

**SGI® Total Performance 9100 (2Gb TP9100)
Storage System User's Guide**

007-4522-003

CONTRIBUTORS

Written by Matt Hoy

Illustrated by Kelly Begley

Production by Glen Traefald and Karen Jacobson

Engineering contributions by Terry Fliflet, David Lucas, Van Tran, Michael Raskie and Ted Wood

COPYRIGHT

© 2002–2003, 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

Silicon Graphics, SGI, the SGI logo, IRIX, and Origin are registered trademarks, and CXFS, FailSafe, Octane2, ProPack, SGI Advanced Linux, and Silicon Graphics Fuel are trademarks of Silicon Graphics, Inc., in the United States and/or other countries worldwide.

Brocade and Silkworm are registered trademarks of Brocade Communications, Inc.

Mylex is a registered trademark of Mylex Corporation, and LSI Logic business unit. QLogic is a trademark of QLogic Corporation.

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

Record of Revision

Version	Description
001	August 2002 Original printing
002	February 2003 Engineering revisions
003	October 2003 Firmware and hardware revisions

Contents

Figures	ix
Tables	xi
About This Guide.	.xiii
Audience	.xiii
Structure of This Document	.xiii
Related Publications	.xiv
Conventions Used in This Guide	. xv
Product Support	. xv
Reader Comments.	.xvi
1. Storage System Overview	1
Overview of Storage System Features	1
RAID Configuration Features	2
Release 6.0 Features	2
RAID Fault Tolerance and Flexibility Features	4
JBOD Configuration Features	5
Availability Features	5
Supported Platforms	5
Compatibility	6
Storage System Enclosure	6

Enclosure Components	8
Operators (Ops) Panel.	11
PSU/Cooling Module.	12
RAID LRC I/O Modules	14
RAID Loopback LRC I/O Modules	16
JBOD LRC I/O Module	18
Drive Carrier Module	19
Enclosure Bay Numbering	22
Storage System Rack	26
Rack Structure	26
Power Distribution Units (PDUs)	29
Opening and Closing the Rear Rack Door	31
Storage System Tower	32
2. Connecting to a Host and Powering On and Off	35
Connecting to a Host	35
Grounding Issues	37
Connecting the Power Cords and Powering On the 2 Gb TP9100 Tower	37
Checking AC Power and Storage System Status for the Tower	39
Connecting the Power Cords and Powering On the 2 Gb TP9100 Rack	40
Checking Grounding for the Rack	42
Powering On the Rack	42
Checking AC Power and System Status for the Rack	43
Powering Off	44
Powering Off the 2 Gb TP9100 Rack.	45
Powering Off the 2 Gb TP9100 Tower or a Single Enclosure	46
3. Features of the RAID Controller.	47
Enclosure Services Interface (ESI) and Disk Drive Control	47
Configuration on Disk (COD).	48
Drive Roaming	49
Data Caching	50
Write Cache Enabled (Write-back Cache Mode)	50
Write Cache Disabled (Write-through or Conservative Cache Mode)	51

RAID Disk Topologies	52
Simplex Single-port RAID Topology	53
Duplex Single-Port RAID Topology	54
Simplex Dual-Port RAID Topology.	55
Duplex Dual-Port RAID Topology	56
Dual-port Duplex Two-Host RAID Configuration	57
Dual-Port Duplex RAID Configuration	58
4. Using the RAID Controller	59
Software Tools for the Controller	59
RAID Levels	60
CAP Strategy for Selecting a RAID Level	61
Configuring for Maximum Capacity	62
Configuring for Maximum Availability	63
Configuring for Maximum Performance	65
Disk Topologies	65
System Drives	68
System Drive Properties	68
System Drive Affinity and Programmable LUN Mapping	69
Drive State Reporting	69
Automatic Rebuild	71
5. Troubleshooting	75
RAID Guidelines	76
Solving Initial Startup Problems	76
Using Storage System LEDs for Troubleshooting	78
ESI/Ops Panel LEDs and Switches	79
Power Supply/Cooling Module LEDs	84
RAID LRC I/O Module LEDs	85
RAID Loopback LRC I/O Module LEDs	89
JBOD LRC I/O Module LEDs	89
Drive Carrier Module LEDs	90
Using the Alarm for Troubleshooting	91

Solving Storage System Temperature Issues 92
Thermal Control 92
Thermal Alarm. 93
Using Test Mode 93
Care and Cleaning of Optical Cables. 94
6. Installing and Replacing Drive Carrier Modules 95
Adding a Drive Carrier Module 95
Replacing a Drive Carrier Module 98
LUN Integrity and Drive Carrier Module Failure 99
Replacing the Disk Drive Module	100
A. Technical Specifications	103
Storage System Physical Specifications	103
Environmental Requirements.	105
Power Requirements	106
LRC I/O Module Specifications	107
Disk Drive Module Specifications	108
SGI Cables for the 2 Gb TP9100 Storage System	108
B. Regulatory Information	109
FCC Warning	109
Attention	109
European Union Statement	109
International Special Committee on Radio Interference (CISPR)	110
Canadian Department of Communications Statement.	110
Attention	110
VCCI Class 1 Statement	110
Class A Warning for Taiwan	111
Index.	113

Figures

Figure 1-1	Front View of Rackmount Enclosure.	7
Figure 1-2	Rear View of Rackmount Enclosure	8
Figure 1-3	Front View of Enclosure Components	9
Figure 1-4	RAID (Base) Enclosure Components, Rear View	10
Figure 1-5	JBOD (Expansion) Enclosure Components, Rear View	11
Figure 1-6	Ops Panel	12
Figure 1-7	PSU/Cooling Module	13
Figure 1-8	PSU/Cooling Module Switches and LEDs	14
Figure 1-9	Dual-port RAID LRC I/O Module	15
Figure 1-10	Single-port RAID LRC I/O Module	16
Figure 1-11	Single-port RAID Loopback LRC I/O Module	17
Figure 1-12	Dual-port RAID Loopback LRC I/O Module	18
Figure 1-13	JBOD LRC I/O Module	19
Figure 1-14	Drive Carrier Module	20
Figure 1-15	Anti-tamper Lock	21
Figure 1-16	Dummy Drive Carrier Module	22
Figure 1-17	Rackmount Enclosure Bay Numbering and Module Locations	23
Figure 1-18	Tower Enclosure Bay Numbering and Module Locations.	25
Figure 1-19	Example of 2 Gb TP9100 Rack (Front View).	27
Figure 1-20	Example of 2 Gb TP9100 Rack (Rear View)	28
Figure 1-21	PDU Locations and Functions	30
Figure 1-22	Opening the Rack Rear Door	31
Figure 1-23	Front View of Tower	32
Figure 1-24	Rear View of Tower	33
Figure 1-25	Tower Storage System Power Cords	34
Figure 2-1	Power Cords for the Tower	38
Figure 2-2	ESI/Ops Panel LEDs and Switches	39

