

TPM Installation Instructions and User's Guide
for SGI® TP9100

007-4382-004

CONTRIBUTORS

Written by Jenn McGee

Illustrated by Chrystie Danzer

Production by Karen Jacobson

Engineering contributions by Dave Lucas, Sammy Wilborn, Rob Novak, and Van Tran.

COPYRIGHT

© 2003 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

IRIX, SGI, the SGI logo, and Silicon Graphics are registered trademarks and Altix and SGI ProPack are trademarks of Silicon Graphics, Inc., in the United States and/or other countries worldwide.

Internet Explorer, Microsoft, and Windows 2000 are trademarks or registered trademarks of Microsoft Corporation. JavaScript is a trademark of Sun Microsystems. Linux is a trademark of Linus Torvalds. Mozilla, Netscape, and Netscape Navigator are trademarks or registered trademarks of Netscape Communications. Qlogic is a trademark of Qlogic Corporation. Red Hat is a trademark of Red Hat, Inc. AcceleRAID and Mylex are trademarks of LSI Logic Corporation. Tivoli is a registered trademark of IBM. Veritas is a registered trademark of Veritas Software Corporation.

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
003	August 2002 Updated to support Release 5.0; TP9100 (1Gb TP9100) and TP9100 (2Gb TP9100)
004	February 2003 Updated to support Release 5.15 (2Gb Single Host Port RAID Controllers and Single Port RAID Loopback LRC), Release 5.4 (SN-McKinley LINUX Support) and TPM 1.4

Contents

List of Figures	ix
List of Tables	xiii
About This Guide	xv
Audience	xv
Structure of this Guide	xv
Product Support	xvi
Related Publications	xvi
Obtaining Publications	xvi
Conventions	xvii
Reader Comments	xvii
1. Introduction	1
Product Overview	1
Features	2
Software Packaging	2
IRIX/Linux Software Packaging	2
2. Installing, Configuring, and Running TPM on IRIX and Linux	5
Host Prerequisites	5
Installing the Software	6
Installing TPM on IRIX Platforms	6
Installing TPM on Linux Platforms	7
Uninstalling TPM on IRIX Platforms	8
Uninstalling TPM on Linux Platforms	8
Configuring Usernames and Passwords	9
Default Security	9
Enhanced Security (Encrypted)	10

Invoking TPM15
Setting the Environment Variable for IRIX and Linux.15
Starting the SGI TP9100 Array Manager for IRIX and Linux15
Setting Up the Browser16
Logon Screen17
Additional Considerations18
3. Using TPM19
Main Menu19
Firmware Feature Information21
Features for 8.40 Firmware Release21
Features for 7.75 Firmware Release on the FFx Controller22
Configuration Functions23
Select Physical Device (Configuration Function Menu)24
Modify Disk Device Status (Configuration Function Menu)26
Create/Destroy/Expand Logical Drives (Configuration Function Menu)29
Initialize LUN(s) (Configuration Function Menu).52
Initialize LUNs in Background (Configuration Function Menu).55
Configure Host-to-LUN Mapping (Configuration Function Menu)58
Installing Device Drivers for New LUNS (Configuration Function Menu)64
Administrative Functions65
View/Modify RAID Controller Configuration.65
Reset Controller(s).69
Gracefully Bring a Controller Off Line71
Gracefully Bring a 2nd Controller On Line73
Perform Data Consistency Operations on a LUN74
Enable/Disable Write Cache for LUN(s)77
Modify/Purge Non-Volatile WWN Tables.78

Reporting Functions	79
Set Default Screen Refresh Rate	80
Display (Dual) Controller Status	80
Topology Query	81
Display SCSI/Fibre Device Information	82
Display FULL Subsystem Configuration Information	84
Display Physical Subsystem Information	86
Display Logical Subsystem Information	88
Display Environmental Subsystem Information	89
Miscellaneous Functions	90
Display Status of Background Jobs	91
Flush Controller(s) Write Cache to Disk	93
Save Current Controller Configuration	94
Load Controller Configuration	96
Flash New Firmware Onto Controller(s)	97
Flash New Firmware Onto Supported Disk(s)	99
Adjust Battery Settings	101
Clear Configuration	103
Set The Real Time Clock	104
View the Controller's Internal Event Log	105
Identify a Disk.	106
Scan for New Enclosures and Disks	107
Save Controller Debug Dump	109
Stop Configurator Service Routine on Host	110
Automatic Restart (Reboot) on Failure Parameter	110
A. TPMWatch Event Monitor and Logger	113
Running the TPM Version of TPMWatch From IRIX	113
B. Event and Error Codes	117
Event and Error Code Table	117

List of Figures

Figure 2-1	Miscellaneous Functions Menu	11
Figure 2-2	passwd.bin File Error Message	12
Figure 2-3	Maintain Security Features Window	13
Figure 2-4	Cannot Log on Screen	14
Figure 2-5	Logon Screen	17
Figure 3-1	Main Menu (partial screen)	20
Figure 3-2	Select Physical Device Dialog Box	24
Figure 3-3	Modify Disk Device Status Dialog Box (partial view taken on 1Gb TP9100 enclosure)	27
Figure 3-4	Create/Destroy/Expand Logical Drives Dialog Box	30
Figure 3-5	New Configuration Screen, 1Gb TP9100 12 drive enclosure	31
Figure 3-6	Drives Assigned into Pack A (12 drive, 1Gb TP9100 RAID Subsystem view).	35
Figure 3-7	All Drive Packs Created (12 drive, 1Gb TP9100 RAID Subsystem view)	36
Figure 3-8	Select Drive Packs to Include in LUN	37
Figure 3-9	LUN is Being Defined	38
Figure 3-10	LUN is Defined	39
Figure 3-11	All LUNs Defined	40
Figure 3-12	Updating the RAID Controller Warning Screen	41
Figure 3-13	LUN Online Screen.	41
Figure 3-14	Create/Destroy/Expand Logical Drives Screen	43
Figure 3-15	Add LUN(s) Screen (12 drive, 1Gb TP9100 RAID Subsystem view)	44
Figure 3-16	Add LUN(s) Warning Screen	45
Figure 3-17	Add LUN(s) Confirmation Screen	45
Figure 3-18	Select LUN for Deletion	47
Figure 3-19	Delete LUN Warning Screen	48
Figure 3-20	Create/Destroy/Expand Logical Drives Screen	50

Figure 3-21	Expand LUN Screen (12 drive, 1Gb TP9100 RAID Subsystem view) .	51
Figure 3-22	Warning Message	52
Figure 3-23	Initialize LUN	53
Figure 3-24	Initialization Confirmation.	54
Figure 3-25	Initialization Progress	55
Figure 3-26	Initialize LUNs in Background	56
Figure 3-27	Initialize LUNs in Background	57
Figure 3-28	Storage Area Network	59
Figure 3-29	SAN Mapping Topology view - 1Gb TP9100 RAID Controller (Single Port)	62
Figure 3-30	Install New Device Drivers Warning Screen.	64
Figure 3-31	Administrative Functions Menu	65
Figure 3-32	View/Modify Controller Configuration, 2Gb TP9100, 8.40 Firmware	66
Figure 3-33	View/Modify Controller Configuration, 2Gb TP9100 Dual Port, 8.40 Firmware.	67
Figure 3-34	View/Modify Controller Configuration, 2Gb TP9100 Dual Port, 8.40 Firmware.	68
Figure 3-35	Reset Controllers Warning Screen.	69
Figure 3-36	Controller Booting Window	70
Figure 3-37	Controller Responding Window	71
Figure 3-38	Warning Screen	72
Figure 3-39	Warning Screen	73
Figure 3-40	Data Consistency Check/Restore	74
Figure 3-41	Operation Started Dialog Box	75
Figure 3-42	Background Initialization Status Screen	76
Figure 3-43	Enable Write Cache For LUN(s) Dialog Box.	77
Figure 3-44	World Wide Name Table for RAID Controller	78
Figure 3-45	Reporting Functions Menu	79
Figure 3-46	Default Screen Refresh Rate	80
Figure 3-47	Controller Status Information	81
Figure 3-48	Topology Query (16 drive, 2Gb TP9100 RAID Subsystem View) .	82
Figure 3-49	Display SCSI/Fibre Device Information	83
Figure 3-50	Controller Information Values (partial screen)	84

Figure 3-51	Controller Tunable Parameters	84
Figure 3-52	GroupConfig Structure	85
Figure 3-53	Physical Subsystem Information (TP9100 1G, 12 bay enclosure view)	86
Figure 3-54	Physical Subsystem Information (TP9100 1G, 12 bay enclosure view)	87
Figure 3-55	Logical Subsystem Information	88
Figure 3-56	Display Environmental Subsystem Information Window	89
Figure 3-57	Miscellaneous Functions Menu	90
Figure 3-58	Display Status of Background Jobs—Background Init in Progress	92
Figure 3-59	Cache Flush Warning	93
Figure 3-60	Cache Flush Successful	94
Figure 3-61	Save Current Controller Configuration	94
Figure 3-62	Configuration Save Successful	96
Figure 3-63	Load Controller Configuration	97
Figure 3-64	Flashing New Firmware	98
Figure 3-65	Flash New Firmware onto Supported Disk(s) Dialog Box	99
Figure 3-66	Flash New Firmware Dialog Box.	100
Figure 3-67	Flash New Firmware Complete Box.	101
Figure 3-68	Adjust Battery Settings Window	102
Figure 3-69	Clear Configuration Screen	103
Figure 3-70	Clear Configuration Message.	104
Figure 3-71	Set Real Time Clock Dialog Box	104
Figure 3-72	Set Real Time Clock Confirmation Dialog Box.	105
Figure 3-73	View Controller’s Internal Event Log Dialog Box	105
Figure 3-74	Identify a Disk Dialog Box	107
Figure 3-75	View Controller’s Internal Event Log Dialog Box	108
Figure 3-76	Dump Debug Configuration to Disk Dialog Box	109
Figure 3-77	TPM Shutdown Message	110

List of Tables

Table 1-1	Files in the Distribution (IRIX and Linux)	2
Table 3-1	Select Physical Device Field Definitions	26
Table 3-2	Modify Disk Device Status Field Definitions	28
Table B-1	Error-Event Codes118

About This Guide

This guide provides information on preparing, installing, configuring, and using the Total Performance Manager (TPM) application for the TP9100 1Gb/s FFX RAID controllers with 7.75 firmware and TP9100 2Gb/s FFX -2 RAID controllers with 8.40 and later firmware.

Note: For information on TP9100 1Gb/s FFX RAID controllers with firmware level 6.14 and 7.03, see previous versions of this manual.

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 the IRIX and Linux 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 and Linux” — 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, “Event and 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.

Related Publications

The following documents contain additional information that may be helpful:

- *SGI Total Performance 9100 Storage System Owner's Guide*
- *SGI Total Performance 9100 Installation and Maintenance Instructions*
- *SGI Total Performance 9100 Storage System User's Guide*

Obtaining Publications

You can obtain SGI documentation in the following ways:

- See the SGI Technical Publications Library at <http://docs.sgi.com>. Various formats are available. This library contains the most recent and most comprehensive set of online books, release notes, man pages, and other information.
- If it is installed on your SGI system, you can use InfoSearch, an online tool that provides a more limited set of online books, release notes, and man pages. With an IRIX system, select **Help** from the Toolchest, and then select **InfoSearch**. Or you can type `infosearch` on a command line.

- You can also view release notes by typing either `grelnotes` or `relnotes` on a command line.
- You can also view man pages by typing `man <title>` on a command line.

Conventions

The following conventions are used throughout this publication:

Convention	Meaning
<code>command</code>	This fixed-space font denotes literal items such as commands, files, routines, path names, signals, messages, and programming language structures.
<i>variable</i>	Italic typeface denotes variable entries and words or concepts being defined.
user input	This bold, fixed-space font denotes literal items that the user enters in interactive sessions. (Output is shown in nonbold, fixed-space font.)
[]	Brackets enclose optional portions of a command or directive line.
...	Ellipses indicate that a preceding element can be repeated.
<code>manpage(x)</code>	Man page section identifiers appear in parentheses after man page names.
GUI element	This font denotes the names of graphical user interface (GUI) elements such as windows, screens, dialog boxes, menus, toolbars, icons, buttons, boxes, fields, and lists.

Reader Comments

If you have comments about the technical accuracy, content, or organization of this document, contact SGI. 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 is located at the bottom of each page.)

You can contact SGI 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 Web page:
<http://docs.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.

SGI values 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 TP9100 external RAID disk subsystems. TPM consists of two programs that run natively on an IRIX or Linux system that is physically attached to a RAID storage subsystem. All communication is through the TP9100 Fibre Channel (FC) RAID 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 1Gb/s TP9100 and 2Gb/s TP9100 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. This includes Microsoft's Internet Explorer 5.x and above, Netscape Navigator 4.x and above, and the Mozilla browser. 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 and Linux. 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 webpages.
<code>tpm</code>	TPM service routine.
<code>dameventlog</code>	Log file (automatically created if not found).

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

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

Installing, Configuring, and Running TPM on IRIX and Linux

This chapter explains how to install and configure the TPM application on the IRIX operating system. 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 5.0, Netscape's 4.x, and Netscape's Mozilla 1.0 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 IRIX versions:
 - IRIX 6.5.8 or later for the 1Gb/s TP9100 (FFX) RAID Controllers
 - IRIX 6.5.16 or later for the 2Gb/s TP9100 (FFX-2) RAID Controllers
- Operating system. TPM supports the following Linux versions:
 - Linux 7.75 firmware only for the 1Gb/s TP9100 (FFX) RAID Controllers

- Linux, SGI environment 7.2 or later with SGI ProPack SGI Altix 3000 series servers using only TPM 1.4 or later
- 2Gb/s TP9100 (FFX-2) RAID Controllers

Please note the following information on operating system support:

- TP9100 support for IA32 Linux has been frozen at Red Hat release 6.2 with ProPack 1.3, using TPM 1.0 with controller firmware 6.14 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 and 7.03. Support for the TP9100 on newer releases of Windows will not be developed.
- TPM version 1.4 supports Single Port Controllers. TPM 1.3 and earlier cannot be used on Single Port Controllers.
- TP9100 support for IA64 Linux is only provided when using SGI Altix 3000 series servers running SGI Linux environment 7.2 or later with SGI ProPack 2.1 or later.
- TPM 1.4 or later supports SGI Linux environment 7.2 or later with SGI ProPack 2.1 or later.

Installing the Software

This section describes how to install and uninstall the TPM software on the IRIX platform.

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. Enter the following command to launch the IRIX installation tool (`inst`) to install the `sgi_tpm` software image:

```
inst -f /CDROM/irix/dist/sgi_tpm
```

4. To specify the package, enter:

```
list
```

5. To install the software, enter:

```
install
```

6. At the Install subsystem prompt, enter:

```
sgi_tpm
```

7. To complete the installation, enter:

```
go
```

8. Enter the following command to exit the install program:

```
quit
```

The exit operation is automatically performed.

9. Enter the following command to unmount the CD-ROM:

```
umount /CDROM
```

10. Remove the CD-ROM from the CD-ROM drive.

Installing TPM on Linux Platforms

Use the following procedure to install TPM 1.4 on SGI Altix 3000 series servers running SGI Linux environment 7.2 or later with SGI ProPack 2.1 or later:

1. Log in as root.
2. Insert the TPM 1.4 or later CD-ROM into the CD-ROM drive.
3. Enter the following command to mount the CD-ROM:

```
mount /dev/cdrom/mnt/cdrom
```

4. To launch the Red Hat Package Manager (`rpm`), which installs the TPM 1.4 or later software, enter the following command followed by the `rpm` package filename. The command below is valid for the TPM 1.4 release only and the filename will change on later releases:

```
rpm -iv /mnt/cdrom/linux/rpm/SGItpm_linux_ia64_tpm1.4.rpm
```

Note: In certain instances, a `--force` flag may have to be used in the `rpm`.

5. Enter the following command to unmount the CD-ROM:

```
umount /mnt/cdrom
```

6. Remove the CD-ROM from the CD-ROM drive.

Uninstalling TPM on IRIX Platforms

Use the following procedure to uninstall TPM on IRIX platforms:

1. Log in as root.
2. Enter the IRIX `versions -remove` command as follows to remove TPM software:

```
versions remove sgi_tpm
```

Note: IRIX users, if you are removing an TPM 1.3 or earlier release, the uninstall will remove your existing `passwd.txt` file. If you have entries in your `passwd.txt` file and are planning on installing TPM again at a later date it is recommended that you save your `passwd.txt` file to another directory.

Uninstalling TPM on Linux Platforms

Use the following procedure to uninstall TPM on Linux platforms:

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

```
rpm -ev SGItpm_LINUX-1.4.0-1
```

Configuring Usernames and Passwords

TPM's security mechanism has been modified on releases 1.4 and later. This was done for the following reasons:

- Tighten control of what users are provided access to the configurator and attached storage.
- Reduce the possibility of several users configuring the same RAID controller configuration at the same time and then causing potential problems.

TPM 1.4 and later releases now provide the following options: use an encrypted password file, allow only one user to be logged in at one time, restrict client access at IP level by blocking out a specific IP address or a group of IP addresses.

The TPM software by default will now bounce the user logged in to TPM, once another user (from another server IP address) does a log in to TPM. TPM will allow multiple users to log in from the same server without a bounce occurring.

New installations of TPM 1.4 or later now require that a default username `admin` and password `sgj` be used after installation in order to log in to TPM. TPM options can be viewed by entering the following command when in the `/opt/dam` directory `./tpm -/`. The following is only the security part of this option:

```
#./tpm -/
```

```
Usage: tpm [-B] [-E "123.*.*.*"] [-I "123.1.2.3"]
```

Where:

- B mode where incoming user cannot bounce off current user
- E *ip* excludes access to given range of ip addresses (for example, "100.*.*.*")
- I *ip* includes access to given range of ip addresses (for example, "100.10.50.*")

Default Security

For current TPM 1.3 and earlier users doing an upgrade installation to TPM 1.4 or later, the upgrade installation will not modify your current `passwd.txt` file.

New users installing TPM 1.4 or later will be provided default security using the `passwd.txt` file.

The system administrator can modify the `passwd.txt` file as necessary. In the `passwd.txt` 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 and everything is case-sensitive. The username, password, or both fields can be blank. By default, the `passwd.txt` file has a default username of `admin` and a default password of `sgi` designated.

