Dump data can be retrieved by going to Miscellaneous Functions and clicking on Save Controller Debug Dump after an abort has completed and the controller is back online.

Stop Configurator Service Routine on Host

In TPM 1.0, when the **Stop Configurator Service Routine on Host** menu selection is made, the screen shown in Figure 3-112 appears.

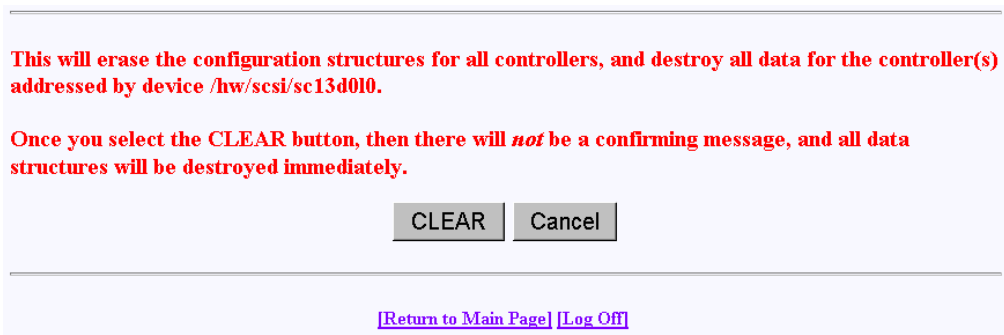


Figure 3-112 TPM 1.0 Stop Service Screen

This screen provides an emergency shutdown routine for TPM 1.0. After you click **CLEAR**, the TPM 1.0 service routine is immediately shut down. You are not prompted with an “Are-You-Sure Message.” Once the service routine has been killed, no commands can be issued, and anyone with a Web browser receives the standard error message indicating the host is not found.

In TPM 1.2, when the **Stop Configurator Service Routine on Host** menu selection is made, an emergency shutdown routine is initiated for TPM 1.2. You are not prompted with an “Are-You-Sure” message. Once the TPM service routine has been killed, no commands can be issued, and anyone with a Web browser receives the standard error message indicating the host is not found. After TPM 1.2 is shut down, the message shown in Figure 3-113 appears.

Service routine has been halted on the host, and a log entry has been made.

Figure 3-113 TPM 1.2 Shutdown Message

Automatic Restart (Reboot) on Failure Parameter

Controller parameters implemented in firmware version 7.75 and later only control the behavior of the Automatic Restart on Failure (ROF) feature. These parameters are:

- ROF Reboot Count: this parameter specifies the maximum number of times a controller attempts automatic restart to recover from firmware detected errors (see the ROF Reboot Count parameter in Figure 3-56 on page 93).
- ROF Rearm Interval: this parameter specifies the minimum time interval that the controller must operate before refreshing the number of restart attempts (see the ROF Rearm Interval parameter in Figure 3-56 on page 93)

When a controller reaches the maximum restart attempts, the automatic restart feature becomes disabled until the value is refreshed. Any subsequent firmware detected errors require manual intervention to recover the controller.

The possible values for the time interval between refreshing the number of restart attempts range from DISABLED to seven days. The values are displayed when you click the ROF Rearm Interval dropdown box. To select a particular value, depress the right mouse button while scrolling to the desired time interval and then release the button to select the value.

To disable ROF, set the ROF Reboot Count to 0 and the ROF Rearm Interval to something other than DISABLED. When the automatic restart feature is disabled, manual intervention is required to recover from fatal firmware detected errors. Manual intervention may involve physically removing and replacing the failed controller.

The ROF parameters take effect immediately, without resetting the controllers.

TPMWatch Event Monitor and Logger

TPMWatch is a support program designed to poll RAID subsystems and report their health to an output file. The file can then be used by a user-supplied program or shell script to provide notification in the event a component fails or goes offline. The program works by issuing commands to the controller to report status information for all LUNS, disk drives, and enclosure components (fans, power supplies, batteries, and so on).

To minimize performance impact, provide the greatest amount of flexibility to incorporate TPMWatch in external routines, the program is designed with the following considerations:

- User-defined polling period (in seconds).
- Generates only 11 I/Os.
- User supplies the status file name upon invocation.
- User supplies the optional history file name upon invocation.
- The status file is pure ASCII text, so the result can easily be interpreted by a shell script.

Running the TPM 1.0 Version of TPMWatch From Linux and IRIX

To run TPMWatch, make sure the \$DAM_HOME environment variable is set, and you invoke the program from root. The syntax is:

```
tpmwatch [-P Freq][[-d][[-H HistFile][[-T LogFile]-S StatusFile  
-D DeviceFile[-M Addresses][[-L]
```

Where:

-P Freq: Frequency in seconds between each poll.