Figure 2-3	Rack Power Cabling 41
Figure 2-4	Rackmount Enclosure ESI/Ops Panel Indicators and Switches 44
Figure 3-1	Simplex Single-port RAID Topology 53
Figure 3-2	Duplex Single-port RAID Topology 54
Figure 3-3	Simplex Dual-port Dual-host RAID Topology 55
Figure 3-4	Duplex Dual-port RAID Configuration 56
Figure 3-5	Dual-port Dual-path Attached Duplex RAID Topology 57
Figure 3-6	Dual-port Quad-path Duplex RAID Topology 58
Figure 4-1	Example of RAID Levels within a Drive Pack (LUN) 61
Figure 4-2	Tower I/O Modules, Channels, and Loops 66
Figure 4-3	Rackmount Enclosure I/O Modules, Channels, and Loops (Front View) 67
Figure 5-1	ESI/Ops Panel Indicators and Switches 79
Figure 5-2	Power Supply/Cooling Module LED 84
Figure 5-3	Dual-port RAID LRC I/O Module LEDs 85
Figure 5-4	Single-port RAID LRC I/O Module LEDs 87
Figure 5-5	JBOD LRC I/O Module LEDs 89
Figure 5-6	Drive Carrier Module LEDs 90
Figure 6-1	Unlocking the Drive Carrier Module 96
Figure 6-2	Opening the Module Handle 97
Figure 6-3	Inserting the Disk Drive Module in a Rackmount Enclosure 97
Figure 6-4	Locking the Drive Carrier Module 98
Figure 6-5	Unlocking the Disk Drive Module 100
Figure 6-6	Removing the Drive Carrier Module 101

Tables

Table i	Document Conventions	xv
Table 4-1	Supported RAID Levels	60
Table 4-2	RAID Level Maximum Capacity	62
Table 4-3	Array Operating Conditions	63
Table 4-4	RAID Levels and Availability.	64
Table 4-5	RAID Levels and Performance	65
Table 4-6	Physical Disk Drive States.	70
Table 5-1	ESI/Ops Panel LEDs	79
Table 5-2	Ops Panel Configuration Switch Settings for JBOD	80
Table 5-3	Ops Panel Configuration Switch Settings for RAID	82
Table 5-4	Dual-port RAID LRC I/O Module LEDs	86
Table 5-5	Single-port RAID LRC I/O Module LEDs	87
Table 5-6	JBOD LRC I/O Module LEDs.	90
Table 5-7	Disk Drive LED Function	90
Table 5-8	Thermal Alarms	93
Table A-1	Dimensions	103
Table A-2	Weights	104
Table A-3	Power Specifications	104
Table A-4	Ambient Temperature and Humidity Requirements	105
Table A-5	Environmental Specifications.	105
Table A-6	Minimum Power Requirements	106
Table A-7	Rack PDU Power Specifications	106
Table A-8	LRC I/O Module Specifications	107
Table A-9	Drive Carrier Module Specifications (1.6-inch 36-GB Drive)	108
Table A-10	SGI Fibre Channel Fabric Cabling Options for the 2 Gb TP9100 Storage System	108

About This Guide

This guide explains how to operate and maintain the SGI 2 Gb Total Performance 9100 (2 Gb TP9100) Fibre Channel storage system. As part of the SGI Total Performance Series of Fibre Channel storage, this storage system provides compact, high-capacity, high-availability RAID and JBOD (“just a bunch of disks”) storage for supported SGI servers.

The 2 Gb TP9100 storage system can be connected to one or more Fibre Channel boards (host bus adapters, or HBAs) in the SGI server separately or in combination (loop). Software interfaces from a third party are shipped with the storage system.

Audience

This guide is written for users of the SGI 2 Gb TP9100 Fibre Channel storage system. It presumes general knowledge of Fibre Channel technology and knowledge of the host SGI server, the HBA, and other Fibre Channel devices to which the storage system might be cabled.

Structure of This Document

This guide consists of the following chapters:

- Chapter 1, “Storage System Overview,” describes storage system formats and the modules in the storage system.
- Chapter 2, “Connecting to a Host and Powering On and Off,” explains how to cable the storage system to a host, how to connect the power cord, and how to power the storage system on and off.
- Chapter 3, “Features of the RAID Controller,” describes SCSI Enclosure Services (SES), configuration on disk (COD), drive roaming, Mylex Online RAID Expansion (MORE), and data caching.

- Chapter 4, “Using the RAID Controller,” introduces software tools for the controller, gives configuration information, and explains RAID levels and criteria for selecting them, storage system drives and drive state management, and automatic rebuild.
- Chapter 5, “Troubleshooting,” describes storage system problems and suggests solutions. It explains how to use storage system LEDs and the storage system alarm for troubleshooting.
- Chapter 6, “Installing and Replacing Drive Carrier Modules,” explains how to add a new disk drive module and how to replace a defective disk drive module.
- Appendix A, “Technical Specifications,” gives specifications for the storage system in general and for specific modules.
- Appendix B, “Regulatory Information,” contains Class A regulatory information and warnings for the product.

An index completes this guide.

Related Publications

Besides this manual and the manuals for the storage system third-party software, locate the latest versions of the user’s guide for the server and for any other Fibre Channel devices to which you are attaching the storage (such as the SGI Fibre Channel Hub or switch).

If you do not have these guides, you can find the information online in the following locations:

- IRIS InSight Library: From the Toolchest, select **Help > Online Books > SGI EndUser** or **SGI Admin**, and select the applicable guide.
- Technical Publications Library: If you have access to the Internet, see: <http://docs.sgi.com>.

Conventions Used in This Guide

Table i contains the conventions used throughout this guide.

Table i Document Conventions

Convention	Meaning
Command	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.
<code>user input</code>	Fixed-space font denotes literal items that the user enters in interactive sessions. Output is shown in nonbold, fixed-space font.
Hardware	This font denotes a label on hardware, such as for a port or LED.
[]	Brackets enclose optional portions of a command or directive line.

Product Support

SGI provides a comprehensive product support and maintenance program for its products. If you are in North America and would like support for your SGI-supported products, contact the Technical Assistance Center at 1-800-800-4SGI or your authorized service provider. If you are outside north America, contact the SGI subsidiary or authorized distributor in your country.

Reader Comments

If you have comments about the technical accuracy, content, or organization of this document, please 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 can be found on the back cover.)