Note: If you do not have any username(s) or password(s) setup, it is recommended that upon installation of TPM 1.4 or later, you take advantage of the new enhanced security option. See the “Enhanced Security (Encrypted)” section below.

Enhanced Security (Encrypted)

The Enhanced Security method in TPM 1.4 and later releases provides the option of using password encryption. This is implemented by adding the provided `passwd.bin` file to the `/opt/dam` (DAM HOME directory).

Note: The `passwd.bin` file is not created on installation of TPM and must be copied from the `security` directory on the TPM CD (TPM 1.4 or later CDs only).

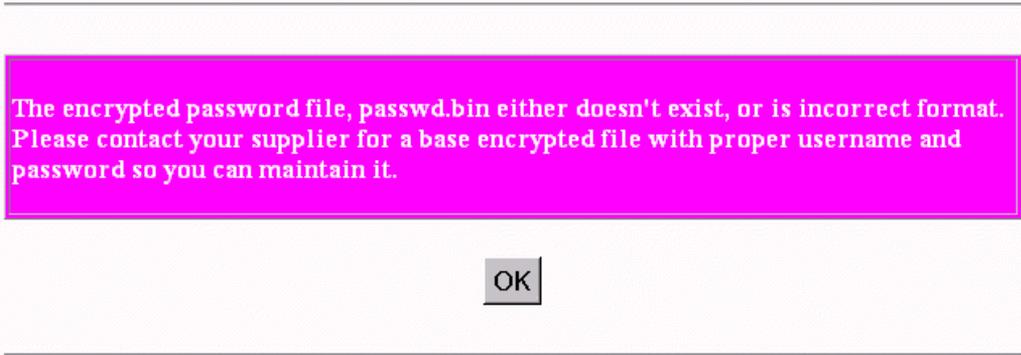
Once the `passwd.bin` file is placed into the `/opt/dam` directory, upon the next start of TPM, the default username `admin` and password `sgi` must be used to log in to TPM. Once logged in to TPM, the capability of adding in multiple users, each with their own unique username and password is provided in the **Maintain Security Features** menu. The **Maintain Security Features** menu is listed under **Miscellaneous Functions** on the TPM main page as shown in Figure 2-1 on page 11.

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.
- [Maintain Security Features](#) - Allows you to maintain program access user accounts and options.
- [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 2-1 Miscellaneous Functions Menu

Note: If you do not have the Enhanced Security `passwd.bin` file in the `/opt/dam` directory and you try to select **Maintain Security Features** from the main page, the screen shown in Figure 2-2 will appear indicating that you only have the `passwd.txt` file in the `/opt/dam` directory.



The encrypted password file, passwd.bin either doesn't exist, or is incorrect format. Please contact your supplier for a base encrypted file with proper username and password so you can maintain it.

OK

Figure 2-2 passwd.bin File Error Message

When using the passwd.bin file, the maximum number of user accounts is 10. If the **User Name** field is blank, then the entry will be ignored. A blank password is acceptable, but not advisable for security reasons. Administrator accounts should be the only ones set with the ability to maintain this screen.

Note: If you are upgrading from a previous release of TPM, once you move the passwd.bin file into your DAM_HOME directory, the passwd.bin file will override all usernames and passwords defined in the passwd.txt file. The default username and password must then be used to log in to TPM. Once logged in to TPM, all username(s) and password(s) that were set up in your passwd.txt file can be added to the **Maintain Security Features** window as shown in Figure 2-3. After you select **update**, all entries made will be saved and encrypted into the passwd.bin file. After this is done and verified, it is recommended that you delete the passwd.txt file or at a minimum its contents.

Security Management

This release has room for 10 user accounts. If the **User Name** field is blank, then the entry will be ignored. A blank password is acceptable, but not advisable for security reasons. Accounts with Admin set have the ability to maintain this screen.

Note: If you delete all of the accounts, you will not be able to log back in!

User Name	Password	Full Name	E-Mail Address	Admin	Shell
admin	sgi			<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>
				<input type="checkbox"/>	<input type="checkbox"/>
				<input type="checkbox"/>	<input type="checkbox"/>
				<input type="checkbox"/>	<input type="checkbox"/>
				<input type="checkbox"/>	<input type="checkbox"/>
				<input type="checkbox"/>	<input type="checkbox"/>
				<input type="checkbox"/>	<input type="checkbox"/>
				<input type="checkbox"/>	<input type="checkbox"/>
				<input type="checkbox"/>	<input type="checkbox"/>
				<input type="checkbox"/>	<input type="checkbox"/>

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

Figure 2-3 Maintain Security Features Window

Note: If you delete all of the accounts, you will not be able to log back in until you remove the `passwd.bin` file from the `DAM_HOME (/opt/dam)` directory.

Note: If you wish to go back to the default `passwd.txt` password scheme, delete (or move) the `passwd.bin` file from the `/opt/dam` (DAM HOME directory). If you moved the `passwd.txt` file you will need to move it back to the `/opt/dam` (DAM HOME directory). The `passwd.txt` file will then be read and used.

Note: When TPM is started with the `-B` option (recommended), only one user at a time will be allowed to log in and use TPM. Users who try to log in when another user is already logged in will get the following message returned:

Only one user allowed at a time currently user username logged in
at server IP address on date, time.

See Figure 2-4.

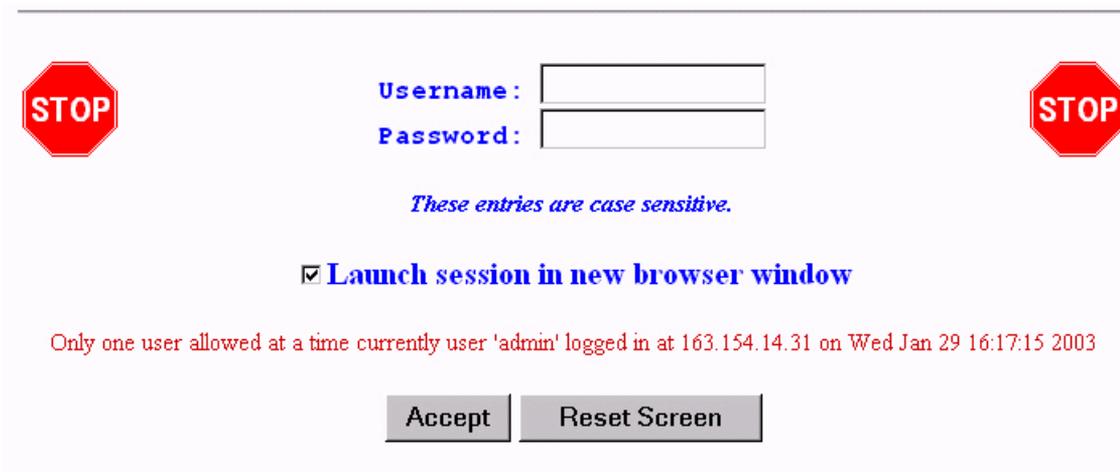


Figure 2-4 Cannot Log on Screen

Note: When using the `-B` flag, the following conditions may result in a user not being able to log in to TPM:

A user just kills the browser window and does not properly log out from tpm.

The browser crashes during a tpm session.
The browser process gets killed.

Workaround: The system administrator must kill the tpm process and restart tpm.

Note: When TPM is started without the `-B` option, the second (incoming) user logging in to TPM will bounce off the first (current) user. This will result in any uncommitted changes that the first user was in the process of doing getting terminated. SGI recommends using the `-B` option to avoid this situation from occurring.

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 for IRIX and Linux, follow these steps:

1. Log on as root (or have root access).

2. For csh, or tcsh shell, enter:

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

3. For sh, bash, or ksh shell, enter:

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

```
export DAM_HOME
```

Starting the 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 recommended 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 -B flag is recommended to use on startup of TPM to limit only one user at a time.

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

Note: When enabled, debug mode may cause the application to run more slowly and is not supported for non-debug useage.

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 its 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 system, 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 2-5.

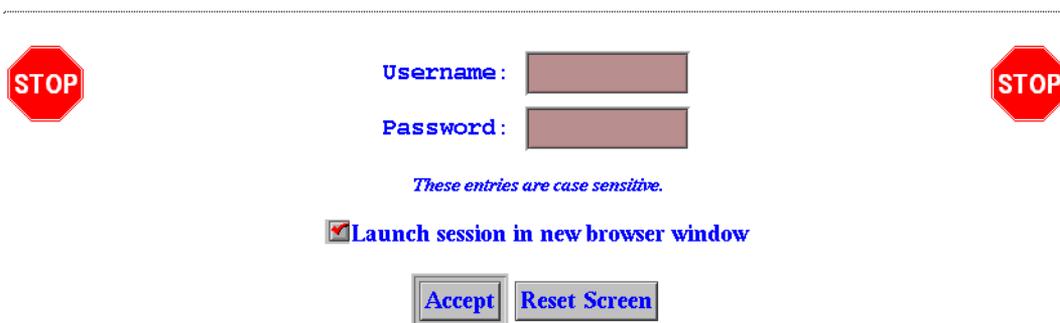


Figure 2-5 Logon Screen

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` or `passwd.bin`. If the username and passwords match the file contents, the main window screen appears (see Figure 3-1 on page 20). 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 his or her 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. 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 multiuser 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 using FTP, 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 TP9100 model (1Gb FFX RAID Controller or 2Gb FFX2 RAID Controller) and the firmware level 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.

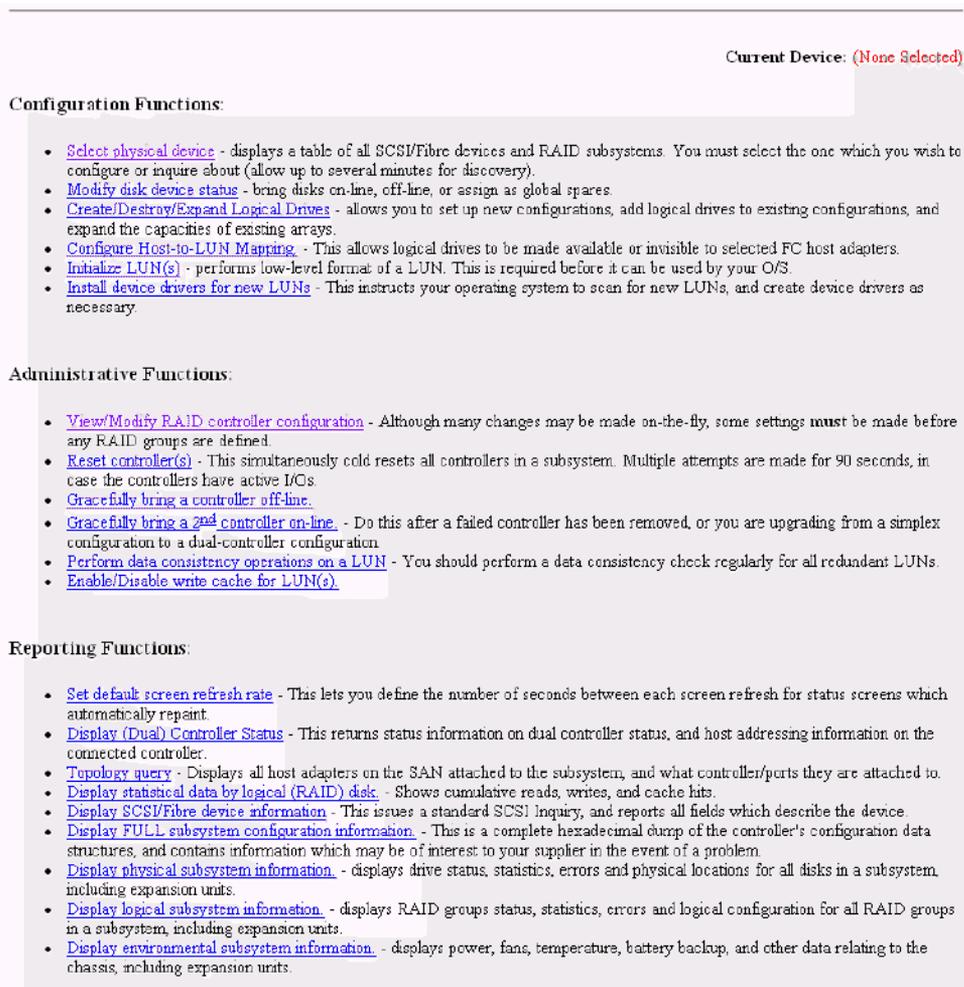


Figure 3-1 Main Menu (partial screen)

Figure 3-1 shows a portion of the main menu window that appears once you log on. The main menu windows are different for 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 23
- “Administrative Functions” on page 66
- “Reporting Functions” on page 80
- “Miscellaneous Functions” on page 91

Firmware Feature Information

This section lists the new firmware and TPM features.

Features for 8.40 Firmware Release

The following lists features for the 8.40 firmware release.

- AutoFlash
If a replacement controller (in a duplex configuration) has a different firmware image than the replaced (failed) controller, the replacement controller is automatically reflashed to match the firmware level of the surviving controller.
- F_Port
Pre-8.x firmware supports controller attachment to an FL_port on a switch (a server or switch port that implements a FC-AL). With 8.x and later firmware, a point-to-point connection to an F_port on a switch and an N_Port on a server is supported. Point-to-point connections are higher performance connections than FC-AL connections.
- SANmapping-256
SANmapping will now support up to 256 host nodes (was 64 host nodes). The number of LUNs supported (across multiple controller pairs) remains unchanged.
- Selectable LUN deletion

Any LUN within an array can be deleted regardless of the order in which it was created.

Note: Selectable LUN deletion requires TPM 1.4 and 8.40 or later RAID controller firmware.

Features for 7.75 Firmware Release on the FFx Controller

The following lists features for the 7.78 firmware release on the FFx Controller.

- 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.

- 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 would then be used to retrieve the debug dump.

- 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 is set to once per 24 hours.

- 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.

- 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.

- No Controller Reset Requirement on LUN Addition or Deletion

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

- 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.

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 24
- “Modify Disk Device Status (Configuration Function Menu)” on page 26
- “Create/Destroy/Expand Logical Drives (Configuration Function Menu)” on page 29
- “Initialize LUN(s) (Configuration Function Menu)” on page 52
- “Initialize LUNs in Background (Configuration Function Menu)” on page 55
- “Configure Host-to-LUN Mapping (Configuration Function Menu)” on page 58
- “Installing Device Drivers for New LUNS (Configuration Function Menu)” on page 65

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 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/20000050cc004b71/lun0/c6p1 Node# 20000050cc004b71				RAID RAID	SGI	TP9100 FFX2	DAC960FFx2 (0) 20-00-00-50-CC-00-4B-71	DAC960FFx2 (1) 20-00-00-50-CC-00-4B-71
<input type="checkbox"/>	/hw/scsi/20000080e5116db1/lun0/c6p1 Node# 20000080e5116db1				RAID RAID	MYLEX	DACARMRB 3_13884	DAC960FFx (0) 20-00-00-80-E5-11-6D-B1	DAC960FFx (1) 20-00-00-80-E5-11-69-A5
<input type="checkbox"/>	/hw/scsi/sc8d010	8	0	0	RAID RAID	SGI	TP9100 F PSEUDO	DAC960FFx2 (0) 20-00-00-50-CC-00-67-66	DAC960FFx2 (1) 20-00-00-50-CC-00-67-66
<input type="checkbox"/>	/hw/scsi/20000050cc004b71/lun0/c6p3 Node# 20000050cc004b71				RAID RAID	SGI	TP9100 FFX2	DAC960FFx2 (1) 20-00-00-50-CC-00-4B-71	DAC960FFx2 (0) 20-00-00-50-CC-00-4B-71
<input type="checkbox"/>	/hw/scsi/20000080e51169a5/lun0/c6p3 Node# 20000080e51169a5				RAID RAID	MYLEX	DACARMRB PSEUDO	DAC960FFx (1) 20-00-00-80-E5-11-69-A5	DAC960FFx (0) 20-00-00-80-E5-11-6D-B1

Select

Reset Screen

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 not 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, followed by the `ioconfig -f /hw` command. For more information, see the `scsiha(1m)` and `ioconfig(1m)` man pages.

If these 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 on page 20) 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	The physical device path is the pass-through device driver name (<code>/hw/scsi/</code>) for that particular device.
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)	<p>The Vendor ID and Product ID parameters are returned by a standard SCSI Inquiry.</p> <p>On TP9100 1Gb/sec FC RAID Controllers with 7.75 firmware:</p> <p>The first field of the Product ID typically starts with DACAMRB. The second field describes the type and size of RAID LUN you have. For example, the selected LUN at <code>/hw/scsi/sc2d010</code> (first row of Figure 3-2 on page 24) points to a 102,996MB RAID-0+1 system disk. The digits before the “B” indicate the 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.</p> <p>On TP9100 2Gb/sec FC RAID Controllers:</p> <p>The Vendor ID will be SGI and The Product ID field will be TP9100 FFX2. Logical Unit 0 information for RAID type and size of LUN are not provided.</p>
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 is assigned by the controller manufacturer.

Note: The Linux select device screen will differ from the IRIX select device screen. Some of the above fields will differ and others are not available and will contain N/A.

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 appears.