-d: turns on status debugging features by sending information to stderr.

-H HistoryFile: Optional history file, which is appended to the last status information record. Do not supply this parameter if you do not wish to create a history file.

-T LogFile: Sends error text to a user-specified file.

-S StatusFile: The health of the subsystem (see record layout in Table A-1).

-D DeviceFile: Device to check (for example, /hw/scsi/sc37d010). This is the raw or pass-through driver attached to any LUN on the subsystem. It doesn't make any difference which LUN you pick. Furthermore, the program still operates correctly if the LUN that it is nailed to goes offline or is even removed.

Note: For Linux users, the device to check (for example, /dev/dam5) can be found by running the "damioscan" program, located in the /opt/dam directory. When running TPM, this information is displayed in the Select Physical Device (Configuration Function Menu) window.

-M Addresses: E-mail addresses used for sending alarm messages.

Note: When specifying more than one e-mail address, the entire list of e-mail addresses must be enclosed in double-quotes (""). For example, -M "user1@domain user2@domain...".

-L: Send events to system log file.

An example syntax is:

```
/opt/dam/tpmwatch -D /hw/scsi/sc37d010 -P 10 -L &
```

Note: A minimum of two parameters must be specified, one of which must be DeviceFile. If only the DeviceFile parameter is specified, TPMWATCH will not start, and the syntax usage message will be reported.

Running the TPM 1.2 Version of TPMWatch From IRIX

To run TPMWatch, make sure the \$DAM_HOME environment variable is set, and you invoke the program from root. The syntax is:

```
tpmwatch [-P Freq] [-d] [-e] [-i] [-s] [-w] [-H HistFile] [-T  
LogFile] -S StatusFile -D DeviceFile [-M Addresses] [-L]
```

Where:

-P Freq: Frequency in seconds between each poll

-C Freq: Frequency in hours to autosave configuration

-d: Enable debug mode to stderr

-e: E-mail error messages (requires -E flag and 7.X FW or above)

-i: E-mail informational messages (requires -E flag and 7.X FW or above)

-s: E-mail severe messages (requires -E flag and 7.X FW or above)

-w: E-mail warning messages (requires -E flag and 7.X FW or above)

-H HistFile: Optional event history file. Do not supply this parameter if you do not wish to create a history file.

-T LogFile: Sends error text to user-specified log file

-S StatusFile: Status file

-D DeviceFile: Device to check (for example, /hw/scsi/sc37d010). This is the raw or pass-through driver attached to any LUN on the subsystem. It doesn't make any difference which LUN you pick.

-M Addresses: E-mail addresses used to send alarm messages

Note: When specifying more than one e-mail address, the entire list of e-mail addresses must be enclosed in double-quotes (""). For example, -M "user1@domain user2@domain...".

-L: Send events to system log file. An example syntax is:

```
/opt/dam/tpmwatch -D /hw/scsi/sc37d010 -P 10 -L &
```

Note: A minimum of two parameters must be specified, one of which must be DeviceFile. If only the DeviceFile parameter is specified, TPMWATCH will not start, and the syntax usage message will be reported.

Running the TPM 1.0 Version of TPMWatch From Windows

To run TPMWatch from Windows, you must be at the directory where the executable resides. The syntax is:

```
tpmwatch [-P Freq][-d][-H HistFile][-T LogFile]-S StatusFile
-D DeviceFile[-E SMPT -M Addresses][-L]
```

Where:

-P Freq: Frequency in seconds between each poll.

-d: turns on status debugging features by sending information to `stderr`.

-H HistoryFile: Optional history file, which is appended to the last status information record. Do not supply this parameter if you do not wish to create a history file.

-T LogFile: Sends error text to a user-specified file.

-S StatusFile: The health of the subsystem (see record layout in Table A-1).

-D DeviceFile: Device to check (for example, `/hw/scsi/sc37d010`). This is the raw or pass-through driver attached to any LUN on the subsystem. It doesn't make any difference which LUN you pick. Furthermore, the program still operates correctly if the LUN that it is nailed to goes offline or is even removed.

-E SMPT: Mail server address (name of mail server, for example, `mail.xyz.com`).

-M Addresses: E-mail addresses used for sending alarm messages. For example:

"<user_name@xyz>" (quotes and angle brackets are required)

-L: Send events to system log file.

To invoke TPMWATCH, bring up a DOS window, change directories to the location where TPMWATCH is installed, and type the command to start TPMWATCH. An example is:

```
tpmwatch -D /dev/h2i00100 -P 10 -L
```

Note: A minimum of two parameters must be specified, one of which must be DeviceFile. If only the DeviceFile parameter is specified, TPMWATCH will not start, and the syntax usage message will be reported.