You can contact us in any of the following ways:

- Send e-mail to the following address:
techpubs@sgi.com
- Use the Feedback option on the Technical Publications Library World Wide Web page:
<http://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

SGI values your comments and will respond to them promptly.

Storage System Overview

The SGI 2 Gb Total Performance 9100 (2 Gb TP9100) Fibre Channel storage system provides you with a high-capacity, high-availability Fibre Channel storage solution. The storage system can be configured for JBOD (“just a bunch of disks”) or RAID (“redundant array of inexpensive disks”) operation, and is available in both rackmount and tower formats. The modular design of the 2 Gb TP9100 expands easily to meet your needs.

The following sections describe the structure and features of the storage system:

- “Overview of Storage System Features” on page 1
- “Storage System Enclosure” on page 6
- “Enclosure Components” on page 8
- “Storage System Rack” on page 26
- “Storage System Tower” on page 32

Overview of Storage System Features

The features of the SGI 2 Gb TP9100 storage system are outlined in the following sections:

- “RAID Configuration Features” on page 2
- “JBOD Configuration Features” on page 5
- “Availability Features” on page 5
- “Supported Platforms” on page 5
- “Compatibility” on page 6

RAID Configuration Features

- 64-drive maximum configuration
- 32 logical units maximum

Release 6.0 Features

Version 9.03 FFx2 controller firmware introduces the following new features over the 8.40 and 8.50 controller firmware versions:

Disk Drive Predictive Failure Analysis

Predictive failure analysis (PFA) monitors events on disk drives, sets the frequency of the interval used to poll individual drives for PFA events, instructs the controller as to what action should be taken regarding a disk that reports a PFA event, and scans all drives as a one time maintenance check. This information is an indication of an error or detected problem with the physical disk drive and the drive should be replaced. Mylex controllers recognize the PFA information and set the device state to offline failed.

The user can specify what to do when a disk drive reports a PFA condition: either place the drive in the offline state and begin a rebuild process, or inform the user via the configuration/monitoring interface that an event has occurred. The user can also launch a single maintenance check and set the interval for repeated polling of PFA events.

Rolling Upgrade Capability

This feature is available in firmware releases 9.03 and higher.

Note: Controller firmware may be upgraded from 9.03 to a later release without interruption. However, rolling upgrades can only be done on dual-active controllers that are configured for failover/failback operation.

The rolling firmware upgrade feature enables the user of a dual-active system to upgrade the controller firmware to a new release level while online and without losing data access. During the process, all I/O is routed first through one controller while the other controller is flashed and restarted with the new firmware image. Once the new firmware image is installed on one controller, I/O is shifted to the upgraded controller, and the

second controller is updated with new firmware image. When the second controller is restarted, I/O is again routed to its normal path.

Note: Rolling upgrades will only be supported on SGI IRIX platforms (as listed in “Supported Platforms” on page 5) and SGI Altix series servers running an SGI Linux environment of 7.2 or later with SGI ProPack 2.1, or SGI Advanced Linux Environment 2.1 or later with SGI ProPack 2.2 or later.

You may experience brief delays in performance, similar to those experienced in a failover/failback operation. During the upgrade, you will lose redundancy due to the fact that each controller is alternately taken offline while it is being flashed with a new firmware image. The cache on both controllers is flushed and both controllers are placed into conservative cache mode during the upgrade process.

If a power cycle occurs during the firmware upgrade and the firmware of one controller is corrupted, the other controller will continue to have access to data when power is returned. Because firmware is flashed sequentially to each controller it is not possible for the firmware of both controllers to be corrupted by a power failure. Check the firmware levels of each controller if you suspect that a controller’s firmware image has been corrupted. If the firmware levels are not at the desired version, repeat the rolling upgrade process.

SAN Mapping 256

Firmware releases 8.40 and later support up to 256 host nodes. TPM 1.5 and later provide changes required for 256 host node support. TPM 1.4 and earlier releases support up to 64 host nodes.

Drive Sizing Algorithm

On new configurations (only), a new drive sizing algorithm reduces the maximum amount of capacity that the RAID controller truncates from disk drives larger than 80GB on 2-Gb/s FFX2 RAID controllers running 9.x or later. On a 146 GB disk drive, the new algorithm increases the usable capacity by about 5 GB per drive.

Note: Upgrading from 8.40 or 8.50 to 9.03 or later does not cause the new drive sizing algorithm to be used. The drive sizing algorithm only applies to new configurations where no logical units or previous configurations have existed.

RAID Fault Tolerance and Flexibility Features

- 1 to 16 disk drives can be combined into a pack (15+1 RAID group)
- 5 RAID Levels (0, 1, 0+1, 3, and 5)

Note: SGI does not support the mixing of any non-redundant (RAID 0) and redundant (fault-tolerant) RAID levels within a drive pack (LUN).

- 32 LUNs (15+1 RAID group)
- 1 Gb/s or 2 Gb/s front end (FE) and back end (BE)
Fibre Channel arbitrated loop (FC-AL)
- Immediate LUN availability (ILA)
- Transparent disk drive rebuilds
- Variable stripe size per controller (8K, 16K, 32K, and 64K)
- Mirrored cache
- Drive roaming during power off
- Cache coherency
- Transparent failover and failback
- Automatic error recovery
- Write through, write back, or read ahead support
- Automatic detection of failed drives
- Automatic drive rebuilds, using “hot spare” drive
- Hot-swappable drives
- SAN mapping server to LUN mapping

- Automatic firmware flashing
In a dual controller configuration, the firmware of the replacement controller is automatically flashed to match the firmware of the surviving controller.

Features supported by release 6.0 are as follows:

- RAID controllers that make use of predictive failure analysis (PFA) on disk drives
- Rolling upgrade capability
- Support for up to 256 host nodes
- New drive sizing algorithm

JBOD Configuration Features

- 96 drive maximum configuration
- 1x16 (more storage) and 2x8 (more bandwidth) disk topologies

Availability Features

- Dual power feeds with dual power supplies
- Redundant cooling
- Battery back-up (BBU) maintains cache in case of power failure
- IRIX path failover
- Dynamic hot-sparing
- Non-disruptive component replacement
- Enclosure services interface (ESI) for SCSI enclosure services (SES)

Supported Platforms

- Software: SGI IRIX operating system 6.5.16 or later, SGI Linux Environment of 7.2 with SGI ProPack 2.1 and SGI Advanced Linux Environment 2.1 or later with SGI ProPack 2.2 or later
- Hardware: SGI Altix series, SGI Origin 200, Origin 300, Origin 350, Origin 2000, and Origin 3000 family servers. Silicon Graphics Octane, Silicon Graphics Octane2, Silicon Graphics Fuel visual workstations and Silicon Graphics Tezro Workstations

Note: Rolling upgrades will only be supported on SGI IRIX platforms (as listed in “Supported Platforms” on page 5) and SGI Altix series servers running an SGI Linux environment of 7.2 or later with SGI ProPack 2.1, or SGI Advanced Linux Environment 2.1 or later with SGI ProPack 2.2 or later.