Modify Disk Device Status for RAID Subsystem at /hw/scsi/sc134010

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 (Backmount view, rotate 90° clockwise if in tower)

<p>861 37538304FC [2705] S/N: 3EL00890 W/N: 200002037653457 100MHz, 2 Ports, FC-AL Drive Speed: 10016 RPM Physical: 37560 Usable: 37530 LoopID=125 (7db) Chan=C ID=125 LUNs: None</p> <p>ON LINE HOT SPARE UNCONFIGURED</p>	<p>861 37538304FC [2705] S/N: 3CB0136L W/N: 20000000002A3709 100MHz, 2 Ports, FC-AL Drive Speed: 10016 RPM Physical: 37500 Usable: 34686 LoopID=124 (7Ch) Chan=L ID=124 LUNs: None</p> <p>ON LINE HOT SPARE UNCONFIGURED</p>	<p>861 37538304FC [2705] S/N: 3EL0094T W/N: 20000002037659786 100MHz, 2 Ports, FC-AL Drive Speed: 10016 RPM Physical: 37500 Usable: 37530 LoopID=123 (7B) Chan=C ID=123 LUNs: None</p> <p>ON LINE HOT SPARE UNCONFIGURED</p>	<p>861 37538304FC [2705] S/N: 3EL0081C W/N: 2000020376598A1 100MHz, 2 Ports, FC-AL Drive Speed: 10016 RPM Physical: 37560 Usable: 37530 LoopID=122 (7Ab) Chan=L ID=122 LUNs: None</p> <p>ON LINE HOT SPARE UNCONFIGURED</p>
<p>861 37538304FC [2705] S/N: 3EL00890 W/N: 2000020376598ED 100MHz, 2 Ports, FC-AL Drive Speed: 10016 RPM Physical: 37560 Usable: 37530 LoopID=121 (79b) Chan=C ID=121 LUNs: None</p> <p>ON LINE HOT SPARE UNCONFIGURED</p>	<p>861 37538304FC [2705] S/N: 3EL00890 W/N: 20000000002765906b 100MHz, 2 Ports, FC-AL Drive Speed: 10016 RPM Physical: 37560 Usable: 37530 LoopID=120 (78h) Chan=L ID=120 LUNs: None</p> <p>ON LINE HOT SPARE UNCONFIGURED</p>	<p>861 37538304FC [2705] S/N: 3EL0096E W/N: 2000000203765972F 100MHz, 2 Ports, FC-AL Drive Speed: 10016 RPM Physical: 37500 Usable: 37530 LoopID=119 (77b) Chan=C ID=119 LUNs: None</p> <p>ON LINE HOT SPARE UNCONFIGURED</p>	<p>861 37538304FC [2705] S/N: 3EL0084V W/N: 2000020376598D6 100MHz, 2 Ports, FC-AL Drive Speed: 10016 RPM Physical: 37560 Usable: 37530 LoopID=118 (76b) Chan=L ID=118 LUNs: None</p> <p>ON LINE HOT SPARE UNCONFIGURED</p>

Legend:

- Vendor ID (Pv-xxxx) (Firmware Revision)
- World Wide Name
- Clock Speed, # of Ports, Interface Type
- Drive 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 (partial view taken on 1Gb TP9100 enclosure)

The **Modify Disk Device Status** dialog box allows you to not only quickly view status of the drives, but also to define hot spares. Each drive in Figure 3-3 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	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.
Fibre Channel Interface Speed	The drive Fibre Channel interface speed in gigabits per second.
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	The spindle speed in RPM of the hard disk drive.

Table 3-2 Modify Disk Device Status Field Definitions (continued)

Field	Definition
Physical	The physical number of blocks and megabytes on the disk. 1 MB = 1024 * 1024 bytes, and one block = 512 bytes.
Usable	The usable number of blocks and megabytes on the disk. These numbers are always less than the physical numbers because the RAID controller allocates a portion of disk space for its configuration on disk (COD). Usable sizes vary, based on the version controller firmware and COD versions.
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.
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, OFFLINE, ONLINE, REBUILD, and EMPTY.

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 3-4 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 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>	<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 any LUN - Retains ALL Data (Except that of the LUN(s) to be destroyed).</p> <p>This will let you delete any LUN. If you click on the Delete LUN, then the next screen will present a list of LUNs to select from, and you will have the opportunity to cancel the procedure.</p> <p>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 device is sixteen. • You can not have 32 LUNs defined. The expansion require one free LUN in order to execute. • You can not perform an expansion if all 32 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, Initiate, 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 • The RAID level must not change, with exception that RAID-1 can be expanded into RAID 0+1 	<p>Expand LUN</p>

Figure 3-4 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 you have 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 3-5.

RAID Array Configuration for system at *hw/ucsl/isc13d010*

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																	0	0	0	0
<i>Stripe Size (All LUNs) is TBD. Totals:</i>																	0	0	0	0	0

RAID Subsystems:

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

17560MB 0-123 <input type="checkbox"/>	17560MB 1-124 <input type="checkbox"/>	17560MB 0-123 <input type="checkbox"/>	17560MB 1-122 <input type="checkbox"/>
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.

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

Figure 3-5 New Configuration Screen, 1Gb TP9100 12 drive enclosure

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 16 physical disk drives

to be configured as a system drive. A system drive can span up to 16 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 16.
- 1 to 16 disk drives can be combined into a pack.
- 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.

Note: Do not mix RAID 0 (non-redundant) and RAID 1,3,5, and 0+1 (redundant) LUNs within a pack (system drive).

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 16 packs.
- The minimum size of a system drive is 8 MB. The maximum size is 2 TB (see the note on the next page).
- Up to 32 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 MORE operations on the IRIX operating system. SGI does not support spanning of non-redundant (RAID 0) and redundant (RAID 1,3,5, and 0+1) LUNs together within a drive pack (system drive). Redundant (fault tolerant) RAID levels may be spanned. See “Expand LUN” on page 48 for more information on MORE operations.

Note: TPM will allow a LUN capacity to be created of 2 terabytes.

On SGI IRIX operating system levels prior to 6.5.15, the disk utility `fx` has limited support of up to 1 terabyte for a single LUN, or volume (.999 terabyte (2147483647 blocks [1099511627264 bytes])).

SGI IRIX levels 6.5.15 and later have the required `fx` and IRIX infrastructure changes (for example, `prtvtoc`, and so on) to support a 2 terabyte LUN (1.999 terabytes (4294967295 blocks [2199023255040 bytes])).

Please take this into account ahead of time and do not create and initialize system drives (LUNs) that will exceed the `fx` support level of your system.

Follow this procedure to create and define new LUNs. More details on how to do this are found starting in Figure 3-6 on page 35 of this manual.

Note: The RAID Subsystems enclosure view will differ depending on the TP9100 system that TPM is configuring.

TP9100 (1Gb TP9100) will contain checkboxes for 12 drives.
TP9100 (2Gb TP9100) will contain checkboxes for 16 drives.

The figures used in this manual are 1Gb TP9100 based.

1. Activate the desired checkboxes in the **RAID Subsystems** area of the screen shown in Figure 3-5 on page 31 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).
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 3-5 on page 31) 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 3-6.

RAID Array Configuration for system at /hw/ses1/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 2-123	17560MB 3-122													70242	68664	34332	51498	34332
B	Next Assign																				
Stripe Size (All LUNs) is 1GB. 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 2-123	17560MB 3-122
17560MB 0-121	17560MB 1-120	17560MB 2-119	17560MB 3-118
17560MB 0-117	17560MB 1-116	17560MB 2-115	17560MB 3-114

- 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 Drives Assigned into Pack A (12 drive, 1Gb TP9100 RAID Subsystem view)

Figure 3-6 shows 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 , 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 unassign 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 3-4 on page 30), 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 31 for rules on drive pack configuration.

When you are finished, the screen looks similar to the one pictured in Figure 3-7. The drives have been grouped into three separate drive packs.

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

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	17550MB 0-125	17550MB 1-124	17550MB 0-123	17550MB 1-122													70242	68664	34332	51408	34332
B	17550MB 0-121	17550MB 1-120	17550MB 0-119	17550MB 1-118													70242	68664	34332	51408	34332
C	17550MB 0-117	17550MB 1-116	17550MB 0-115	17550MB 1-114													70242	68664	34332	51408	34332
D	Next Assign																				
Stripe Size (All LUNs) is TBD. Totals: 210728 205992 102396 154434 102396																					

RAID Subsystems:			
Enclosure #0 (Rackmount view, rotate 90° clockwise if in tower)			
17550MB 0-125	17550MB 1-124	17550MB 0-123	17550MB 1-122
17550MB 0-121	17550MB 1-120	17550MB 0-119	17550MB 1-118
17550MB 0-117	17550MB 1-116	17550MB 0-115	17550MB 1-114

Clear All Select Previous Pack Define LUNs Select Marked Cancel

- 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 3-7 All Drive Packs Created (12 drive, 1Gb TP9100 RAID Subsystem view)

Defining a LUN Using the Drive Packs

To define a LUN, follow these steps:

1. Click **Define LUNs**.

The screen shown in Figure 3-8 appears.

Logical Unit Configuration for system at `/dev/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
<input type="checkbox"/> A	17560MB 0-125	17560MB 1-124	17560MB 0-123	17560MB 1-122												70242	68664	0 (0.0%)						
<input type="checkbox"/> B	17560MB 0-121	17560MB 1-120	17560MB 0-119	17560MB 1-118												70242	68664	0 (0.0%)						
<input type="checkbox"/> C	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 4,284,967,295 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>				

Figure 3-8 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 3-9 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	# 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

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

Figure 3-9 LUN is Being Defined

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

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

The screen shown in Figure 3-10 appears.

Logical Unit Configuration for system at `/hw/scsi/sc3a010`

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

Figure 3-10 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 3-11 shows how the screen appears after using all the drive packs to define three separate LUNS.

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	1-124	0	35155968	17166
																			0	0-123	0	35155968	17166
																			0	1-122	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
																			1	1-118	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
																			2	1-114	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 3-11 All LUNs Defined

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.

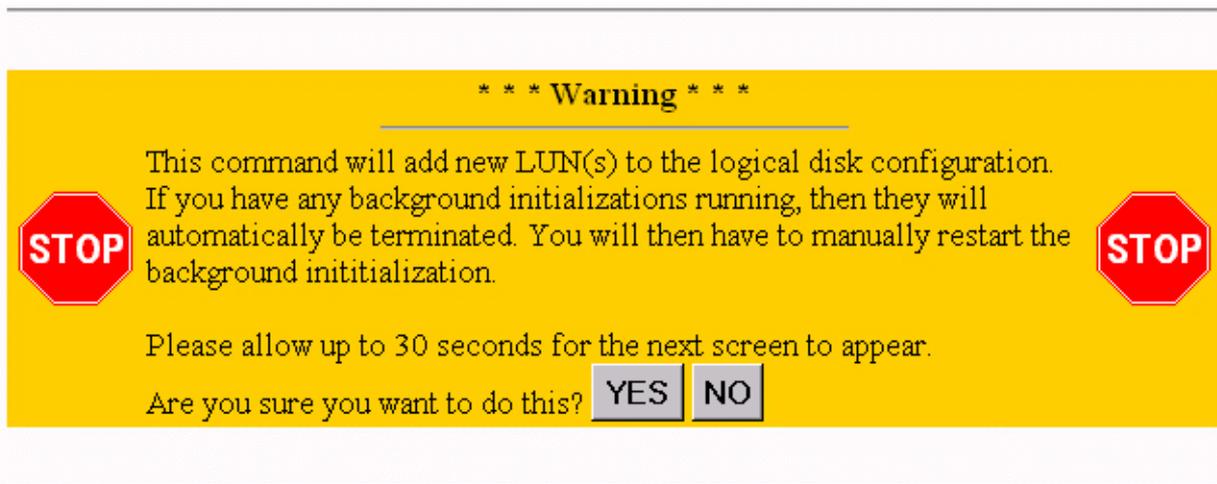


Figure 3-12 Updating the RAID Controller Warning Screen

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

If all goes well, the screen shown in Figure 3-13 appears indicating that the new LUN(s) are online.

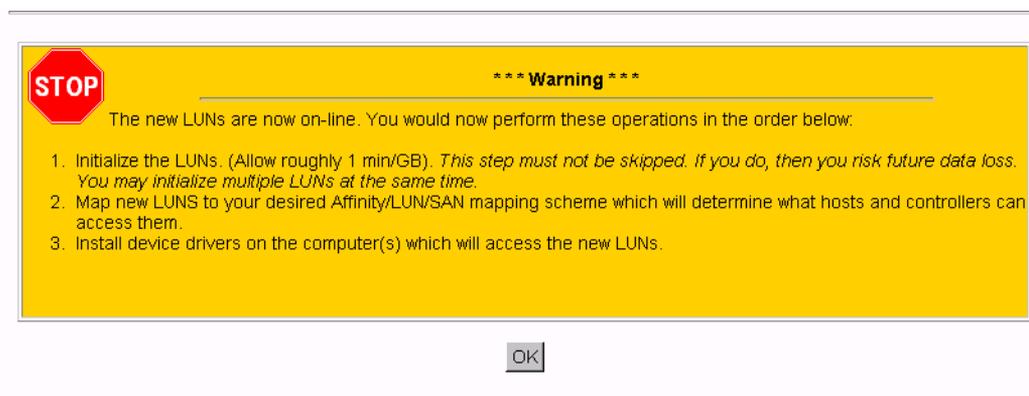


Figure 3-13 LUN Online Screen

3. After you click **OK** on the screen shown in Figure 3-13, 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 65).
 - b. Initialize the LUNs (see “Initialize LUN(s) (Configuration Function Menu)” on page 52).

Add LUN(s)

Back at the main menu, select **Create/Destroy/Expand Logical Drives** under the **Configuration Function** menu. The screen shown in Figure 3-14 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 type. Or choose 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>	<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 any LUN - Retains ALL Data (Except that of the LUN(s) to be destroyed).</p> <p>This will let you delete any LUN. If you click on the Delete LUN, then the next screen will present a list of LUNs to select from, and you will have the opportunity to cancel the procedure.</p> <p>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 32 LUNs defined. The expansion requires one free LUN in order to execute. • You can not perform an expansion if all 32 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, Initiate, 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 • The RAID level must not change, with exception that RAID-1 can be expanded into RAID 0+1 	<p>Expand LUN</p>

Figure 3-14 Create/Destroy/Expand Logical Drives Screen

To add one or more LUNs, follow these directions:

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

A screen similar to the one in Figure 3-15 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 0-118 <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 3-15 Add LUN(s) Screen (12 drive, 1Gb TP9100 RAID Subsystem view)

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 31.

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 3-14 on page 43) to perform an online RAID expansion procedure.

Note: SGI does not support LUN Expansion and MORE operations.

3. When you have finished defining the new LUN, click **Next Screen**.

The warning screen shown in Figure 3-16 appears.

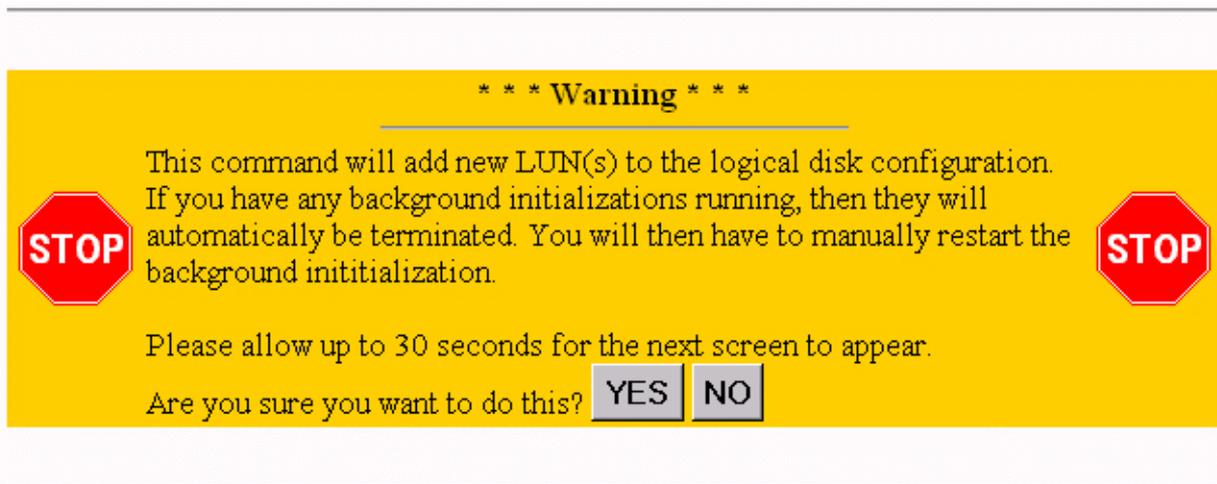


Figure 3-16 Add LUN(s) Warning Screen

4. Click **YES** to add the LUN, or **NO** to cancel and go back to the main menu.

If you select **YES**, the confirmation screen in Figure 3-17 appears after a brief waiting period.

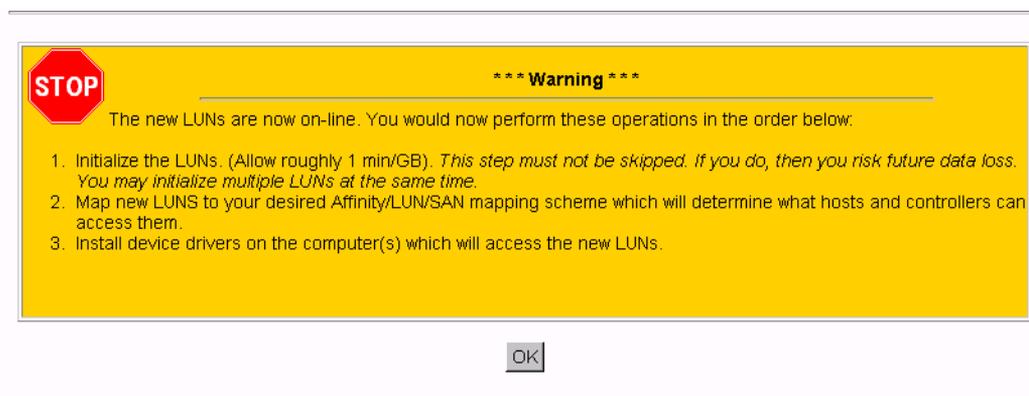


Figure 3-17 Add LUN(s) Confirmation Screen

5. After you click **OK** on the screen shown in Figure 3-17, you are returned back to the main page. You may then follow the steps shown in the screen of Figure 3-17.

Delete LUN

Use the **Delete LUN** button on the **Create/Destroy/Expand Logical Drives** screen (see Figure 3-14 on page 43) to delete a LUN.

Note: Any LUN can be deleted.

To delete a LUN, follow these directions:

Note: The following LUN deletion directions are valid for users using TPM 1.4 or later with 8.40 or later firmware only. Users with TPM 1.3 or firmware releases that do not support selective LUN deletion should refer to prior versions of this manual.