Table A-1 Record Layout

Byte #	Size	Description
0	8	Polling date in YYYYMMDD format.
8	6	Polling time in HHMMSS format (24-hour clock).
14	1	Subsystem summary status. Set to + if all is OK, - otherwise.
15	1	Fan Status: + if All On-Line, otherwise a single digit representing number of failed fans.
16	1	PSU Status: + if All On-Line, otherwise a single digit representing number of failed power supply units.
17	1	Temperature status. + if within specifications, - if over-temperature warning (or temperature sensor failed).
18	1	+ if UPS found and on-line, - if found and off-line, blank if no UPS.
19	1	- if battery is currently discharging, + otherwise.
20	1	+ if no further enclosure alarms, - otherwise (this tests all ESS pages for any alarm state and detects such things as a FC path failure).
21	2	Master/Slave controller status (see Table A-2).
23	4	Partner State (see Table A-3).
27	32 x 1	Status of each possible System Drive (LUN): + If On-line - If Off-line (and defined) C if Critical (and defined) ? if in unknown state

Table A-1 (continued) Record Layout

Byte #	Size	Description
		Blank if no none defined
59	8 X 16 X 1	Status of each Physical Device Disk[Channel][ID] (0:0, 0:1 .. 7:15) + If On-line D if drive is DEAD R if drive is Rebuilding S if drive is in Stand By state ? if drive in unknown state Blank if empty slot
187	1	New line character 0Ah.

Master/Slave Controller Status Table

Table A-2 Master/Slave Controller Status Table

Value	Meaning
00	Slave disabled, or simplex.
Master States During Redundancy	
10	Initial state during dual-active boot.
11	Default master got first ping.
12	Master in negotiation.
13	Negotiation done; waiting for ping.
14	Negotiation done; received ping.
15	CC nexus established.

Table A-2 (continued) Master/Slave Controller Status Table

Value	Meaning
16	Insufficient memory.
17	Waiting for debounce ping.
18	Waiting for replacement notification.
Failed Slave States	
20	Slave failed.
21	Slave controller is ejected.
22	Slave controller removal is detected.
23	Slave controller is inserted.
24	Command slave controller is inserted.
25	Command slave controller to pause.
26	Paused slave ready to resume.
27	Paused slave controller failed.

Table A-2 (continued) Master/Slave Controller Status Table

Value	Meaning
States During Redundancy	
40	Default slave is slave.
41	Slave entered monitoring mode.
Failed Master States	
80	Master failed.
81	Master controller is ejected.
82	Master controller removal detected.
83	Master controller is inserted.
84	Command master controller insert.
85	Relinquish control.

Partner Status Table

Table A-3 Partner Status Table

Value	Partner State	Notes
0000	No Partner	Controller running in simplex mode.
0100	Booting	From power-up.
0101	Booting	Partner is replacement controller.
0200	Active	Controller-Controller nexus.
0201	Active	Partner is survivor.
0300	Failed	Ping time-out.
0301	Failed	Negotiation - get chunk failure.

Table A-3 (continued) Partner Status Table

Value	Partner State	Notes
0302	Failed	Negotiation - SCSI communication failed or wrong cables, or firmware versions/builds are different.
0303	Failed	Negotiation - host ID mismatch.
0304	Failed	Negotiation - SLIP/DIFFL/FBR mismatch.
0305	Failed	Negotiation - disk channels available mismatch.
0306	Failed	Negotiation - host channels available mismatch.
0307	Failed	Negotiation - firmware version mismatch.
0308	Failed	Negotiation - firmware type mismatch.
0309	Failed	Negotiation - memory size mismatch.
030a	Failed	Negotiation - memory read of partner failed.
030b	Failed	Negotiation - MS_INTNEG command to partner failed.
030c	Failed	Kill Partner command received.
030d	Failed	Partner failed during failback TID handover.
030e	Failed	Partner didn't enter nexus after negotiation complete.
030f	Failed	Partner failed for unknown reason.
0310	Failed	Failed Write Back Synchronization Failed on Channel 0.
0311	Failed	Failed Write Back Synchronization Failed on Channel 1.
0312	Failed	Failed Write Back Synchronization Failed on Channel 2.
0313	Failed	Failed Write Back Synchronization Failed on Channel 3.
0314	Failed	Failed Write Back Synchronization Failed on Channel 4.
0315	Failed	Failed Write Back Synchronization Failed on Channel 5.
0316	Failed	Negotiation - firmware build mismatch.
0317	Failed	Negotiation - device channel cables are crossed.