Compatibility

Note: Copper Fibre Channel host bus adapters (HBAs) are not supported by the TP9100 (2Gb TP9100).

- QLogic 2200 optical 33/66-MHz HBA
- QLogic 2310 optical 66-MHz HBA
- QLogic 2342 optical 66-MHz dual channel HBA
- Brocade family SAN switches
 - SilkWorm 2400 1-Gb/s 8-port switch
 - SilkWorm 2800 1-Gb/s 16-port switch
 - SilkWorm 3200 2-Gb/s 8-port switch
 - SilkWorm 3800 2 Gb/s 16-port switch
 - SilkWorm 3900 2-Gb/s 32-port switch
 - SilkWorm 12000 2-Gb/s 64-port switch
- IRIX release level 6.5.16 or later

Storage System Enclosure

The enclosure is the basic unit of the SGI 2 Gb TP9100 storage system. Each enclosure contains a minimum of 4 and maximum of 16 disk drives and the component modules that handle I/O, power and cooling, and operations.

The enclosure is available in two formats: RAID (Redundant Array of Inexpensive Disks) and JBOD (Just a Bunch of Disks). An enclosure with single or dual RAID modules is a

RAID (base) enclosure. An enclosure without a RAID module is a *JBOD* or *expansion enclosure*. The expansion enclosure can be cabled to a RAID enclosure and provides additional disk modules. The RAID controller can address up to 64 disk drives; thus, three expansion enclosures can be cabled to it.

Enclosures can be installed in industry standard 19-in. racks or be configured as a stand-alone tower. Figure 1-1 shows the front view of a rackmount enclosure fully populated with drives.




Figure 1-1 Front View of Rackmount Enclosure

Figure 1-2 shows the rear view of a rackmount enclosure.




Figure 1-2 Rear View of Rackmount Enclosure

Enclosure Components

The enclosure contains the following component modules (see Figure 1-3 and Figure 1-4):

- Integrated operators panel (ops panel)
- Two power supply cooling modules (PSU/cooling modules)
- One or two loop resiliency circuit input/output (LRC I/O) modules with optional integrated Mylex FFX-2 RAID controllers

Note: In simplex RAID configurations, the enclosure will contain a RAID loopback LRC module in place of one of the RAID LRC I/O modules.

- Up to 16 disk drive carrier modules
- Dummy drive carrier modules

Figure 1-3 shows a front view of the enclosure components.




Figure 1-3 Front View of Enclosure Components

Figure 1-4 shows a rear view of the RAID (base) enclosure components.





Figure 1-4 RAID (Base) Enclosure Components, Rear View

Figure 1-5 shows a rear view of the JBOD (expansion) enclosure components.




Figure 1-5 JBOD (Expansion) Enclosure Components, Rear View

These components are discussed in the following sections:

- “Operators (Ops) Panel” on page 11
- “PSU/Cooling Module” on page 12
- “RAID LRC I/O Modules” on page 14
- “RAID Loopback LRC I/O Modules” on page 16
- “JBOD LRC I/O Module” on page 18
- “Drive Carrier Module” on page 19
- “Enclosure Bay Numbering” on page 22

Operators (Ops) Panel

The operators panel (ops panel) contains an enclosure services processor that monitors and controls the enclosure (see Figure 1-6). The ops panel contains LEDs which show the status for all modules, an audible alarm that indicates a fault state is present, a push-button alarm mute switch, and a thumb-wheel enclosure ID address range selector switch. When the 2 Gb TP9100 is powered on, the audible alarm sounds for one second, and the power-on LED illuminates.

Figure 1-6 shows the ops panel and identifies its components. For more information about the LEDs and configuration switches, see “ESI/Ops Panel LEDs and Switches” in Chapter 5.




Figure 1-6 Ops Panel

PSU/Cooling Module

Two power supply cooling modules (PSUs) are mounted in the rear of the enclosure (see Figure 1-7). These modules supply redundant cooling and power to the enclosure. Voltage operating ranges are nominally 115 V or 230 V AC, selected automatically.

Note: If a power supply fails, do not remove it from the enclosure until you have a replacement power supply. The cooling fans in the power supply will continue to operate even after the power supply fails. Removing a failed power supply and not replacing it immediately can result in thermal overload.




Figure 1-7 PSU/Cooling Module

Four LEDs mounted on the front panel of the PSU/cooling module (see Figure 1-8) indicate the status of the power supply and the fans. Module replacement must be completed within 10 minutes after removal of the failed module. For more information, see “Power Supply/Cooling Module LEDs” in Chapter 5.




Figure 1-8 PSU/Cooling Module Switches and LEDs

RAID LRC I/O Modules

The storage system enclosure includes two loop resiliency circuit (LRC) I/O modules with optional integrated RAID controllers. There are two RAID LRC I/O modules available: a dual-port version and a single-port version (see Figure 1-9 and Figure 1-10). The enclosure is available with or without RAID LRC I/O modules. An enclosure with one or two RAID LRC I/O modules is a RAID *base* enclosure. An added enclosure with JBOD LRC I/O modules is called an *expansion* enclosure, which must be cabled to a RAID LRC I/O enclosure. The base and expansion enclosures can be connected with the copper SFP cables that are included with the expansion enclosure or with optical SFP cables.

The FC-AL backplane in the enclosure incorporates two independent loops formed by port bypass circuits within the RAID LRC I/O modules. The RAID LRC I/O modules use FC-AL interfacing with the host computer system. Processors in the RAID LRC I/O modules communicate with the enclosure services interface (ESI) to devices on the backplane, PSU, LRC and ops panel, to monitor internal functions. These processors operate in a master/slave configuration to allow failover




Figure 1-9 Dual-port RAID LRC I/O Module




Figure 1-10 Single-port RAID LRC I/O Module

The RAID LRC I/O modules can address up to 64 disk drives. A maximum of two fully populated JBOD expansion enclosure can be cabled to a RAID base enclosure. The disk drives in each enclosure can be of different capacities, but all of the disk drives in an individual LUN must be of the same capacity.

For information about the LEDs on the rear of the RAID LRC I/O modules, see “RAID LRC I/O Module LEDs” in Chapter 5.

RAID Loopback LRC I/O Modules

A RAID loopback LRC I/O module may be installed in slot B to create an simplex RAID configuration. The loopback LRC I/O modules do not contain the FFX-2 circuitry and connect RAID LRC I/O module A to the B-side of the disk drives. These modules are sometimes referred to as a RAID wrap LRC I/O modules. There are two version of the the RAID loopback LRC I/O module available: a single-port version and a dual-port version. (See Figure 1-11 and Figure 1-12).