1. Select the LUN that you want to delete.

Logical Disk Information for RAID Subsystem at `/hw/scsi/2000050cc004b71/1un0/c6p1`

Logical Drive #0 (138848 MB
145,592,680,448 Usable Bytes) RAID 3 (Right Asymmetric) [Write Cache Enabled, Stripe=64KB]

Affinity Map: **SAN Mapping**

Background Tasks: No background jobs active

DELETE This LUN

Drive#	Span	Physical Enclosure	Physical Row	Physical Column	Channel	ID	Starting Block #	Span Size In blocks
0	0	0	0	0	1	4	0	142180352
1	0	0	1	0	1	8	0	142180352
2	0	0	2	0	1	12	0	142180352

Logical Drive #1 (69424 MB
72,796,340,224 Usable Bytes) RAID 1 (Mirror) [Write Cache Enabled, Stripe=64KB]

Affinity Map: **SAN Mapping**

Background Tasks: No background jobs active

DELETE This LUN

Drive#	Span	Physical Enclosure	Physical Row	Physical Column	Channel	ID	Starting Block #	Span Size In blocks
0	0	0	1	3	0	11	0	142180352
1	0	0	2	3	0	15	0	142180352

Logical Drive #2 (208272 MB
218,389,020,672 Usable Bytes) RAID 5 (Right Asymmetric) [Write Cache Enabled, Stripe=64KB]

Affinity Map: **SAN Mapping**

Background Tasks: No background jobs active

DELETE This LUN

Drive#	Span	Physical Enclosure	Physical Row	Physical Column	Channel	ID	Starting Block #	Span Size In blocks
0	0	0	0	1	0	5	0	142180352
1	0	0	1	1	0	9	0	142180352
2	0	0	2	1	0	13	0	142180352
3	0	0	2	2	1	14	0	142180352

CANCEL

Figure 3-18 Select LUN for Deletion

2. Click **Delete LUN**.

The warning screen shown in Figure 3-19 is displayed. This screen shows the size and characteristics of the last LUN that was created.

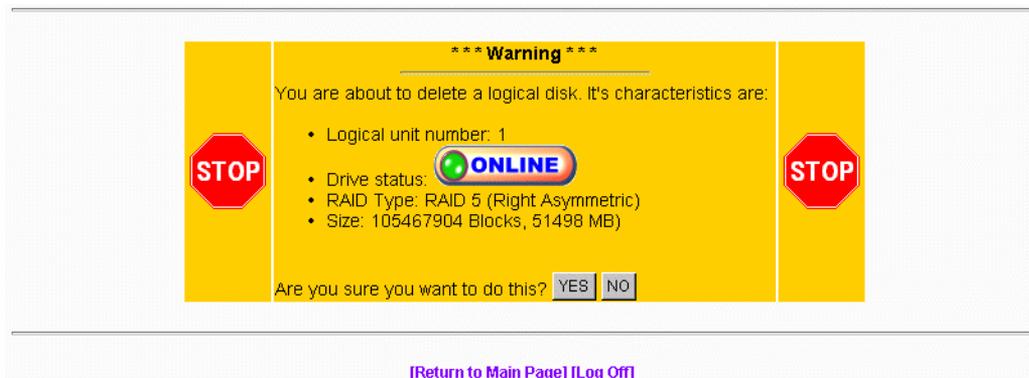


Figure 3-19 Delete LUN Warning Screen

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

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 or MORE operations.

- If you wish to attempt an online expansion, and your environment is not properly configured, the request will fail, and you will receive an appropriate error message.
- 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.
- 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.

- One to six drives can be added to a set at a time. The maximum number of physical disk drives is 16.
- You can not have 32 LUNs defined. The expansion requires one free LUN in order to execute.
- You cannot perform an expansion if all 32 LUNs are defined.
- The disk drives being added must be in STANDBY, and 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) the process automatically resumes 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.
- MORE, 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 online.

To expand a LUN, follow these steps:

1. Click **Expand LUN** in the **Create/Destroy/Expand Logical Drives** screen (see Figure 3-20).

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. <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 any LUN - Retains ALL Data (Except that of the LUN(s) to be destroyed).</p> <p>This will let you delete any LUN. If you click on the Delete LUN, then the next screen will present a list of LUNs to select from, and you will have the opportunity to cancel the procedure.</p> <p>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 32 LUNs defined. The expansion requires one free LUN in order to execute. • You can not perform an expansion if all 32 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 sub-system, (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, Initiate, 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 • The RAID level must not change, with exception that RAID-1 can be expanded into RAID 0+1 	<p>Expand LUN</p>

Figure 3-20 Create/Destroy/Expand Logical Drives Screen

A screen similar to that shown in Figure 3-21 appears.

RAID Array Configuration for system at */hw/scsi/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	17550MB 0-125	17550MB 1-124	17550MB 0-123	17550MB 1-122													70242	0	RAID 0 (Stripe)
B	17550MB 0-121	17550MB 1-120	17550MB 0-119	17550MB 1-118	17550MB 0-117	<input checked="" type="checkbox"/>	Next Assign										87803	1	RAID 0+1 (Mirrored Stripe)

RAID Subsystems:			
Enclosure #0 (Rackmount view, rotate 90° clockwise if in tower)			
17550MB 0-125	17550MB 1-124	17550MB 0-123	17550MB 1-122
17550MB 0-121	17550MB 1-120	17550MB 0-119	17550MB 1-118
17550MB 0-117	17550MB 1-116 <input type="checkbox"/>	17550MB 0-115 <input type="checkbox"/>	17550MB 1-114 <input type="checkbox"/>

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

Figure 3-21 Expand LUN Screen (12 drive, 1Gb TP9100 RAID Subsystem view)

Note: In the screen shown in Figure 3-21, 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 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.

3. 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 3-22 appears.

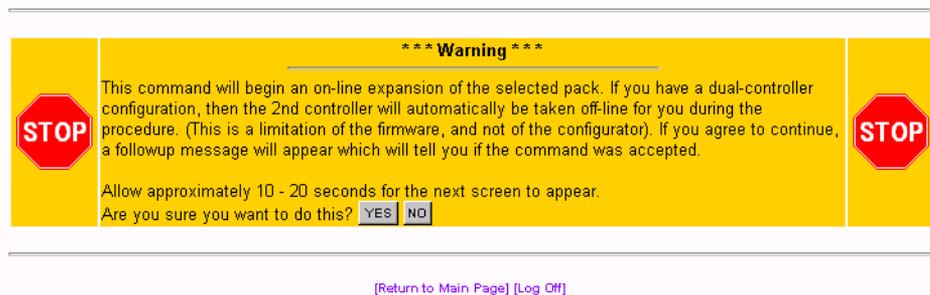


Figure 3-22 Warning Message

4. 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 LUNs can be initialized in the foreground or background. 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. To initialize LUNS in the foreground, click **Initialize LUNs** on the main menu.

The screen shown in Figure 3-23 appears.

To initialize LUNS in the background (instant LUN availability), click **Initialize LUNs in the background** from the main menu.

(Initialize LUN) 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	Initialize? <input type="text" value="NO"/>	UNINITIALIZED
Logical Drive #1 (35060 MB, 36,763,074,560 Bytes)	RAID 0+1 (Mirrored Stripe)	
Affinity Map: SAN Mapping	Initialize? <input type="text" value="NO"/>	UNINITIALIZED
Logical Drive #2 (35060 MB, 36,763,074,560 Bytes)	RAID 5 (Right Asymmetric)	
Affinity Map: SAN Mapping	Initialize? <input type="text" value="NO"/>	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 3-23 Initialize LUN

As shown in Figure 3-23, 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 3-24 appears.

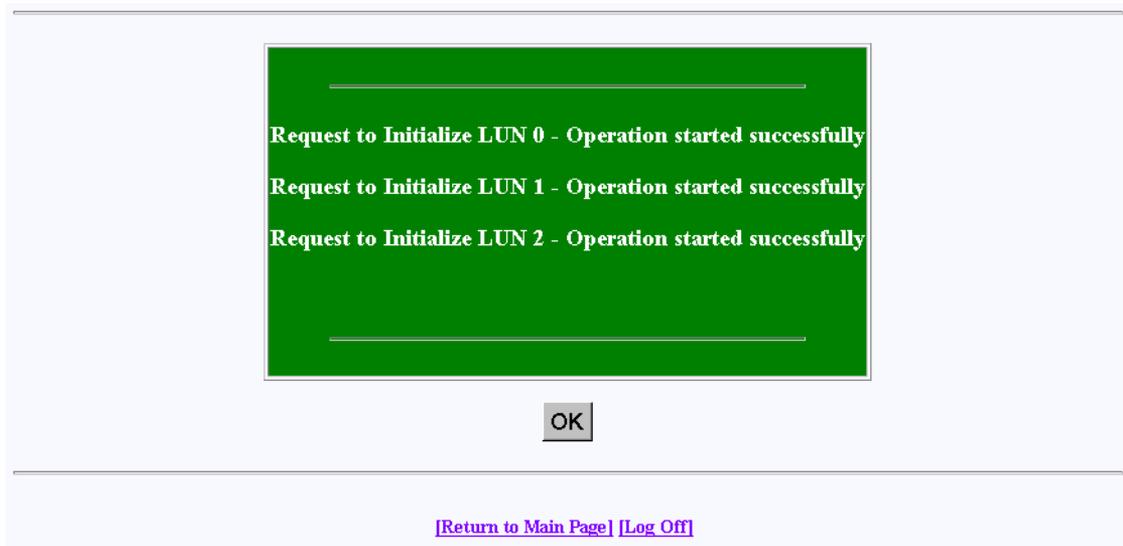


Figure 3-24 Initialization Confirmation

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 3-25 displays how the procedure is progressing, and updates every 10 seconds.

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 <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 3-25 Initialization Progress

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.

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

To initialize LUNs in the background, follow these steps:

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

The screen shown in Figure 3-26 appears.



Figure 3-26 Initialize LUNs in Background

2. To select an individual Logical Drive to be initialized, select the **Initialize** radio button next to the desired Logical Drive.

To select all Logical Drives created to be initialized, select the **Initialize All** radio button.

3. Once the Logical Drive or all Logical Drives are selected, click **INITIALIZE**.

The screen shown in Figure 3-27 appears.

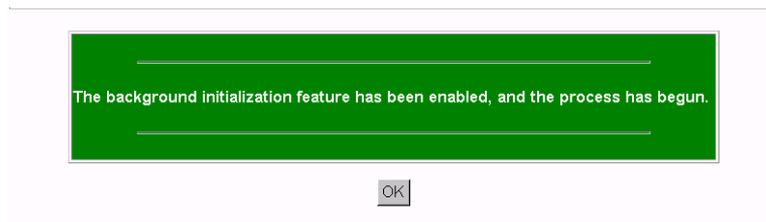


Figure 3-27 Initialize LUNs in Background

4. 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. To avoid this, let the background initialization process complete on all logical drives created, prior to putting the RAID array online.

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 20) 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 3-29 on page 63 appears.

The mapping strategy currently used by the controller is indicated with (CURRENT). In the example of Figure 3-29 on page 63, the SAN Map is the current mapping strategy.

Note: SAN MAP is the only mapping strategy supported by the TP9100 RAID system. On FFX2 RAID Controllers, TPM 1.4 will not by default set the LUN affinity to all (allow all hosts), as prior releases of TPM did.

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:

Note: A RAID controller 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 mappings for one port, then select **Apply**. After you select **OK** and go back to the main page, select **Configure Host to LUN Mapping** again and set mapping for the next port. This will allow you to postpone the controller reset until you have all mapping configured.

Note: 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.

Note: Only the most experienced administrators should make changes in an online 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.

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 3-28.

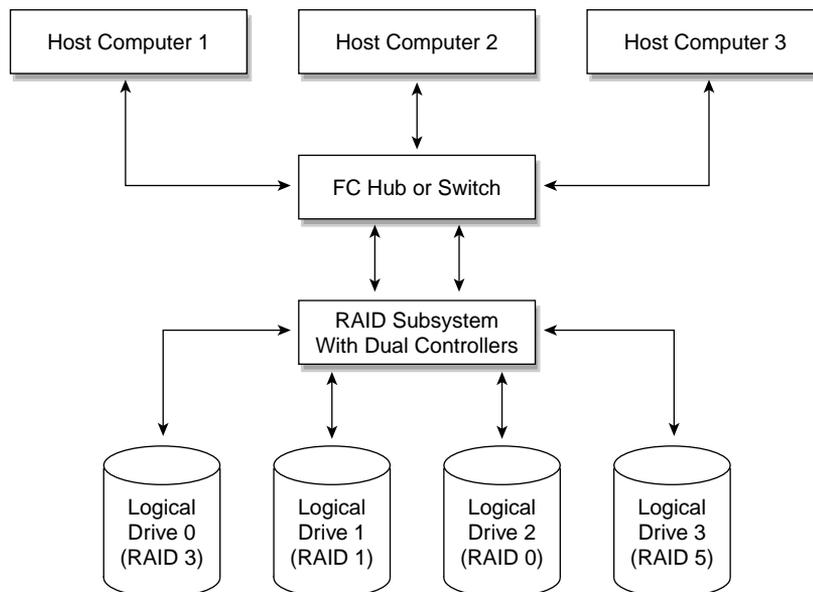


Figure 3-28 Storage Area Network

Without host to LUN mapping, each host computer (1 through 3) has complete access to all four system drives. When a host system boots, IRIX operating systems might not automatically attempt to mount all of these system drives, 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:

- 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.
- 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.
- 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.

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.
- **MultiPort.** 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.

Note: The SGI supported topology for multi-path failover is Multi-Port. Use TPM software to set the topology.

- 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.

Caution: If two systems independently access the same volume of data and the operating system does not support file locking, data corruption may occur. To avoid this, create two or more volumes (or LUNs) and configure each volume to be accessed by only one system.

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.

SAN Logical Drive Map for DAC960FFx2 RAID Subsystem
at /twe/scsi/20000050ce004b71/Lun0/c6p1

Topology Multi-TID	Topology Descriptions (Current Topology is Multi-TID)																																
Portmap Current: C0P0 C1P0 C0P0	<ul style="list-style-type: none"> Inactive Port - Ctrl0/Port0 and Ctrl1/Port0 are active. During failover, the inactive port on the partner takes over for the active port on the failed controller. Master/Slave - C0 active, C1 passive. When both controllers are operational, controller 0 (C0) is active on the loop, while C1 is inactive. If C0 fails, then C1 joins the loop(s) and resumes the I/O for C0. MultiPort - All ports are active. This topology does <i>not</i> provide transparent failover/failback and requires the hosts to handle controller failures through the use of DMP driver/HEA or similar host based path-failover technology. Multi-TID - All ports are active. This topology provides transparent failover/failback by allowing a failed controllers partner to impersonate the IDs of the failed controller as well as itself. Multi-TID can be usedwith (as an extra level of protection) or without host based path-failoversolutions. Hosts that track targets by Fabric ID (DID) must use a host based path-failover solution even with Multi-TID. <p style="color: red; font-weight: bold;">If you change the topology or LUN Number, then you must RESET the controller before changed are applied. However, a reset is not required to change any of the below mapping.</p> <p style="color: green; font-weight: bold; font-size: 1.2em;">On-screen changes will be lost unless you click the Apply button.</p>																																
<table border="1" style="width: 100%; border-collapse: collapse;"> <thead> <tr> <th colspan="2"></th> <th>SD #00</th> <th>SD #01</th> <th>SD #02</th> </tr> </thead> <tbody> <tr> <td colspan="2" style="text-align: center;">Logical Disks</td> <td></td> <td></td> <td></td> </tr> <tr> <td style="text-align: center;">SD #</td> <td style="text-align: center;">Type</td> <td style="text-align: center;">MB</td> <td></td> <td></td> </tr> <tr> <td style="text-align: center;">0</td> <td style="text-align: center;">RAID 3</td> <td style="text-align: center;">277,696</td> <td></td> <td></td> </tr> <tr> <td style="text-align: center;">1</td> <td style="text-align: center;">RAID 1</td> <td style="text-align: center;">69,424</td> <td></td> <td></td> </tr> <tr> <td style="text-align: center;">2</td> <td style="text-align: center;">RAID 5</td> <td style="text-align: center;">208,272</td> <td></td> <td></td> </tr> </tbody> </table>						SD #00	SD #01	SD #02	Logical Disks					SD #	Type	MB			0	RAID 3	277,696			1	RAID 1	69,424			2	RAID 5	208,272		
		SD #00	SD #01	SD #02																													
Logical Disks																																	
SD #	Type	MB																															
0	RAID 3	277,696																															
1	RAID 1	69,424																															
2	RAID 5	208,272																															
Host World Wide Name		SD #00	SD #01	SD #02																													
Drive to LUN Mapping		*LUN0	*LUN1	*LUN2																													
Allow ANY host access (currently defined hosts and any other host)		<input type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>																													
10-00-00-60-69-20-15-72		<input type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>																													
20-00-00-E0-6E-00-0E-ED		<input type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>																													
20-00-00-E0-8E-04-F0-C5		<input type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>																													

Note: each SD# above refers to a unique logical disk number (System Disk) number. This screen lets you assign LUN numbers to each SD, as well as the WWNs for each HEA which will be allowed to have access to them. Warning: Some operating systems such as IRIX and LINUX require that a LUN#0 is mapped to them or they won't probe for additional ones.

Figure 3-29 SAN Mapping Topology view - 1Gb TP9100 RAID Controller (Single Port)

Note: Reset button is not a controller reset button. It is used if you make a mistake and need to reset all select buttons.

Regardless of the topology, the bottom portion of the screen shown in Figure 3-29 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:

Note: A RAID controller 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 mappings for one port, then select **Apply**. After you select **OK** and go back to the main page, select **Configure Host to LUN Mapping** again and set mapping for the next port. This will allow you to postpone the controller reset until you have all mapping configured.

Note: 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.

Note: Only the most experienced administrators should make changes in an online environment.

Installing Device Drivers for New LUNS (Configuration Function Menu)

Selecting this item on the main menu (see “Main Menu” on page 19) 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 3-30 appears.

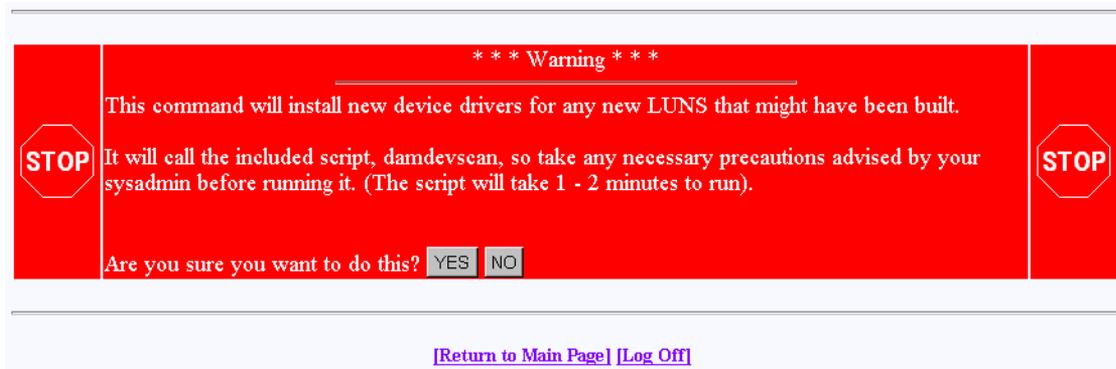


Figure 3-30 Install New Device Drivers Warning Screen

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:

- IRIX: TPM calls the `$DAM_HOME/damdevscan` shell script.