Table A-3 (continued) Partner Status Table

Value	Partner State	Notes
0320	Failed	Hot pull of partner detected while nexus active.
0321	Failed	Partner absent at boot.
0322	Failed	Power failed before failover finished.
0323	Failed	Power failed before relinquish finished.
0341	Failed	Controller-controller locking had unrecoverable SCSI error.
0400	Removed	Partner is removed.
0500	Inserted	Partner is inserted.

Error Codes

Table B-1 lists error codes associated with the direct commands (MDACIOCTL). These codes can be reported when using TPM and are provided as reference information.

Table B-1 Error Codes

Key	ASC	ASQ	Description
05	24	00	Illegal field in SCSI CDB.
09	00	02	Bounds error.
09	80	50	Parity error detected. Internal firmware error.
09	81	00	Last check had a good completion.
09	81	01	Drive is write-protected.
09	81	02	Specified device not present. Invalid enclosure ID specified.
09	81	03	Host block count is zero.
09	81	04	Unimplemented opcode from host.
09	81	05	Invalid system drive number specified.
09	81	06	if (CommandCode == MDACIOCTL_MORE), a rebuild or check is already in progress. Else controller is busy.
09	81	09	Invalid parameter (reserved bytes) in CDB.
09	81	0B	Enclosure device not ready.
09	81	0C	Initialization is in progress.
09	81	0D	Consistency check is in progress.
09	81	0F	Command issued to C1 in a duplex system. OFM is disabled.
09	81	10	Check consistency or expansion is in progress.
09	81	11	Maximum allowable number of system drives is already defined.
09	81	12	System drive to migrate is in CRITICAL mode.
09	81	13	System drive to migrate is not in ONLINE mode.
09	81	14	Controllers are in dual-active mode for Simplex MORE.
09	81	15	Failover still in progress.

Table B-1 (continued) Error Codes

Key	ASC	ASQ	Description
09	81	16	Migrate command issued to system drive with more than one span.
09	81	17	Disk drive to add is already part of a system drive.
09	81	18	Disk drive to add is not in STANDBY (Hot Spare) mode.
09	81	19	Multiple system drives defined on PDs to enlarge.
09	81	1A	Invalid parameter in data list.
09	81	1B	The rebuild/migrate rate is improperly set to 0xff (255).
09	81	1C	A drive in system drive to expand is also in other system drive with more than one span.
09	81	1D	The drive to add is too small.
09	81	1E	Cannot get temporary memory.
09	81	1F	COD write to disk failed.
09	81	20	Controller not ready.
09	81	21	Controller not ready - waiting for start unit.
09	81	22	Controller not ready - can't get ready.
09	81	23	Controller not ready - command didn't complete.
09	81	24	No more devices to report.
09	81	25	More than 32 enclosures attached, cannot return data.
09	81	26	Insufficient buffer space to return all data. Allocation length too small in CDB.
09	81	27	Specified request not supported. Invalid page code requested.
09	81	28	Device scan in progress for new direct command.
09	81	29	Invalid RAID type.
09	81	2A	Specified device not found.
09	81	2B	Maximum # of COD groups (64) already used.
09	81	30	Invalid range for Config2 parameter.

Table B-1 (continued) Error Codes

Key	ASC	ASQ	Description
09	81	31	A Config2 parameter can not be changed at this time.
09	81	40	Program image failed do to invalid image file.
09	81	41	Program image failed do to EEPROM write failure.
09	81	42	Program image failed do to EEPROM read/compare failure.
09	81	50	Device specified in UDD not previously defined.
09	81	51	Maximum number of physical drives reached.
09	81	52	More than 32 SDs configured.
09	81	53	Requested data larger than allocation length.
0	9	81	54 invalid COD group ID.
09	81	55	Reserved field used or invalid value in field.
09	81	56	New RDN already in use.
09	81	57	Specified transfer size too small.
09	81	58	Top level only can be deleted.
09	81	59	Last defined device only can be deleted.
09	81	5A	Physical device already specified.
09	81	5B	Specified physical device is not configured.
09	81	5C	Cannot change RAID type.
09	81	5D	Cannot change stripe size.
09	81	5E	Cannot change device number.
09	81	5F	Physical device no available.
09	81	60	No groups present.
09	81	61	Bad number of drives to add.
09	81	62	User sent IDD (spanned LUN).
09	81	63	PDD must be configured prior to MORE request.

Table B-1 (continued) Error Codes

Key	ASC	ASQ	Description
09	81	64	Start LBA must be 0.
09	81	65	Entire PDD must be used.
09	81	66	SDD sent with MORE has bad field(s); illegal state change for logical device.
09	81	67	No SES device present (Operational Fault Management MUST be enabled for this command to work).
09	81	68	Invalid SAN map.
09	81	69	Skipped an XLDD number.
09	81	6A	Drive exists but is unconfigured.
09	81	6B	Invalid stripe size in configuration.