Note: The RAID LRC I/O modules in an enclosure must both be single-port controllers, or they must both be dual-port controllers. SGI does not support single-port and dual-port controllers in the same enclosure.




Figure 1-11 Single-port RAID Loopback LRC I/O Module




Figure 1-12 Dual-port RAID Loopback LRC I/O Module

JBOD LRC I/O Module

The JBOD LRC/IO module uses a Fibre Channel arbitrated loop (FC-AL) to interface with the host computer system. The FC-AL backplane incorporates two independent loops formed by port bypass circuits within the LRC I/O modules. Processors housed on the LRC modules provide enclosure management and interface to devices on the backplane, PSU/cooling module, and ops panel, to monitor internal functions. These processors operate in a master/slave configuration to allow failover.

Note: The JBOD LRC I/O module can address up to 96 disk drives; thus, six JBOD enclosures can be cabled together.

The enclosure may be configured with either one or two LRC I/O modules. If only one module is installed, an I/O blank module must be installed in the unused bay.






Figure 1-13 JBOD LRC I/O Module

For information about the LEDs on the rear of the JBOD LRC I/O module, see “RAID Loopback LRC I/O Module LEDs” in Chapter 5.

Drive Carrier Module

The disk drive carrier module consists of a hard disk drive mounted in a die-cast aluminum carrier. The carrier protects the disk drive from radio frequency interference, electromagnetic induction, and physical damage and provides a means for thermal conduction. For more information about drive carrier modules, see Chapter 6, “Installing and Replacing Drive Carrier Modules”.



Note: Ensure that the handle always opens from the left.

Figure 1-14 Drive Carrier Module

Drive Carrier Handle

The drive carrier module has a handle integrated into its front face. This handle cams the carrier into and out of the drive bay, holds the drive to the backplane connector, and prevents the unauthorized removal of the drive by means of an anti-tamper lock (see

Figure 1-15). For more information about operating the anti-tamper lock, see “Replacing a Drive Carrier Module” on page 98.




Figure 1-15 Anti-tamper Lock

For information about the drive carrier module LEDs, see “Drive Carrier Module LEDs” in Chapter 5.

Dummy Drive Carrier Modules

Dummy drive carrier modules must be installed in all unused drive bays. They are designed as integral drive module front caps with handles and must be fitted to all unused drive bays to maintain a balanced airflow. For information about replacing the dummy drive carrier modules, see “Replacing the Disk Drive Module” on page 100.




Figure 1-16 Dummy Drive Carrier Module

Enclosure Bay Numbering

This section contains information about enclosure bay numbering in the following sections:

- “Rackmount Enclosure Bay Numbering” on page 22
- “Tower Enclosure Bay Numbering” on page 24


Rackmount Enclosure Bay Numbering

The rackmount enclosure is 4 bays wide and 4 bays high, and the bays are numbered as follows:

- The disk drive bays, located in front, are numbered 1 to 4 from left to right and 1 to 4 from top to bottom. Drives in bays 1/1 and 4/4 are required for storage system management; these bays must always be occupied.
- The rear bays are numbered 1 to 5 from right to left.

The location of a disk drive module is identified by combining the column and row numbers (top and side numbers in Figure 1-20). For example, the disk drive in the upper left corner of the enclosure is disk 1-1. A module located in the rear of the enclosure is identified by its bay number. For example, the PSU/cooling module on the far left side of the enclosure is in bay 5.

Figure 1-17 shows the enclosure bay numbering convention and the location of modules in the rackmount enclosure.



Note: Each enclosure must have drives installed in position 1/1 and 4/4 to enable the SES monitor functions.

Figure 1-17 Rackmount Enclosure Bay Numbering and Module Locations


Tower Enclosure Bay Numbering

The tower enclosure is 4 bays wide by 4 bays high, and the bays are numbered as follows:

- The disk drive bays, located in front, are numbered 1 to 4 from right to left and 1 to 4 from top to bottom. Drives in bays 1/1 and 4/4 are required for storage system management; these bays must always be occupied.
- The rear bays are numbered 1 to 5 from top to bottom.

The location of a disk drive module is identified by combining the column and row numbers (top and side numbers in Figure 1-18). For example, the disk drive in the upper right corner of the enclosure is disk 1-1. A module located in the rear of the enclosure is identified by its bay number. For example, the PSU/cooling module on the bottom of the enclosure is in bay 5.

Figure 1-18 shows the correct positions of the modules and the enclosure bay numbering convention for the tower.



Note: Each enclosure must have drives installed in position 1/1 and 4/4 to enable the SES monitor functions.

Figure 1-18 Tower Enclosure Bay Numbering and Module Locations

Storage System Rack

This section contains information about the 2 Gb TP9100 storage system rack in the following sections:

- “Rack Structure” on page 26
- “Power Distribution Units (PDUs)” on page 29
- “Opening and Closing the Rear Rack Door” on page 31

Rack Structure

The 2 Gb TP9100 rack is 38U high and is divided into 12 bays. Eleven of these bays can house enclosures; the 2U bay at the top of the rack accommodates the SGI Fibre Channel Hub or one or more Fibre Channel switches. All eleven of the enclosure bays can be occupied by independent RAID enclosures or a combination of RAID enclosures and expansion enclosures. (Each RAID enclosure can support up to three expansion enclosures). Unoccupied bays must contain a 3U filler panel to provide proper airflow.

Caution: Equipment must be installed in the bays only as described above.

Figure 1-19 shows the front of a 2 Gb TP9100 rack with two enclosures installed.




Figure 1-19 Example of 2 Gb TP9100 Rack (Front View)

Figure 1-20 is a rear view of the 2 Gb TP9100 rack.




Figure 1-20 Example of 2 Gb TP9100 Rack (Rear View)

Power Distribution Units (PDUs)

The power distribution units (PDUs) mounted in the rear of the rack provide power to the enclosure and switch bays. The breakers on the PDUs also provide a power on/off point for the rack and enclosures. See Figure 1-21 for socket and breaker locations and functions. All sockets in the PDUs are rated at 200 to 240 VAC, with a maximum load per bank of outlet sockets of 8 A, and are labeled as such. The sockets are connected to equipment in the bays as follows:

- Socket 1 at the top of each PDU is for the 2U bay at the top of the rack that houses the SGI Fibre Channel hub or one or more Fibre Channel switches.
- Sockets 2 through 12 on each PDU are for the 11 3U bays, which accommodate 2 Gb TP9100 enclosures.