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

Administrative Functions

The **Administrative Functions** menu is located on the main screen and is shown in Figure 3-31.

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.

Figure 3-31 Administrative Functions Menu

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

- “View/Modify RAID Controller Configuration” on page 66
- “Reset Controller(s)” on page 70
- “Gracefully Bring a Controller Off Line” on page 72
- “Gracefully Bring a 2nd Controller On Line” on page 74
- “Perform Data Consistency Operations on a LUN” on page 75
- “Enable/Disable Write Cache for LUN(s)” on page 78
- “Modify/Purge Non-Volatile WWN Tables” on page 79

View/Modify RAID Controller Configuration

The **View/Modify RAID controller configuration** 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 3-32 through Figure 3-34. A partial dialog box is shown in each figure.

DAC960FFx2 Controller Configuration (Firmware Type 5 Rev 8.40 Build #0) 512 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 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 type="text" value="Auto negotiate"/>	Host Channel Data Rate	On-the-fly	This sets the data rate on all host channels. Once you save a change, then you may have to reboot the subsystem, host, and/or make configuration changes to your SAN before the subsystem will be recognized. This is a safe operation in the sense that it will not affect your data; it will change the interface speed, which can cause all logical disks and the subsystem to disappear until you force everything to renegotiate the interface speed. <i>Note: By design, this will set the speed settings on all host channels on controllers 0 and 1, even if you have only one controller attached. WARNING: Changing this setting from the factory default setting of "Auto negotiate" is not recommended. Do not change this parameter unless specifically instructed by your supplier.</i>
<input type="text" value="Auto negotiate"/>	Disk Channel Data Rate	On-the-fly	This sets the data rate on all disk channels. Once you save a change, then you may have to reboot the subsystem, host, and/or make configuration changes to your SAN before the subsystem will be recognized. This is a safe operation in the sense that it will not affect your data; it will only affect the interface speed, which will cause all logical disks and the subsystem to disappear until you force everything to renegotiate the interface speed. <i>Note: By design, this will set the speed settings on all disk channels on controllers 0 and 1, even if you have only one controller attached. WARNING: Changing this setting from the factory default setting of "Auto negotiate" is not recommended. Do not change this parameter unless specifically instructed by your supplier.</i>
<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 used for applications with a high degree of sequential access. WARNING: Do not enable this feature unless told to do so by your supplier!
<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="40"/>	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. <i>If you wish to monitor the rebuild status through the configurator, set this value no higher than 45.</i>

Figure 3-32 View/Modify Controller Configuration, 2Gb TP9100, 8.40 Firmware

<input type="text" value="2"/>	Disk Queue Limit	Reset	<p>Sets the maximum allowed queue depth for tagged commands to all attached disk drives. This value is further limited to the disk drives own tag limit, when that limit is reached. A setting of one is similar to no tags. If using device combing (Queueing - Coalescing Optimization), set the queue tag limit to 2.</p> <p>Do not change this value unless specifically directed to do so. Range is (1-255).</p>
<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 confused by QUEUE FULL.
<input checked="" type="checkbox"/>	Fallover 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 fallover 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 fallover, 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 checked="" 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).
<input type="text" value="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 makes duplex firmware work better in a simplex environment by disabling some internal operations. (A single controller running Simplex firmware will provide best performance). Note, if this parameter is enabled in a duplex environment, on resetting the controller, it will be disabled again. It is not possible to run a dual controller configuration in simplex mode.
<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.
<input type="text" value="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.
<input type="text" value="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.

Figure 3-33 View/Modify Controller Configuration, 2Gb TP9100 Dual Port, 8.40 Firmware

Note: Figure 3-34 shows a Dual Port Duplex RAID Controller Configuration.

<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"/> [Active] Ctrl0 Port1 <input type="text" value="1"/> [Active] Ctrl1 Port0 <input type="text" value="2"/> [Active] Ctrl1 Port1 <input type="text" value="3"/> [Active]	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. <i>Do not adjust these numbers unless instructed to do so by your vendor's technical support team.</i>
<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="12 Hours"/>	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="2"/>	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"/>	Background Init	On-the-fly	If enabled, background initialization will start automatically. When the controller restarts up or after a configuration change on any system drives that require initialization. If disabled background initialization will NOT start automatically when the controller restarts up or after a configuration change. Regardless of this parameter being enabled or disabled. Background initialization must be initiated to immediately use the LUN(s) for whatever I/O you wish to perform. Background initialization makes uninitialized system drives consistent by setting the parity while concurrently allowing the host to have read and write access to the system drive. <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.

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.**

Figure 3-34 View/Modify Controller Configuration, 2Gb TP9100 Dual Port, 8.40 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 3-35 appears.

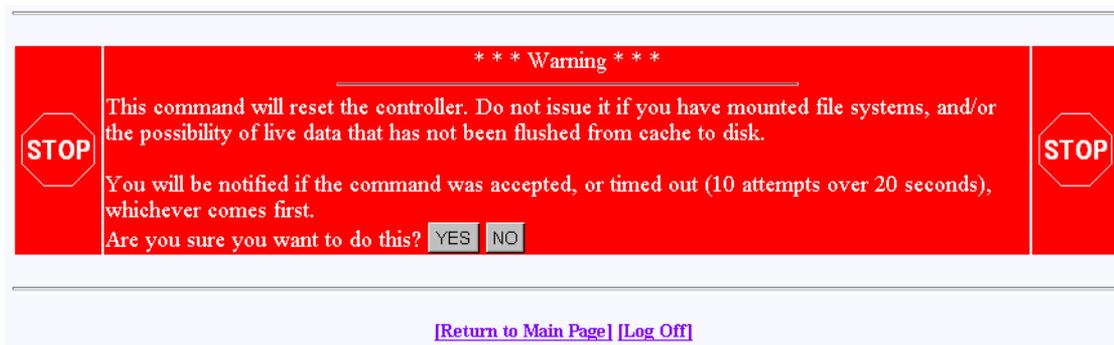


Figure 3-35 Reset Controllers Warning Screen

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.

You will see the screens shown in Figure 3-36 and Figure 3-37. 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 3-36 Controller Booting Window



Figure 3-37 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 3-38 appears.



Figure 3-38 Warning Screen

Click the **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 3-39 appears.

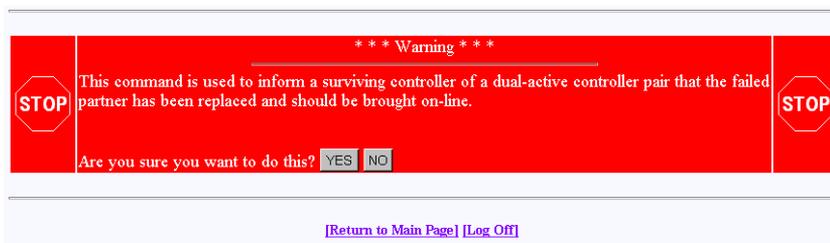


Figure 3-39 Warning Screen

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 3-40 appears.

Check and/or Restore Consistency of LUNs for RAID Subsystem at /hw/scsi/sc5d010

Logical Drive #0 (71618560 Blocks, 36,668,702,720 Bytes)	RAID 0+1 (Mirrored Stripe)	
Affinity Map: SAN Mapping	Check Only? <input type="radio"/>	Check and Restore? <input type="radio"/>
Logical Drive #1 (107427840 Blocks, 55,003,054,080 Bytes)	RAID 3 (Parity Stripe)	
Affinity Map: SAN Mapping	Check Only? <input type="radio"/>	Check and Restore? <input type="radio"/>
Logical Drive #2 (107427840 Blocks, 55,003,054,080 Bytes)	RAID 5 (Parity Stripe)	
Affinity Map: SAN Mapping	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 3-40 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 executing any of the check or restore actions, the dialog box of Figure 3-41 appears, indicating that the process has started.

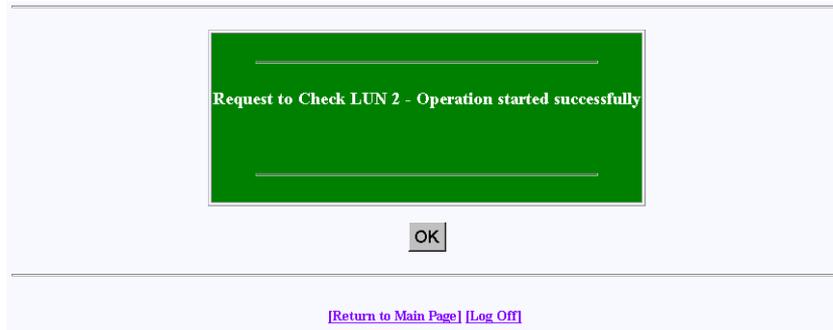


Figure 3-41 Operation Started Dialog Box

Click **OK** to proceed.

The window shown in Figure 3-42 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
 /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 3-42 Background Initialization Status Screen

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 3-43 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)	 ONLINE
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)	 ONLINE
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)	 ONLINE
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 3-43 Enable Write Cache For LUN(s) Dialog Box

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 3-44 appears.

World Wide Name Table for RAID controller at device: `/hw/scsi/sc3d110`

- This is a list of the host ports and what controller ports they are connected to.
- If a fibre channel controller (HBA) does not appear below, then the RAID 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.
- The alphanumeric descriptions are stored in the ASCII file, `wwntext.txt`, located in the home directory, which may be manually edited.
- 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	Description	Controller 0 Port 0	Controller 0 Port 1	Controller 1 Port 0	Controller 1 Port 1
20-00-00-E0-8B-05-AE-64	{Undefined}	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
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					

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

Figure 3-44 World Wide Name Table for RAID Controller

The World Wide Name (WWN) table includes a delete WWN from WWN table and a purge old entries option.

- 1Gb TP9100 RAID controllers support a maximum of 64 host WWNs.
- 2Gb TP9100 RAID controllers support a maximum of 256 host WWNs.

Note: All SAN mapping assignments created using firmware version 7.75 or later 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 3-45.

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.

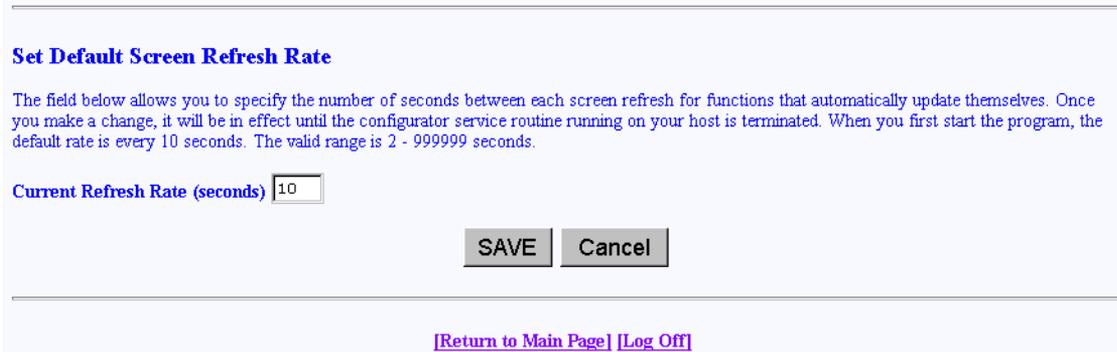
Figure 3-45 Reporting Functions Menu

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

- “Set Default Screen Refresh Rate” on page 81
- “Display (Dual) Controller Status” on page 81
- “Topology Query” on page 82
- “Display SCSI/Fibre Device Information” on page 83
- “Display FULL Subsystem Configuration Information” on page 85
- “Display Physical Subsystem Information” on page 87
- “Display Logical Subsystem Information” on page 89
- “Display Environmental Subsystem Information” on page 90

Set Default Screen Refresh Rate

When the **Set Default Screen Refresh Rate** menu selection is made, the screen shown in Figure 3-46 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)

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

Figure 3-46 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 3-47 appears.

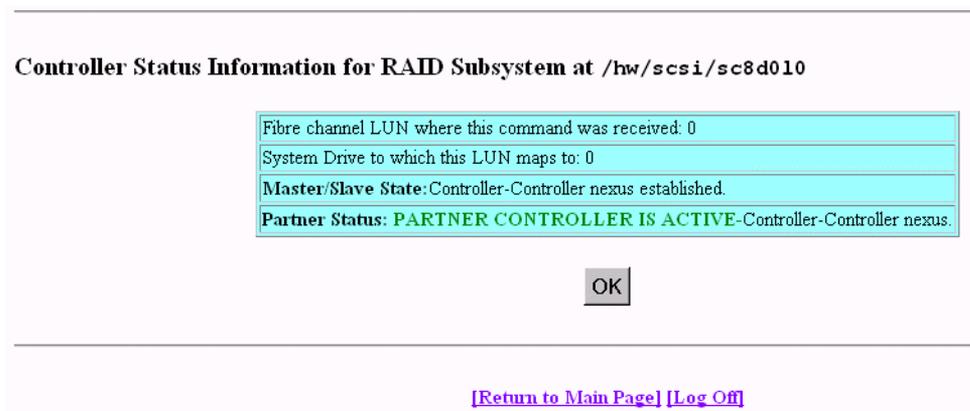


Figure 3-47 Controller Status Information

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 3-48 appears.

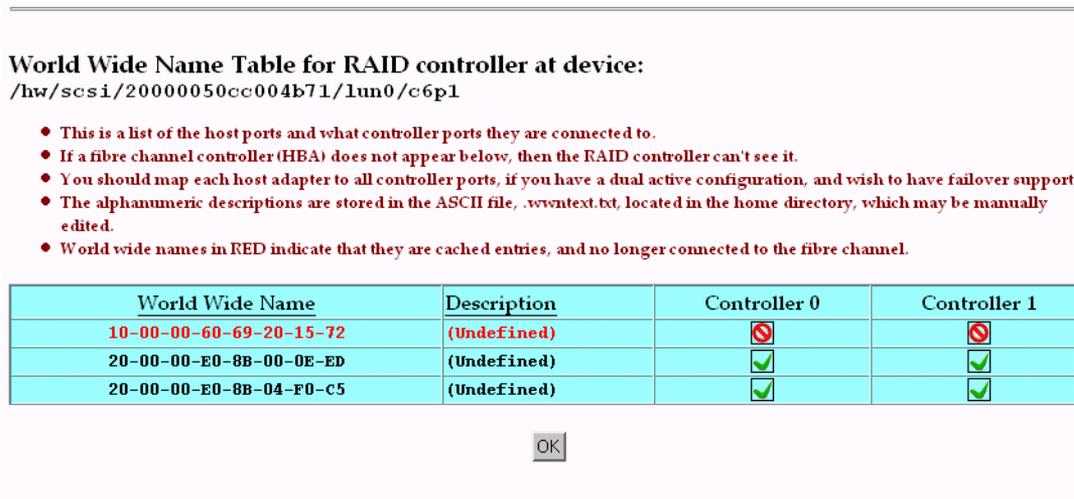


Figure 3-48 Topology Query (16 drive, 2Gb TP9100 RAID Subsystem View)

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 SCSI/Fibre Device Information

When the **Display SCSI/Fibre Device Information** menu selection is made, the screen shown in Figure 3-49 appears.

SCSI Inquiry Dump for Device at /hw/scsi/20000050cc004b71/Lun0/c6p1

Device Characteristics	Value
Vendor (Manufacturer)	SGI
Product Identifier	TP9100 PFX2
Microcode Revision Level	5840
Unit Serial Number	[See below]
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 checked="" type="checkbox"/> 32-Bit Transfers	<input checked="" type="checkbox"/> 32-Bit Addressing	<input checked="" type="checkbox"/> Terminate Task Management
<input checked="" type="checkbox"/> 16-Bit Transfers	<input checked="" type="checkbox"/> 16-Bit Addressing	<input checked="" type="checkbox"/> Normal ACA (NormACA)
<input checked="" type="checkbox"/> Synchronous Data Transfers	<input checked="" type="checkbox"/> Relative Addressing Mode	<input checked="" type="checkbox"/> Linked Commands on LUN
<input checked="" type="checkbox"/> Transfer Disable Messages Supported	<input checked="" type="checkbox"/> Enclosure Services Supported	<input checked="" type="checkbox"/> Dual-ported Device
<input checked="" type="checkbox"/> Handshake on Q Cable Supported	<input checked="" type="checkbox"/> Removable Media	<input checked="" type="checkbox"/> Tagged Command Queuing

* Note - This does NOT mean that SES is enabled, or even guarantee that the device can be accessed via SES software.

```

0000h: 00 00 02 12 33 00 10 02 53 47 49 20 20 20 20 20  ...3...SGI
0010h: 54 50 39 31 30 30 20 46 46 58 32 20 20 20 20 20  TP9100 PFX2
0020h: 35 38 34 30 00 00 8C 11 20 00 00 80 E5 12 8C 11  5840...
0030h: 00 00 00 00 00 00 00 00 - - - - - - - - - -
    
```

ANSI Serial number page (80h)

```

0000h: 00 80 00 28 30 30 30 30 38 63 31 31 32 30 30 30  ... (00008c112000
0010h: 30 30 38 30 65 35 31 32 38 63 31 31 30 30 30 30  0080e5128c110000
0020h: 30 30 30 30 30 30 30 30 30 30 30 - - - - 000000000000
    
```

ANSI Device Identification Page (83h)

```

0000h: 00 83 00 3C 01 01 00 1C 53 47 49 20 20 20 20 20  ...<...SGI
0010h: 00 00 8C 11 20 00 00 80 E5 12 8C 11 00 00 00 00  .....
0020h: 00 00 00 00 01 02 00 08 00 80 E5 00 00 12 8C 11  .....
0030h: 01 14 00 04 01 00 00 00 01 10 00 04 00 00 00 00  .....
    
```

OK

Figure 3-49 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 3-50), Controller Tunable Parameters (see Figure 3-51), and GroupConfig Structure (see Figure 3-52) 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 00 ..k.d.....
0010h: 52 41 49 44 42 52 49 43 4B 00 00 00 00 00 00 00 RAIDBRICK.....
0020h: 4D 59 4C 45 58 20 44 41 43 46 46 78 20 20 20 20 MYLEX DACFFx
0030h: 00 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 00 .F.....
0050h: 00 00 00 00 00 00 00 00 7E 10 F8 7F 20 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 .....
0080h: 00 00 00 00 4D 59 4C 45 58 20 20 20 20 20 20 20 ...MYLEX
0090h: 20 20 20 20 01 00 00 00 00 00 00 00 20 00 80 00 .....
00A0h: 02 00 00 00 00 00 0D 00 0C 00 00 00 00 00 F4 00 .....
00B0h: 02 00 02 00 7E 7E 00 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 00 .....
00D0h: 20 00 80 00 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 00 .....
```

Figure 3-50 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 45 .....MYLE
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 3-51 Controller Tunable Parameters

```

GroupConfig Structure: [D:\ACI\CTL_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 3-52 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 3-53 and Figure 3-54 appear.