Warning: The power distribution units (PDUs) contain hazardous voltages. Do not open the PDUs under any circumstances.

Figure 1-21 shows the PDUs and describes the function of the sockets and breakers.




Figure 1-21 PDU Locations and Functions

Opening and Closing the Rear Rack Door

To open the rear rack door, follow these steps:

1. Locate the latch on the rear rack door.
2. Push up the top part of the latch, as shown in the second panel of Figure 1-22.




Figure 1-22 Opening the Rack Rear Door

3. Press the button as shown in the third panel of Figure 1-22. This action releases the door lever.
4. Pull the door lever up and to the right, to approximately the 2 o'clock position, as shown in the fourth panel of Figure 1-22. The door opens.

To close the door, lift the locking brace at the bottom. Then reverse the steps shown in Figure 1-22 to latch the door.

Storage System Tower

The tower (deskside) version of the storage system houses one RAID enclosure. The tower is mounted on four casters for easy movement. The enclosure in the tower system is rotated 90 degrees from the rackmount orientation.

Figure 1-23 shows the front of the tower.




Figure 1-23 Front View of Tower

Figure 1-24 shows a rear view of the tower.




Figure 1-24 Rear View of Tower

The tower storage system receives power from standard electrical sockets. Figure 1-25 shows the power cords attached to the rear of the tower.




Figure 1-25 Tower Storage System Power Cords

The tower enclosure can be adapted for rackmounting; contact your service provider for more information.

Connecting to a Host and Powering On and Off

This chapter explains cabling the storage system and powering it on and off in the following sections:

- “Connecting to a Host” on page 35
- “Grounding Issues” on page 37
- “Connecting the Power Cords and Powering On the 2 Gb TP9100 Tower” on page 37
- “Connecting the Power Cords and Powering On the 2 Gb TP9100 Rack” on page 40
- “Powering Off” on page 44

Note: For instructions on opening the rear door of the rack, see “Opening and Closing the Rear Rack Door” on page 31.

Connecting to a Host

The 2 Gb TP9100 supports only Fibre Channel optical connectivity to the front-end host or switch. Small form-factor pluggables (SFPs) provide the optical connection to the LRC I/O module.

Note: Copper connections to hosts and/or switches are not supported for either RAID or JBOD enclosures.

A pair of copper cables is packaged with 2 Gb TP9100 JBOD enclosures. These cables are manufactured with copper SFPs on each end of the cable. Use the copper cable/SFP assembly to connect JBOD enclosures used either as capacity expansion enclosures for a RAID system or to connect cascaded JBOD enclosures. When the JBOD enclosure is used

as a host-attached JBOD enclosure, the copper cable/SFP assembly can be replaced with optical SFPs and optical cables.

To connect the storage system to a host, insert an optical cable (with SFP) into the connector labeled “Host 0.” Connect the other end of the optical cable to the FC-AL port on the host.

In addition to cabling directly to an HBA in a host, you can connect the storage system to an SGI Fibre Channel 8-port or 16-port switch (using an optical cable and an optical GBIC). See Table A-10 on page 108 for information on these cables.

Note: The I/O module current limit is 1.5 A.

The host ports of the RAID controller can be connected to a switched-fabric Fibre Channel Arbitrated Loop (FC-AL) or directly to a server in a point-to-point configuration. A FC-AL provides shared bandwidth among the attached nodes; as additional nodes are added to a loop, the bandwidth available to each node decreases. Fibre Channel switched fabrics are interconnected with switches that increase bandwidth as nodes and switch ports are added to the system. The bandwidth available to each node in a switched fabric always remains constant. Unlike previous versions of the TP9100, which only support FC-AL topologies, the FFX-2 RAID controller host ports of the 2 Gb TP9100 can implement the behavior of a N_Port when connected in a point-to-point topology with a server, or when connected to a F_Port on a switch.

In FC_AL topologies, the FFX-2 RAID controller uses NL_Port behavior to connect to FL_Ports on hosts or switches. After a 2Gb TP9100 boots up, it initiates a log-in sequence and automatically determines which topology and the protocol should be used, as dictated by the environment. The topology and protocol are determined by the preferences of the connecting devices and the internal topology of the 2Gb/s TP9100. For example, if the system is in multi-target id mode (MTID) then it will connect as an FC_AL device to ensure that bandwidth is shared equally across the loop. If the system is in multi-port mode it will attempt to connect as a point-to-point topology in order to provide the largest amount of bandwidth possible to each host. When the system is in MTID mode, it can also connect as a FC-AL device depending on the other devices are connected.

Note: Host FC-AL topologies on both the 1 Gb TP9100 with the FFX RAID Controller and 2 Gb TP9100 with FFX-2 RAID controller support fabric.

This transparent flexibility protects investments in existing infrastructure, enhances storage area network (SAN) robustness, and simplifies SAN configuration management.

The 2Gb TP9100 with FfX-2 RAID controller features a host side hub function which is configured by the switches on the ops panel. When the system is in hub mode, FC-AL is the only supported topology. If the system is in point to point mode because of the host hub functionality, the system must be power cycled before connecting to a HBA or switch in an arbitrated loop topology.

For more information on configurations, see “Disk Topologies” on page 65.

Grounding Issues

Each chassis—storage or host—must be well-grounded through its power connector. If you have any doubts about the quality of the ground connection, consult with a qualified electrician.

The branch circuit wiring should include an insulated grounding conductor that is identical in size, insulation material, and thickness to the earthed and unearthed branch-circuit supply conductors.

The grounding conductor should be green, with or without one or more yellow stripes. This grounding or earthing conductor should be connected to earth at the service equipment or, if supplied by a separately derived system, at the supply transformer or motor-generator set. The power receptacles in the vicinity of the systems should all be of an earthing type, and the grounding or earthing conductors serving these receptacles should be connected to earth at the service equipment.



Warning: The rack power distribution units (PDUs) must be connected only to power sources that have a safe electrical earth connection. For safety reasons, this earth connection must be in place at all times.

Connecting the Power Cords and Powering On the 2 Gb TP9100 Tower

The tower requires 115-220 V (autoranging), and is shipped with two power cords, shown in Figure 2-1.




Figure 2-1 Power Cords for the Tower



Caution: Use the power cords supplied with the storage system or power cords that match the specification shown in Table A-7 on page 106.

Geography-specific power cords are available from SGI.

To install the power cords and power on the storage system, follow these steps:

1. Ensure that all modules are firmly seated in the correct bays and that blank plates are fitted in any empty bays.
2. Ensure that the ambient temperature is within the specified operating range of 10 °C to 40 °C (50 °F to 104 °F). If any drives have been recently installed, allow them to acclimatize before operating the system.
3. Connect an AC power cord to each PSU/cooling module. To ensure that your system is properly grounded, test for continuity between the ground pins of the power plugs and a metal component of the enclosure frame.