Physical Disk Information for RAID Subsystem at /hw/scsi/sc13d010 <i>(See legend below)</i>			
Enclosure #0 (Rackmount view, rotate 90° clockwise (f'm tower))			
SGI ST318304FC [2705] S/N: 3ELO098Q WVN: 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 WVN: 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: 3ELO0W4T WVN: 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: 3ELO0W1C WVN: 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: 3ELO0HNG WVN: 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: 3ELO0KDQ WVN: 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: 3ELO0W88 WVN: 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: 3ELO0XDT WVN: 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: 3ELO0VAD	SGI ST318304FC [2705] S/N: 3ELO06JW	SGI ST318304FC [2705] S/N: 3ELO0769	SGI ST318304FC [2705] S/N: 3ELO0W3B

Figure 3-53 Physical Subsystem Information (TP9100 1G, 12 bay enclosure view)

SGI ST318304FC [2705] S/N: 3E000VAD 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: 3E0008JW 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: 3E000769 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: 3E0003B 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 3-54 Physical Subsystem Information (TP9100 1G, 12 bay enclosure view)

These screens show 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 3-55 appear.

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 3-55 Logical Subsystem Information

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 Figure 3-56 on page 90 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 was not 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 3-56 appears.

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: SGI • Product ID: TP6100 • 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	Over/Temp 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 3-56 Display Environmental Subsystem Information Window

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 3-57.

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.
- [Maintain Security Features](#) - Allows you to maintain program access user accounts and options.
- [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 3-57 Miscellaneous Functions Menu

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

- “Display Status of Background Jobs” on page 92
- “Flush Controller(s) Write Cache to Disk” on page 94
- “Save Current Controller Configuration” on page 95
- “Load Controller Configuration” on page 97
- “Flash New Firmware Onto Controller(s)” on page 98

- “Flash New Firmware Onto Supported Disk(s)” on page 100
- “Adjust Battery Settings” on page 102
- “Maintain Security Features” on 9
- “Clear Configuration” on page 104
- “Set The Real Time Clock” on page 105
- “View the Controller’s Internal Event Log” on page 106
- “Identify a Disk” on page 107
- “Scan for New Enclosures and Disks” on page 108
- “Save Controller Debug Dump” on page 110
- “Stop Configurator Service Routine on Host” on page 111
- “Automatic Restart (Reboot) on Failure Parameter” on page 111

Display Status of Background Jobs

When the **Display Status of Background Jobs** menu selection is made, the windows appear as shown in Figure 3-58.

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 3-58 Display Status of Background Jobs—Background Init in Progress

These windows show you how initialization, rebuilds, or checks are progressing. The browser title bar (not shown in Figure 3-58 on page 93) 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 3-59 appears.

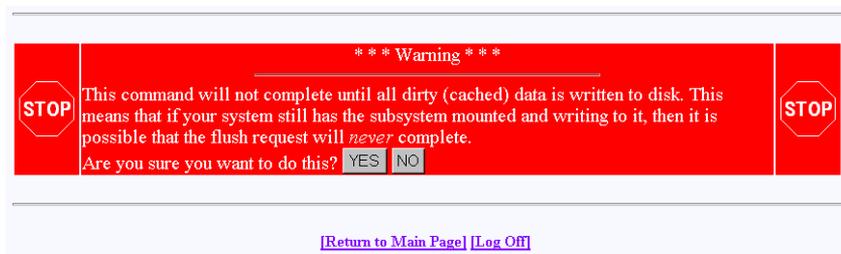


Figure 3-59 Cache Flush Warning

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 3-59 appears.

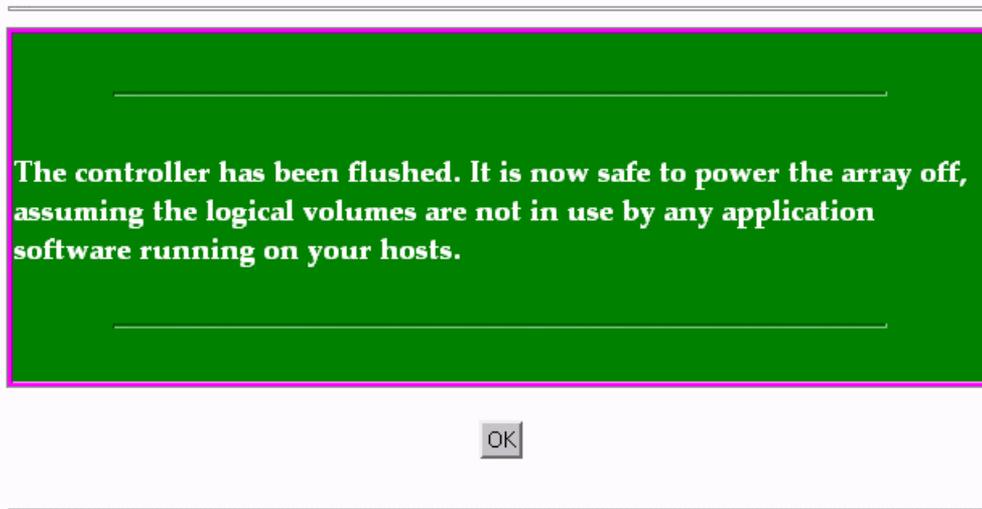


Figure 3-60 Cache Flush Successful

Save Current Controller Configuration

When the **Save Current Controller Configuration** menu selection is made, the screen shown in Figure 3-61 appears.

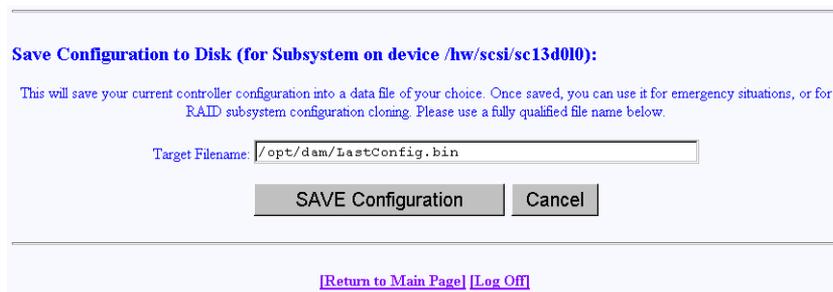


Figure 3-61 Save Current Controller Configuration

Note: It is highly recommended that users save the configuration at the following times:

After initial install, once all system drives are created.

After any LUN additions.

Prior to and after any LUN deletions.

Prior to any major service operations.

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.

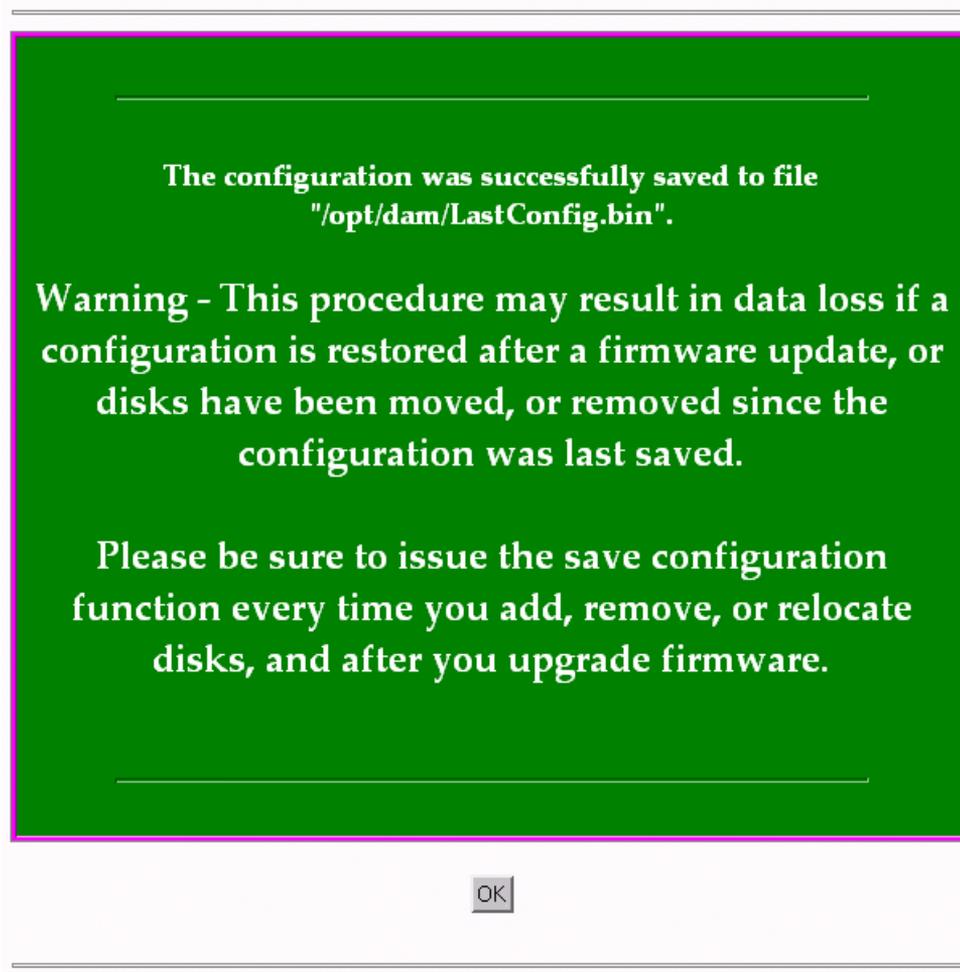


Figure 3-62 Configuration Save Successfull

Load Controller Configuration

When the **Load Controller Configuration** menu selection is made, the screen shown in Figure 3-63 appears.

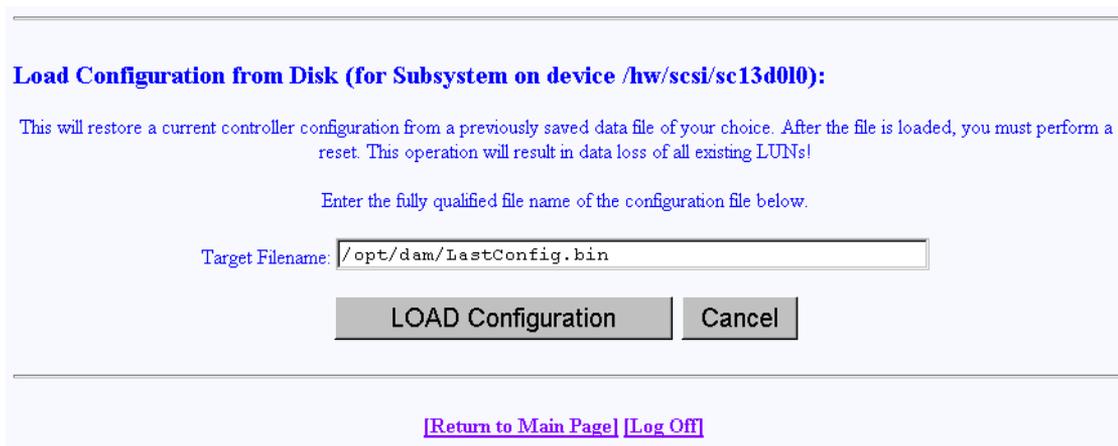


Figure 3-63 Load Controller Configuration

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 loaded from the server 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 3-64 appears.

Flash New Firmware onto Controller(s) (for Subsystem on device /hw/scsi/20000080e5115f52/lun0/c3p1):

This will upgrade/downgrade firmware into all controllers within this subsystem. Please perform the necessary safety precautions before proceeding.

1. Back up all data files.
2. **Make sure no background operations such as consistency checks are running.**
3. Unmount all logical disks from host computers.
4. Flush controller's cache.
5. Copy the firmware to this host computer, not the one attached to the browser.

Firmware Filename:

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

Figure 3-64 Flashing New Firmware

Caution: 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.

Note: The RAID controller firmware to be flashed must be placed in the /opt/dam directory.

Save the current configuration using TPM, prior to any flashing of controller firmware. See "Save Current Controller Configuration" on page 95 for more information.

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 3-65 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 3-32 on page 67). 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 3-66 appears.

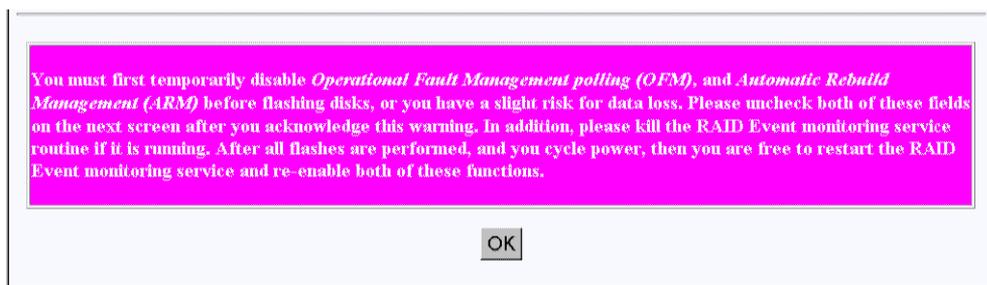


Figure 3-65 Flash New Firmware onto Supported Disk(s) Dialog Box

Caution: 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.

Caution: 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.

Note: The disk drive firmware to be flashed must be placed in the `/opt/dam` directory.

When you click **OK**, the **View/Modify RAID Controller Configuration** dialog box appears (see Figure 3-32 on page 67). 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 3-35 on page 70). After the reset is complete, click the **Flash New Firmware Onto Supported Disk(s)** menu selection.

The dialog box shown in Figure 3-66 appears.

Disk Drive Flashing for RAID Subsystem at `/hw/scsi/sc3d010`

Firmware Filename: `/opt/dam/DiskFw.img`

Select	Channel	ID Dec / Hex	Make	Model	Original Firmware Revision	Serial #	# of Blocks	Start Time	Complete Time
<input type="checkbox"/>	0	5.05h	SGI	ST373405FC	2702	3EK116NF	143374741		
<input type="checkbox"/>	0	7.07h	SGI	ST373405FC	2700	3EK16HRH	143374741		
<input type="checkbox"/>	0	9.09h	SGI	ST373405FC	2702	3EK103VF	143374741		
<input type="checkbox"/>	0	11.0Bh	SGI	ST373405FC	2702	3EK116FK	143374741		
<input type="checkbox"/>	0	13.0Dh	SGI	ST373405FC	2702	3EK119FJ	143374741		
<input type="checkbox"/>	0	15.0Fh	SGI	ST373405FC	2702	3EK0L84B	143374741		
<input type="checkbox"/>	0	17.11h	SGI	ST373405FC	2702	3EK119F8	143374741		
<input type="checkbox"/>	0	19.13h	IBM	DNEF-318350	F90F	AKOKA623	35843670		
<input type="checkbox"/>	1	4.04h	IBM	DNEF-318350	FYG4	AKDL0049	35843670		
<input type="checkbox"/>	1	6.06h	SGI	ST373405FC	2702	3EK10T1G	143374741		
<input type="checkbox"/>	1	8.08h	SGI	ST373405FC	2702	3EK117KA	143374741		
<input type="checkbox"/>	1	10.0Ah	SGI	ST373405FC	2702	3EK119FL	143374741		
<input type="checkbox"/>	1	12.0Ch	SGI	ST373405FC	2702	3EK10XE2	143374741		
<input type="checkbox"/>	1	14.0Eh	SGI	ST373405FC	2702	3EK113W3	143374741		
<input type="checkbox"/>	1	16.10h	SGI	ST373405FC	2702	3EK107B3	143374741		
<input type="checkbox"/>	1	18.12h	SGI	ST373405FC	2702	3EK10R74	143374741		

Comments:

1. Select a device, then press the FLASH button to initiate the firmware flash.
2. Please dismount LUNs on this array from your hosts before continuing.
There is extreme risk of data loss if you do NOT dismount the LUNs from your operating system before continuing.
3. This function is designed for CE's only, and should never be run unless specifically instructed to do so.
4. Each drive will typically flash within a few minutes, depending on the model. Larger drives can sometimes take over 5 minutes.
5. Do not run any I/Os through the subsystem while a flash is being performed.
6. Be sure to flush disk cache before performing this operation.
7. **Make sure no background operations such as consistency checks are running.**
8. You must recycle power to the array after drive(s) are flashed.
9. This feature works on supported disk drives only.

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

Figure 3-66 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 3-67 appears.

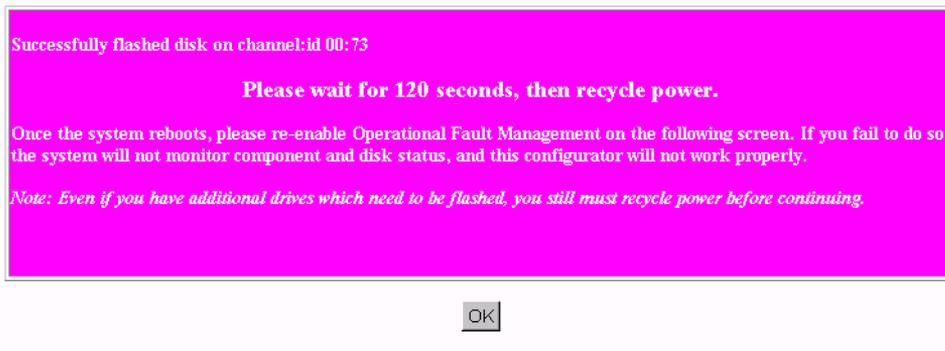


Figure 3-67 Flash New Firmware Complete Box

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 3-68 appears.

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 3-68 Adjust Battery Settings Window

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 3-56 on page 90).

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

- **SAVE Threshold** Saves the battery threshold value entered in the textbox.

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** 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 or later 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 3-69 appears.

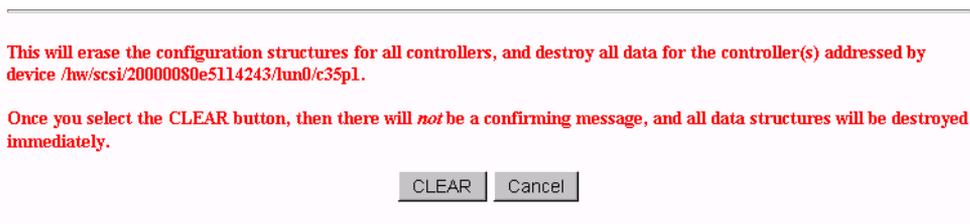


Figure 3-69 Clear Configuration Screen

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 3-70 appears.

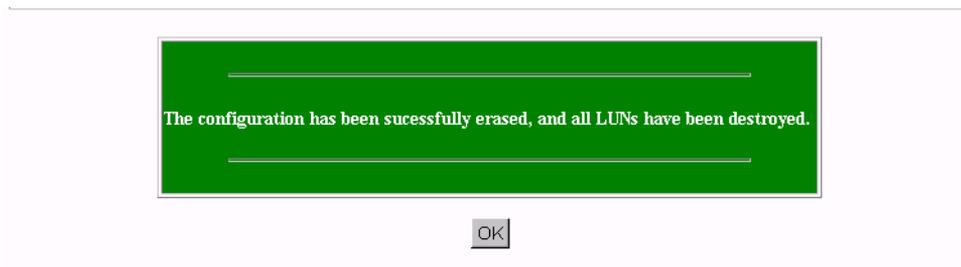


Figure 3-70 Clear Configuration Message

Set The Real Time Clock

When the **Set Real Time Clock** menu selection is made, the dialog box shown in Figure 3-71 appears.

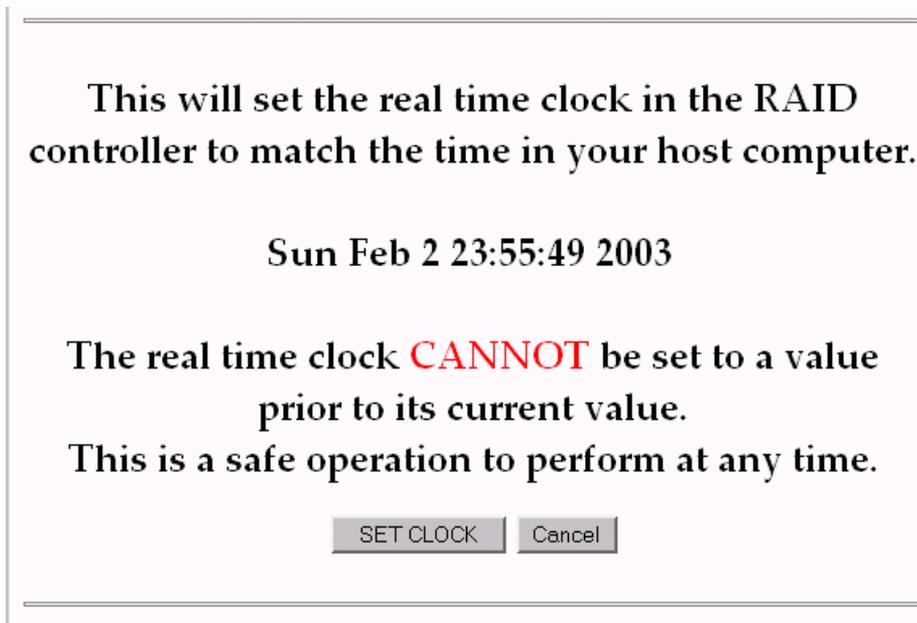


Figure 3-71 Set Real Time Clock Dialog Box

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

The confirmation screen shown in Figure 3-72 appears.



Figure 3-72 Set Real Time Clock Confirmation Dialog Box.

View the Controller's Internal Event Log

When the **View the Controller's Internal Event Log** menu selection is made, the dialog box shown in Figure 3-73 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>

Save All Save New Append All Append New View All Events Cancel

Figure 3-73 View Controller's Internal Event Log Dialog Box

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.

Identify a Disk

For the **Identify a Disk** menu selection to work properly, Operational Fault Management (OFM) must be enabled. Also, at least one of the SES disk drives must be present and operational. When the **Identify a Disk** menu selection is made, the dialog box shown in Figure 3-74 appears.

Physical Devices in RAID Subsystem at `/hw/scsi/sc3d010`

Select	Channel Decimal/Hex	ID Decimal/Hex	Make	Model	Enclosure	Row	Col
<input type="radio"/>	0 / 00h	5 / 05h	SGI	ST373405FC	0	0	1
<input type="radio"/>	0 / 00h	7 / 07h	SGI	ST373405FC	0	0	3
<input type="radio"/>	0 / 00h	9 / 09h	SGI	ST373405FC	0	1	1
<input type="radio"/>	0 / 00h	11 / 0Bh	SGI	ST373405FC	0	1	3
<input type="radio"/>	0 / 00h	13 / 0Dh	SGI	ST373405FC	0	2	1
<input type="radio"/>	0 / 00h	15 / 0Fh	SGI	ST373405FC	0	2	3
<input type="radio"/>	0 / 00h	17 / 11h	SGI	ST373405FC	0	3	1
<input type="radio"/>	0 / 00h	19 / 13h	IBM	DNEF-318350	0	3	3
<input type="radio"/>	1 / 01h	4 / 04h	IBM	DNEF-318350	0	0	0
<input type="radio"/>	1 / 01h	6 / 06h	SGI	ST373405FC	0	0	2
<input checked="" type="radio"/>	1 / 01h	8 / 08h	SGI	ST373405FC	0	1	0
<input type="radio"/>	1 / 01h	10 / 0Ah	SGI	ST373405FC	0	1	2
<input type="radio"/>	1 / 01h	12 / 0Ch	SGI	ST373405FC	0	2	0
<input type="radio"/>	1 / 01h	14 / 0Eh	SGI	ST373405FC	0	2	2
<input type="radio"/>	1 / 01h	16 / 10h	SGI	ST373405FC	0	3	0
<input type="radio"/>	1 / 01h	18 / 12h	SGI	ST373405FC	0	3	2
Total Drives: 16							

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.

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

Figure 3-74 Identify a Disk Dialog Box

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, the dialog box shown in Figure 3-75 appears.

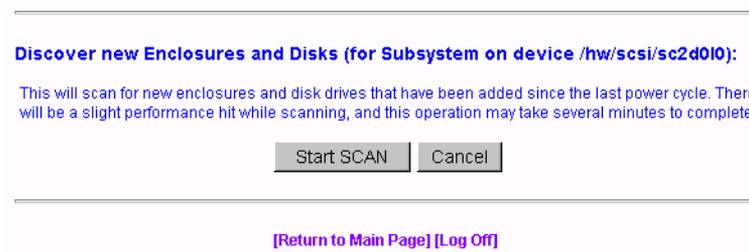


Figure 3-75 View Controller's Internal Event Log Dialog Box

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 3-75 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, the dialog box shown in Figure 3-75 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 3-76 Dump Debug Configuration to Disk Dialog Box

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

When the **Stop Configurator Service Routine on Host** menu selection is made, an emergency shutdown routine is initiated for TPM. 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 is shut down, the message shown in Figure 3-77 appears.

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

Figure 3-77 TPM 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 3-34 on page 69).
- **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 3-34 on page 69).