Caution: Some electrical circuits could be damaged if external signal cables are present during the grounding checks. Do not connect any signal cables to the enclosure until you have completed the ground test

4. Connect the AC power cords to properly grounded outlets.
5. Turn the power switch on each PSU/cooling module to the “on” position (“I”=on, “O”=off).

Checking AC Power and Storage System Status for the Tower

The “Power on” LED on the ESI/ops panel (see Figure 2-2) turns green if AC power is present.




Figure 2-2 ESI/Ops Panel LEDs and Switches

At power-on, check ESI/ops panel LEDs for system status. Under normal conditions, the “Power on” LED should illuminate constant green. If a problem is detected, the ESI processor in the operator panel will illuminate the “System/ESI fault” LED in amber. See “Solving Initial Startup Problems” on page 76 and “Using Storage System LEDs for Troubleshooting” on page 78.

Other modules in the storage system also have LEDs, which are described in “Using Storage System LEDs for Troubleshooting” on page 78.

Connecting the Power Cords and Powering On the 2 Gb TP9100 Rack

The rack requires 220 V and is shipped with a country-specific power cord for each power distribution unit (PDU) that the rack contains. Each power supply of each enclosure in the rack is cabled to the rack PDU on the appropriate side; Figure 2-3 shows an example.

The PDU has double-pole circuit breakers and can be connected to either a phase-to-neutral power source or to a phase-to-phase power source.



Warning: The power distribution units (PDUs) at the sides in the rear of the rack contain hazardous voltages. Do not open the PDUs under any circumstances.

A qualified SGI system support engineer (SSE) will set up the rack and cable it to power. The information in this section is provided for reference and safety reasons only.

Additional rackmountable enclosures that you order after your rack is set up are shipped with two IEC 320 power cords for cabling to the rack PDUs. Qualified SGI SSEs will install and cable the enclosures in the rack.



Warning: The rack PDUs must be connected only to power sources that have a safe electrical earth connection. For safety reasons, this earth connection must be in place at all times.




Figure 2-3 Rack Power Cabling

Checking Grounding for the Rack

If necessary, follow these steps to ensure that a safe grounding system is provided:

1. Note the information in “Grounding Issues” on page 37.
2. For the grounding check, ensure that the rack PDU power cords are not plugged in to a power source.



Caution: Some electrical circuits could be damaged if external signal cables or power control cables are present during the grounding checks.

3. Ensure that each power supply/cooling module of each enclosure in the rack is cabled to a PDU on the appropriate side of the rack.
4. Check for continuity between the earth pin of the enclosure power cords and any exposed metal surface of the enclosures in the rack.
5. Check the earth connection of the power source.



Warning: The rack PDUs must be connected only to power sources that have a safe electrical earth connection. For safety reasons, this earth connection must be in place at all times.

Powering On the Rack

When the rack is set up, it is usually powered on and ready to be operated. If it has been turned off, follow these steps to power it back on:

1. Ensure that the ambient temperature is within the specified operating range of 10 °C to 40 °C (50 °F to 104 °F). If drives have been recently installed, make sure that they have had time to acclimatize before operating them.
2. Ensure that each power supply/cooling module of each enclosure in the rack is cabled to a PDU on the appropriate side of the rack.
3. If they have not already been connected, connect each PDU power cord to a power source. The PDU power cords can be routed through an opening at the top or the bottom of the rack. See Figure 2-3 on page 41.



Warning: The rack PDUs must be connected only to power sources that have a safe electrical earth connection. For safety reasons this earth connection must be in place at all times.

Be careful not to touch the pins on the PDU plug when you insert it into a power source.

4. Press the rack breaker switch at the bottom of each PDU so that the word **ON** shows.
5. Ensure that all of the socket group breakers on each PDU are turned on (position “**I**”=on, “**O**”=off). These breakers are identified by illuminated green buttons.
6. Move the power switch on the rear of each PSU/cooling module (2 per enclosure) to the “On” position (position “**I**”=on, “**O**”=off).

Checking AC Power and System Status for the Rack

When you power on the system, the “Power on” LED on each ESI/ops panel (see Figure 2-4) in each enclosure you are operating should illuminate. If it does not, check that the power supply/cooling modules in the enclosure are correctly cabled to the rack PDUs and turned on.




Figure 2-4 Rackmount Enclosure ESI/Ops Panel Indicators and Switches

At power-on, check the ESI/ops panel LEDs for system status. Under normal conditions, the “Power on” LED should illuminate constant green. If a problem is detected, the ESI processor in the ops panel will illuminate the “System/ESI fault” LED in amber. See “Solving Initial Startup Problems” on page 76 and “Using Storage System LEDs for Troubleshooting” on page 78.

Other modules in the storage system also have LEDs, which are described in “Using Storage System LEDs for Troubleshooting” on page 78.

Powering Off

This section covers powering off the 2 Gb TP9100 in the following sections:

- “Powering Off the 2 Gb TP9100 Rack” on page 45
- “Powering Off the 2 Gb TP9100 Tower or a Single Enclosure” on page 46

Powering Off the 2 Gb TP9100 Rack

Besides the main breaker switch at the bottom of each PDU, the rack PDUs have breaker switches at each 12U of space so that you can power off the enclosures in groups of four and leave the others powered on. Figure 2-3 shows their locations.

To power off the entire rack, follow these steps:

1. Ensure that users are logged off of the affected systems.
2. Move the power switch on the rear of each PSU/cooling module (2 per enclosure) to the "Off" position (position "I"=on, "O"=off).
3. Turn off all of the socket group breakers on each PDU (position "I"=on, "O"=off). These breakers can be identified by the illuminated green switches.
4. Push down the main breaker switch at the bottom of each PDU so that the word **OFF** shows.
5. If appropriate, disconnect the PDU power cords from the power sources.

Powering Off the 2 Gb TP9100 Tower or a Single Enclosure

Besides the main breaker switch at the bottom of each PDU, the rack PDUs have breaker switches at each 12U of space so that you can power off three enclosures and leave others powered on.

To power off a single enclosure or tower storage system, follow these steps:

1. Ensure that users are logged off of the affected systems.
2. Move the power switch on the rear of each PSU/cooling module to the “Off” position (position “I”=on, “O”=off).
3. If appropriate, disconnect the PDU power cords from the power sources.