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, press 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 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] [-C Freq] [-d] [-e] [-i] [-s] [-w] [-c] [-z] [-H
HistFile] [-T LogFile] [-S StatusFile] -D DeviceFile [-M Addresses] [-L] (or)
TPMWATCH -Z -D Device File
```

Where:

-P *Freq*: Frequency in seconds between each poll

-C *Freq*: Frequency in hours to autosave configuration

- d: Enable debug mode to stderr
- z: Set the controllers clock and exit (requires -D option)
- z: Do NOT set the real time clock every polling interval
- e: Log and optionally e-mail error messages (requires 7.x FW or later)
- i: Log and optionally e-mail informational messages (requires 7.x FW or later)
- s: Log and optionally e-mail severe messages (requires 7.x FW or later)
- w: Log and optionally e-mail warning messages (requires 7.x FW or later)
- c: Log and optionally e-mail critical messages (requires 7.x FW or later)
- H *HistFile*: Optional event history file (obsolete in 7.x FW and later)
- T *LogFile*: Sends error text to user-specified log file
- S *StatusFile*: Status file (obsolete in 7.x FW and later)
- 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 does not 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 &`
- v: Display version number and exit

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.

Event and Error Codes

Event and error codes that are generated by the controller modules are displayed in the controller event log and can be viewed using TPM. Information entries identify an event or error that has occurred. Error codes help to identify the cause of a problem, the failing modules, and the service actions that might be needed to resolve the problem.

Event and Error Code Table

The event and error code table headings are defined as follows:

- **Error-Event Number**

A number that identifies the event or error. This number is displayed in the TPM Array Manager Event Log.

- **Type**

C – Critical. Controller failure.

S – Serious. The failure of a major component within the array enclosure.

E – Error Messages, such as a consistency check fails, or a rebuild on a physical disk stopped because of errors.

W – Warning Messages, such as a physical disk, or failed to start a rebuild.

I – Information Messages, such as system startups. When a consistency check or a rebuild has finished, for example, or a physical disk has been put on standby.

- **Description**

A brief description of the event or error.

- **Additional Details**

Additional details on the event or error.

- **Action**

If this item is an event that is displayed for your information only, no action is required. If this item is an error, this column defines the action you must take to resolve the problem.

Table B-1 Error-Event Codes

Error-Event #	Type	Description	Additional Details	User Action Required
1	I	A hard disk has been placed online.	Rebuild completed. Device was configured. Manual online was done.	None
2	I	A hard disk added as hot spare.	Device was configured. Manual hot spare was done. Automatic hot spare was done. 'Raidbld' made it hot spare.	None
3	W	Hard disk error found.	A bad sector was found on the physical media. Mechanical failure on the device. Host SCSI device detected illegal instruction. Target device generated unknown phase sequence.	If problem occurs frequently, replace the device or contact service representative.
4	W	Hard disk PFA condition found, this disk may fail soon.	Physical device predicted some future failure. External RAID logical device may have become critical.	Follow device vendor's defined action.
5	I	An automatic rebuild has started.	A physical device failed and spare was available. A physical device failed and no spare was available. A spare was added.	None
6	I	A rebuild has started.	Client started the rebuild on user's request. User replaced the failed device and 'Raidbld' started the rebuild.	None
7	I	Rebuild is over.	Rebuild completed successfully.	None
8	W	Rebuild is cancelled.	User cancelled the rebuild. Higher priority rebuild started.	Restart the rebuild if required.
9	E	Rebuild stopped with error.	Due to some unknown error on the controller, rebuild failed.	Try rebuild again.

Table B-1 Error-Event Codes (**continued**)

Error-Event #	Type	Description	Additional Details	User Action Required
10	E	Rebuild stoped with error. New device failed.	New physical device failed. New physical device may not be compatible with MDCA hardware/firmware.	Replace the physical device.
11	E	Rebuild stopped because logical drive failed.	At least one more physical device failed in the array. Bad data table overflow.	It may not be possible to recover from this error. Contact your service representative.
12	S	A hard disk has failed.	A physical device failed. A user action caused the physical device to fail.	Replace the physical device.
13	I	A new hard disk has been found.	A physical device has been powered on. A new physical device has been added. Controller was powered on. Controller was added. System has rebooted.	None
14	I	A hard disk has been removed.	User removed an unconfigured physical device. An unconfigured physical device failed. A controller was removed. A controller powered off.	Replace the device if needed.
15	I	A previously configured disk is now available.	User set the physical device to unconfigured.	None
16	I	Expand Capacity started.	User started the RAID Expansion operation. A suspended RAID Expansion operation was started.	None
17	I	Expand Capacity completed.	RAID Expansion finished.	None
18	E	Expand Capacity stopped with error.	Multiple physical devices failed.	It may not be possible to recover from this error. Contact your service representative.
19	W	SCSI command time out on hard device.	Physical device has been removed. Physical device failed. Command timeout value is not correct.	None

Table B-1 Error-Event Codes (**continued**)

Error-Event #	Type	Description	Additional Details	User Action Required
20	C	SCSI command abort on hard disk.	User may have requested to abort the command. Firmware may have aborted the command to recover from error. The device may have aborted the command.	None
21	W	SCSI command retried on hard disk.	The command may have timed out. Bus reset may have occurred. Device reset may have occurred.	None
22	W	Parity error found.	A physical device did not generate proper parity. The controller failed, did not check parity properly. Cable failed. Improper cable length. Another physical device interfered. Some outside environment affected the data on the cable (for example, radio frequency signal). Terminator is not connected. Improper termination.	It may not be possible to recover from this error. Contact your service representative.
23	W	Soft error found.	An error was detected by physical device and data was recovered.	Run consistency check. If problem occurs frequently, replace the physical device.
24	W	Misc error found.	A physical device reported some error, which does not fit in any category. Read/Write command timeout. Data overrun. Physical device was busy when host attempted to send command.	If problem occurs frequently, replace the physical device.
25	I	SCSI device reset.	Firmware has done reset to recover from error. User has done a reset.	None
26	I	Active spare found.	Device was configured. Manual active spare was done. Automatic active spare was done.	None
27	I	Warm spare found.	Device was configured. Manual warm spare was done. Automatic warm spare was done.	None
28	E	Request sense data available.	A physical device reported an error. Firmware reported an operational error.	Read the request sense data to understand the root cause.

Table B-1 Error-Event Codes (**continued**)

Error-Event #	Type	Description	Additional Details	User Action Required
29	I	Initialization started.	Host started the initialization.	Wait until the initialization is completed. If the system is shutdown prior to this process being completed, the physical device can be made useful only by reinitializing it.
30	I	Initialization completed.	Physical device initialization completed successfully.	None
31	W	Initialization failed.	Physical device could have problems supporting the SCSI format command.	Try to initialize again. Contact your service representative.
32	S	Initialization cancelled.	User cancelled the operation.	Hard disk must be initialized again or the hard disk cannot be used.
33	S	A hard disk failed because write recovery failed.	Write recovery process failed.	Replace hard disk and rebuild it.
34	S	A hard disk failed because SCSI bus reset failed.	SCSI bus reset failed.	Replace hard disk and rebuild it.
35	S	A hard disk failed because double check condition occurred.	Double check condition occurred.	Replace hard disk and rebuild it.
36	S	A hard disk failed because device is missing.	Access to the hard disk failed.	Replace hard disk and rebuild it.
37	S	A hard disk failed because of gross error on SCSI processor.	Gross error occurred to the on-board SCSI processor.	Replace hard disk and rebuild it.
38	S	A hard disk failed because of bad tag from the device.	The device responded with an invalid tag.	Replace hard disk and rebuild it.

Table B-1 Error-Event Codes (**continued**)

Error-Event #	Type	Description	Additional Details	User Action Required
39	S	A hard disk failed because command to the device timed out.	SCSI command timed out on the device.	Replace hard disk and rebuild it.
40	S	A hard disk failed because of the system reset.	System reset occurred.	Replace hard disk and rebuild it.
41	S	A hard disk failed because of busy status or parity error.	The device returned busy status. The SCSI transaction with the device met with parity error.	Replace hard disk and rebuild it.
42	S	A hard disk set to failed state by host.	Command from host set the hard disk to failed state.	Replace hard disk and rebuild it.
43	S	A hard disk failed because access to the device met with a selection time out.	Device disconnected or powered off. Bad device.	Replace hard disk and rebuild it. Check power and cabling.
44	S	A hard disk failed because of a sequence error in the SCSI bus phase handling.	Bad hard disk.	Replace hard disk and rebuild it.
45	S	A hard disk failed because returned an unknown status.	Bad hard disk or incompatible device.	Replace hard disk or the device and rebuild it.
46	S	A hard disk failed because device is not ready.	Device not spinning, just turned bad. Power to the device failed.	Replace hard disk and rebuild it. Check power and rebuild device.
47	S	A hard disk failed because device was not found on startup.	Device not connected. Device not responding. Clear configuration suspend mode command was invoked.	Check setup. Check the startup option parameters on the system.
48	S	A hard disk failed because write operation of the 'Configuration On Disk' failed.	Bad hard disk. Device write protected.	Replace hard disk and rebuild it. Check the startup option parameters on the system.

Table B-1 Error-Event Codes (continued)

Error-Event #	Type	Description	Additional Details	User Action Required
49	S	A hard disk failed because write operation of 'Bad Data Table' failed.	Bad hard disk. Device write protected.	Replace hard disk and rebuild it. Check the startup option parameters on the system.
50	W	Physical device status changed to offline.		None
51	I	Physical device status changed to Hot Spare.		None
52	W	Physical device status changed to rebuild.		None
53	W	Physical device ID did not match.		None
54	E	Physical device failed to start.		Reset the physical disk. Replace the disk.
55	W	Physical device negotiated different offset than configuration.		None
56	W	Physical device negotiated different bus width than configuration.		None
57	S	Physical drive missing on startup.	Physical drive missing.	Replace the physical drive or power on all enclosures.
58	E	Rebuild startup failed due to lower disk capacity.	Device capacity not sufficient for doing rebuild.	Replace with a disk having sufficient capacity.
59	W	Physical drive is switching from a channel to the other channel.	Physical drive removed or channel failed.	None
60	E	Temporary-Dead physical drive is automatically made online.	Temporary-Dead state caused because of transient errors.	Analyze event log to find out why the drive was marded DEAD.

Table B-1 Error-Event Codes (**continued**)

Error-Event #	Type	Description	Additional Details	User Action Required
61	I	A standby rebuild was started.	A physical device failed and spare was available.	None
62	I	Hot spare replaced with a smaller capacity physical disk.	The new hot spare may have a smaller physical capacity than the physical disk it replaced. The controller coercion setting may have reduced the configurable size of the new hot spare.	None
67	E	Physical disk found on only one disk channel.	Physical disk is connected on only one channel.	Inspect disk channel cables and related hardware for proper operation.
69	E	Physical disk has acquired an inappropriate loop ID. Enclosure disk-slot operations are disabled while this condition persists.	Enclosure selector switch conflict. Physical disk hardware failure. Enclosure disk slot hardware failure.	Ensure that each disk enclosure selector switch is set to a unique number per enclosure manufacturer specification. Inspect physical disk connector. Power-cycle entire system. Replace physical disk. Replace disk enclosure.
70	E	Physical disk port has failed or cannot operate at the configured channel speed.	Physical disk hardware failure. Physical disk is not compatible with system. Enclosure disk slot hardware failure.	Replace physical disk. Replace disk enclosure.
71	E	Mirror race recovery failed for logical drive.	A read or write operation to a physical disk failed while restoring redundancy.	Run consistency check and restore consistency.
72	E	Controller parameters checksum verification failed; restored default.	NVRAM battery low. NVRAM hardware failure. Improper shutdown of the controller during controller parameter update.	Restore correct controller parameter settings. If problem persists, replace controller.
73	I	Online controller firmware upgrade has started.	User has initiated an online firmware upgrade.	None

Table B-1 Error-Event Codes (continued)

Error-Event #	Type	Description	Additional Details	User Action Required
74	I	Online controller firmware upgrade has completed successfully.	Online controller firmware upgrade has completed without error. The partner controller will now be auto flashed.	None
75	E	Online controller firmware upgrade has failed.	Online controller firmware upgrade has failed. The original firmware will be reloaded.	Use the offline method to load the new firmware.
96	S	Device loop ID conflict (soft addressing) detected.	Device loop ID conflict detected on disk channel resulting in soft addressing; potential data corruption.	Change index selector to enable hard addressing per enclosure manufacturer's specification.
128	N	Consistency check is started.	User started a consistency check. 'Raidbl' started consistency check.	None
129	I	Consistency check is finished.	Consistency check completed successfully without detecting any errors.	None
130	E	Consistency check is cancelled.	User cancelled the consistency check.	Restart consistency check, if required.
131	E	Consistency check on logical drive error.	Inconsistent data was found. Bad sectors were found. A physical device reliability problem.	See bad block and request sense table for more information.
132	E	Consistency check on logical drive failed.	A logical device became critical. A logical device failed.	See request sense data for more information.
133	S	Consistency check failed due to physical device failure.	A physical device failed.	See request sense data for more information.
134	S	Logical drive has been made offline.	One/multiple physical device(s) failed.	It may not be possible to recover from this error. Contact your service representative.
135	E	Logical device is critical.	One physical device failed.	Replace the physical device. Start the rebuild, if required.

Table B-1 Error-Event Codes (**continued**)

Error-Event #	Type	Description	Additional Details	User Action Required
136	I	Logical drive has been placed online.	Rebuild completed. User set the physical device online. New configuration was added.	None
137	I	An automatic rebuild has started on logical drive.	A physical device failed and a spare device was available. A spare physical device was found and replaced the failed device.	None
138	I	A manual rebuild has started on logical drive.	Client started the rebuild on user's request. User replaced the failed device and 'Raidbld' started the rebuild.	None
139	I	Rebuild on logical drive is over.	Rebuild completed successfully only for this logical drive.	None
140	W	Rebuild on logical drive is cancelled.	User cancelled rebuild. Higher priority rebuild started.	Restart the rebuild if required.
141	E	Rebuild stopped with error.	Due to an unknown error on the controller rebuild failed.	Try rebuild again.
142	E	Rebuild stopped with error. New device failed.	New physical device failed. New physical device is not compatible with MDAC hardware/firmware.	Replace the new device.
143	E	Rebuild stopped because logical drive failed.	At least one more physical device failed in the array.	It may not be possible to recover from this error. Contact your service representative.
144	I	Logical drive initialization started.	User started the initialization.	Any previous data is lost.
145	I	Logical drive initialization done.	Initialize operation completed successfully.	None
146	W	Logical drive initialization cancelled.	User cancelled the initialization.	Restart initialization if required.
147	E	Logical drive initialization failed.	One/multiple physical device(s) failed. Controller has been removed. Controller has been powered off.	Refer to the device failure event.

Table B-1 Error-Event Codes **(continued)**

Error-Event #	Type	Description	Additional Details	User Action Required
148	I	A logical drive has been found.	A new configuration has been added. MORE completed. A new controller has been plugged in. Controller has been powered on. System has rebooted.	None
149	I	A logical drive has been deleted.	A new configuration has been added. A new logical device has been deleted. Controller has been removed. Controller has been powered off.	None
150	I	Expand Capacity started.	User started the Online RAID Expansion operation.	None
151	I	Expand Capacity completed.	Online RAID Expansion completed.	None
152	E	Expand Capacity stopped with error.	Multiple physical devices failed.	It may not be possible to recover from this error. Contact your service representative.
153	C	Bad blocks found.	Bad sector was found on a physical device during consistency check/rebuild/RAID expansion operation.	Run a Consistency Check with the Restore option. Restore data from a backup.
154	I	System drive (LUN) size changed.	A new configuration has been added. RAID Expansion has added extra capacity.	None
155	I	System drive type changed.	A new configuration has been added. RAID Expansion completed on RAID 1.	None
156	S	Bad data blocks found. Possible data loss.	Bad blocks were found on multiple physical devices in same zone.	Restore data from a backup.
157	W	System drive LUN mapping has been written to config.		None
158	S	Attempt to read data from block that is marked in Bad Data Table.	Potential data loss.	Restore data from a backup.

Table B-1 Error-Event Codes (**continued**)

Error-Event #	Type	Description	Additional Details	User Action Required
159	E	Data for Disk Block has been lost due to Logical Drive problem.	Data retained in RAID Cache for a Write-back Logical Drive cannot be stored to the physical medium because of Logical Drive problem. The Logical Drive problem could be because of multiple physical devices offline or other reasons.	Insure that all the Physical Drives related to the Logical Drive, disk channel, enclosure, or cabling are functional and accessible. Repair or replace them if necessary.
160	E	Temporary-Offline RAID5/RAID3 array is available to the user again with the possibility of data loss in the array.	Temporary-Offline state caused because of transient errors in physical drives.	Verify data from backup.
161	E	Temporary-Offline RAID0+1/RAID1/RAID0/JBOD array is available to the user again.	Temporary-Offline state caused because of transient errors in physical drives.	None
162	I	A standby rebuild has started on logical drive.	A physical drive failed and a spare device was available.	None
176	I	Logical drive background initialization started.	User may have started background initialization. Firmware may have automatically started background initialization.	None
177	I	Logical drive background initialization stopped.	User may have stopped background initialization. Firmware may have automatically stopped background initialization.	None
178	I	Logical drive background initialization paused.	Background initialization paused due to a higher priority operation.	None
179	I	Logical drive background initialization restarted.	Background initialization started after being paused.	None
180	I	Logical drive background initialization failed.	Background initialization failed.	None.

Table B-1 Error-Event Codes (continued)

Error-Event #	Type	Description	Additional Details	User Action Required
181	I	Logical drive background initialization completed.	Background initialization completed successfully.	None
182	E	Low battery charge level. Logical drive may have lost data.	Controller was powered off for duration longer than battery capacity. User connected a new controller. User connected a new BBU battery.	Run consistency check to verify logical drive consistency. If needed, restore data from backup.
256	S	Fan failure.	Cable connection broken. Bad fan.	Replace fan.
257	I	Fan has been restored.	Faulty fan has been replaced. Cable is connected properly.	None
258	S	Fan failure.	Cable connection broken. Bad fan.	Replace fan.
259	I	Storage cabinet fan is not present.	Enclosure Management Connection is broken. Management hardware is bad. Fan is not present.	Follow enclosure management vendor's diagnostics and repair procedures.
272	S	Power supply failure.	Cable connection is broken. Bad power supply.	Reconnect cable or replace the power supply as required.
273	I	Power supply has been restored.	Faulty power supply has been replaced.	None
274	S	Power supply failure.	Cable connection is broken. Bad power supply.	Replace power supply.
275	I	Storage cabinet power supply is not present.	Management connection is broken. Management hardware is bad. Power supply is not present.	Follow enclosure management vendor's diagnostics and repair procedures.
288	S	Over temperature. Temperature is above 70 degrees Celsius.	Room temperature is too high. Bad fan. Bad sensor.	Turn off the system and allow it to cool down. Adjust the room temperature.
289	W	Temperature is above 50 degrees Celsius.	Room temperature is high. Bad fan.	Replace fan. Turn off the system. Adjust the room temperature.

Table B-1 Error-Event Codes (continued)

Error-Event #	Type	Description	Additional Details	User Action Required
290	I	Normal temperature has been restored.	Faulty fan has been replaced. Room temperature was reduced.	None
291	S	Over temperature.	Room temperature is too high. Bad fan.	Turn off the system and allow it to cool down. Adjust the room temperature.
292	I	Storage cabinet temperature sensor is not present.	Enclosure management connection is broken. Management hardware is bad. Sensor is not present.	Follow enclosure management vendor's diagnostics and repair procedures.
304	S	Enclosure reported failure state.	Power supply failed. Fan failed. Cabinet is too hot.	Follow enclosure management vendor's diagnostics and repair procedures.
305	W	Enclosure reported critical state.	Not available.	Not available.
306	I	Storage Works enclosure reported normal state.	Problem has been rectified.	None
307	I	Uninterruptible power supply disabled.		None
308	I	Uninterruptible power supply AC failed.		None
309	W	Uninterruptible power supply battery low.		None
310	S	Uninterruptible power supply failed.		None
311	I	Uninterruptible power supply normal.		None
320	S	Fan failure.	Cable connection broken. Bad fan.	Replace fan.
321	I	Fan has been restored.	Faulty fan has been replaced. Cable is connected properly.	None

Table B-1 Error-Event Codes (**continued**)

Error-Event #	Type	Description	Additional Details	User Action Required
322	I	Fan is not present.	Enclosure Management Connection is broken. Management hardware is bad. Fan is not present.	Follow enclosure management vendor's diagnostics and repair procedures.
323	S	Power supply failure.	Cable connection is broken. Bad power supply.	Replace the power supply.
324	I	Power supply has been restored.	Faulty power supply has been replaced.	None
325	I	Power supply is not present.	Management connection is broken. Management hardware is bad. Power supply is not present.	Follow enclosure management vendor's diagnostics and repair procedures.
326	S	Temperature is over safe limit. Failure imminent.	Room temperature is too high. Bad fan. Bad sensor.	Turn off the system and allow it to cool down. Adjust the room temperature.
327	W	Temperature is above working limit.	Room temperature is high. Bad fan.	Replace fan. Turn off the system. Adjust the room temperature.
328	I	Normal temperature has been restored.	Faulty fan has been replaced. Room temperature was reduced.	None
329	I	Temperature sensor is not present.	Enclosure management connection is broken. Management hardware is bad. Sensor is not present.	Follow enclosure management vendor's diagnostics and repair procedures.
330	W	Enclosure access critical.	Enclosure management connection is broken. Management hardware is bad.	Follow enclosure management vendor's diagnostics and repair procedures.
331	I	Enclosure access has been restored.	Enclosure has been fixed or replaced.	None

Table B-1 Error-Event Codes (**continued**)

Error-Event #	Type	Description	Additional Details	User Action Required
332	S	Enclosure access is offline.	Enclosure management connection is broken. Management hardware is bad.	Follow enclosure management vendor's diagnostics and repair procedures.
333	S	Enclosure Soft Addressing Detected.	Enclosure has duplicate loop ids (Soft Addressing). Potential data corruption.	Change index selector to enable hard addressing as per enclosure manufacturer's specification.
334	I	Enclosure services ready.		None
335	I	Access to temperature sensor has been lost.	Switch card or temperature sensor has been removed.	None
336	I	Access to power supply status information has been lost.	Switch card or connectivity has been removed.	None
337	I	Access to fan status information has been lost. Switch card or connectivity has been removed.		None
384	I	Array management server software started successfully.	The server system (or array management utility server) started.	If you did not expect a system reboot, investigate.
385	E	Write back error.	Data cache write failed.	The data may have been lost. Restore the data from a backup.
386	W	Internal log structures getting full; PLEASE SHUTDOWN AND RESET THE SYSTEM IN THE NEAR FUTURE.	Too many configuration changes occurred since the last boot.	Reboot the system by power cycling whenever convenient.

Table B-1 Error-Event Codes (**continued**)

Error-Event #	Type	Description	Additional Details	User Action Required
388	C	Controller is dead. System is disconnecting from this controller.		Contact your service representative.
389	W	Controller has been reset.	Controller failed. Controller was removed from the system. Controller has been powered off.	None
390	I	Controller is found.	Driver has reset the controller to recover from an error. Driver has reset the controller to activate new firmware.	None
391	C	Controller is gone. System is disconnecting from this controller.	New controller has been installed. Controller has been powered on. System has rebooted.	None
392	I	BBU Present.	Controller is dead. Controller has been removed. Controller has been powered off.	None
393	W	BBU Power Low.	A BBU unit was found on the controller.	If this message occurs without power failure, replace the BBU.
394	I	BBU Power OK.	BBU does not have enough power to enable the write data cache.	None
395	C	Controller is gone. System is disconnecting from this controller.	BBU has enough power to enable the write data cache.	None
396	I	Controller powered on.	Controller was removed from the system. Controller has been powered off.	None
397	I	Controller is online.	New controller has been installed.	None
398	C	Controller is gone. System is disconnecting from this controller.	Controller was set online.	None
399	W	Controller's partner is gone, controller is in failover mode now.	Controller was set offline.	If you did not expect this, investigate.

Table B-1 Error-Event Codes (**continued**)

Error-Event #	Type	Description	Additional Details	User Action Required
400	I	BBU reconditioning is started.	User started a BBU reconditioning.	None
401	I	BBU reconditioning is finished.	BBU reconditioning completed successfully.	None
402	I	BBU reconditioning is cancelled.	User cancelled the BBU reconditioning.	Restart the BBU reconditioning, if required.
403	S	Installation aborted.		Installation aborted.
404	S	Controller firmware mismatch.	Replacement controller with downlevel firmware installed.	Reload controller firmware.
405	W	BBU removed.	BBU physically removed.	Reinstall BBU.
406	S	WARM BOOT failed.	Memory error detected during WARM boot scan. Possible data loss.	Restore data from a backup.
407	I	BBU calibration cycle started.	New battery detected.	None
408	I	BBU calibration cycle finished.	BBU calibration completed successfully.	None
409	I	BBU calibration cycle is cancelled.	User cancelled the BBU calibration cycle.	None
410	I	BBU battery not present.	A BBU is present, but the battery is not detected.	Install or connect the battery.
411	W	Controller entered Conservative Cache Mode.		None
412	W	Controller entered normal cache mode.		None
413	W	Controller device start complete.		None
414	W	Soft ECC error corrected.	Faulty memory module.	Replace memory module.

Table B-1 Error-Event Codes (continued)

Error-Event #	Type	Description	Additional Details	User Action Required
415	W	Hard ECC error corrected.	Faulty memory module.	Replace memory module.
416	S	BBU recondition needed.		None
417	W	Controller's partner has been removed.		None
418	E	BBU out of service.	BBU will not be able to power the cache if AC power fails. Firmware will switch WriteBack logical drives to WriteThrough.	Replace BBU.
419	W	Updated partner's status.		None
420	W	Relinquished partner.		None
421	W	Inserted partner.		None
422	W	Dual controllers enabled.		None
423	W	Killed partner.		None
424	W	Dual controllers entered nexus.		None
425	S	Controller boot ROM image needs to be reloaded.	Wrong firmware image file downloaded. MAC address changed.	Contact your service representative to reload the boot ROM image.
426	C	Controller is using default non-unique world-wide name.	MAC address lost or not set.	Contact your service representative to set the controller MAC address.
427	E	Mirror Race recovery failed.	Some physical devices could have failed.	Run consistency check and restore consistency.
428	C	Mirror Race on critical drive.	Logical device is critical.	Replace dead drive and rebuild.
512	I	System started.	The server system (or array management utility server) started.	If you did not expect a system reboot, investigate.

Table B-1 Error-Event Codes (**continued**)

Error-Event #	Type	Description	Additional Details	User Action Required
513	I	Size table full.	Too much physical device size information is defined.	Remove unused device information for this system.
514	I	User logged in.	An array management utility user logged in on the server system.	Not available.
515	I	User logged out.	An array management utility user logged out of the server system.	Not available.
516	I	Server alive.	Reconnected to server. Server rebooted.	None
517	S	Lost connection to server, or server is down.	Lost network connection to server. Server shutdown.	None
518	I	Automatic reboot count has changed.	Controller has rebooted. Automatic reboot has rearmed itself or was reconfigured.	None
640	W	Channel failed.	Cable disconnected.	Plug in cable.
641	W	Channel online.	Cable reconnected.	None
642	S	Back end SCSI bus dead.	Lost access to data on SCSI bus.	LRC module may need replacing.
643	I	Back end SCSI bus alive.	Regained access to data on SCSI bus.	None
644	S	Back end fibre dead.	Lost access to data on fibre channel.	LRC module may need replacing.
645	I	Back end fibre alive.	Regained access to data on fibre channel.	None
700	W	Event log empty.	Tried to read past last entry.	None
701	W	Event log entries lost.	Tried to read an entry that does not exist in the event log.	None
702	W	Request sense.	A physical drive has generated an error.	Interpret the Key/ASC/ASCQ and take appropriate action.
703	W	Set real time clock.	Real time clock was set.	None

Table B-1 Error-Event Codes (continued)

Error-Event #	Type	Description	Additional Details	User Action Required
800	W	New configuration received.	A new configuration was downloaded to controller.	None
801	W	Configuration cleared.	Controller was told to clear the configuration.	None
802	W	Configuration invalid.	The controller found an invalid configuration.	None
803	W	Configuration on disk access error.	The controller could not read the configuration off of the disk.	None
804	W	Configuration on disk converted.	The controller converted a down level configuration on disk.	None
805	W	Configuration on disk import failed.	The controller could not import the configuration.	None
806	I	A debug dump exists on this system.	The controller aborted and created debug dump information.	Contact field support for assistance in retrieving the data.
807	I	A debug dump exists on this system.	The partner controller aborted and created debug dump information.	Contact field support for assistance in retrieving the data.
896	S	Internal controller is in the hung state.		Power controller off and on.
897	S	Internal controller has encountered a firmware breakpoint.		Power controller off and on.
912	S	Internal controller has encountered i960 processor specific error.		Power controller off and on.
928	S	Internal controller has encountered Strong-ARM processor specific error.		Power controller off and on.