Features of the RAID Controller

This chapter describes features and operation of the RAID controller in the following sections:

- “Enclosure Services Interface (ESI) and Disk Drive Control” on page 47
- “Configuration on Disk (COD)” on page 48
- “Drive Roaming” on page 49
- “Data Caching” on page 50
- “RAID Disk Topologies” on page 52

Enclosure Services Interface (ESI) and Disk Drive Control

Both the JOBD and RAID LRC I/O modules use enclosure services interface (ESI) commands to manage the physical storage system. ESI provides support for disk drives, power supply, temperature, door lock, alarms, and the controller electronics for the enclosure services. The storage system ESI/ops panel firmware includes SES.

Note: These services are performed by drives installed in bays 1/1 and 4/4; these drives must be present for the system to function. See Figure 1-17 on page 23 for diagrams of their location.

ESI is accessed through an enclosure services device, which is included in the ESI/ops module. SCSI commands are sent to a direct access storage device (namely, the drives in bays 1/4 and 4/4), and are passed through to the SES device.

During controller initialization, each device attached to each loop is interrogated, and the inquiry data is stored in controller RAM. If ESI devices are detected, the ESI process is started. The ESI process polls and updates the following data:

- Disk drive insertion status

- Power supply status
- Cooling element status
- Storage system temperature

The LEDs on the ESI/ops panel show the status of these components.

Configuration on Disk (COD)

Configuration on disk (COD) retains the latest version of the saved configuration at a reserved location on every physical drive. The RAID Controller in the 2Gb TP9100 (Mylex FFX-2) uses COD version 2.1. Previous versions of the TP9100 use COD version 1.0.

Controller firmware versions prior to 7.0 use COD 1.0 format. Firmware versions 7.0 and later use COD 2.1 format. FFX-2 RAID controller support started on version 8.0 firmware.

The COD information stored on each drive is composed of the following:

- Device definition, which contains the following information.
 - The logical device definition/structure for those logical devices dependent on this physical device. This information should be the same for all physical devices associated with the defined logical device.
 - Any physical device information pertaining to this physical device that is different for different physical devices even though they may be part of the same logical device definition.
 - Data backup for data migration. This area also includes required information for the Background initialization feature.
- User device name information and host software configuration parameters. This information is defined by the user and should be the same on all physical drives that are associated with the defined logical drive.
- COD 2.1 locking mechanism. This feature is designed to provide a locking mechanism for multiple controller systems. If any of the controllers is allowed to update COD information independently of the other controllers, this feature allows the controller to lock the COD information for write access before updating the that drive. This feature prevents multiple controllers from updating the COD at the same time.

COD plays a significant role during the power-on sequence after a controller is replaced. The replacement controller tests the validity of any configuration currently present in its NVRAM. Then, it test the validity of the COD information on all disk drives in the storage system. The final configuration is determined by the following rules:

1. The controller will use the most recent COD information available, no matter where it is stored. The most recent COD information is updated to all configured drives. Unconfigured drives are not updated; all COD information on these drives is set to zero.
2. If all of the COD information has an identical timestamp, the controller will use the COD information stored in its NVRAM.



Caution: Any existing COD on a disk drive that is inserted after the controller has started (STARTUP COMPLETE) will be overwritten.



Caution: Mixing controllers or disk drives from systems running different versions of firmware presents special situations that may affect data integrity. If a new disk drive containing configuration data is added to an existing system while power is off, the controller may incorrectly adopt the configuration data from the new drive. This may destroy the existing valid configuration and result in potential loss of data. Always add drives with the power supplied to the system to avoid potential loss of data.

Drive Roaming

Drive roaming allows disk drives to be moved to other channel/target ID locations while the system is powered down. Drive roaming allows for easier disassembly and assembly of systems, and potential performance enhancement by optimizing channel usage.

Drive roaming uses the Configuration on Disk (COD) information stored on the physical disk drive. When the system restarts, the controller generates a table that contains the current location of each disk drive and the location of each drive when the system was powered down. This table is used to remap the physical disk drives into their proper location in the system drive. This feature is designed for use within one system environment, for example, a single system or a cluster of systems sharing a simplex or dual-active controller configuration. Foreign disk drives containing valid COD information from other systems must not be introduced into a system. If the COD

information on a replacement disk drive is questionable or invalid, the disk drive will be labeled unconfigured offline or dead.

If a drive fails in a RAID level that uses a hot spare, drive roaming allows the controller to keep track of the new hot spare, which is the replacement for the failed drive.



Caution: Mixing controllers or disk drives from systems running different versions of firmware presents special situations that may affect data integrity. If a new disk drive containing configuration data is added to an existing system while power is off, the controller may incorrectly adopt the configuration data from the new drive. This may destroy the existing valid configuration and result in potential loss of data. Always add drives with the power supplied to the system to avoid potential loss of data.

Data Caching

RAID controllers can be operated with write cache enabled or disabled. This section describes the modes in the following subsections:

- “Write Cache Enabled (Write-back Cache Mode)” on page 50
- “Write Cache Disabled (Write-through or Conservative Cache Mode)” on page 51

Write caching is set independently for each system drive in the system management software.

Write Cache Enabled (Write-back Cache Mode)

If write cache is enabled (write-back cache mode), a write completion status is issued to the host initiator when the data is stored in the controller’s cache, but before the data is transferred to the disk drives. In dual-active controller configurations with write cache enabled, the write data is always copied to the cache of the second controller before completion status is issued to the host initiator.

Enabling write cache enhances performance significantly for data write operations; there is no effect on read performance. However, in this mode a write complete message is sent to the host system as soon as data is stored in the controller cache; some delay may occur

before this data is written to disk. During this interval there is risk of data loss in the following situations:

- If only one controller is present and this controller fails.
- If power to the controller is lost and its internal battery fails or is discharged.

Write Cache Disabled (Write-through or Conservative Cache Mode)

If write cache is disabled (write-through data caching is enabled), write data is transferred to the disk drives before completion status is issued to the host initiator. In this mode, system drives configured with the write cache enabled policy are treated as though they were configured with write cache disabled, and the cache is flushed.

Disabling write cache (enabling write-through or conservative mode) provides a higher level of data protection after a critical storage system component has failed. When the condition disabling write cache is resolved, the system drives are converted to their original settings.

Conditions that disable write cache are as follows:

- The Enable Conservative Cache controller parameter is enabled in TPM for a dual-active controller configuration, and a controller failure has occurred.
- A power supply has failed (not simply that a power supply is not present).

In this case the SES puts the RAID into conservative cache mode. This condition also triggers the audible alarm.

- An out-of-limit temperature condition exists.

In this case the SES puts the RAID into conservative cache mode. This condition also triggers the audible alarm.

- The controller receives an indication of an AC failure.

To protect against single-controller failure, certain releases of the storage system support dual controllers. To protect against power loss, an internal battery in the controller module maintains the data for up to 72 hours.

RAID Disk Topologies

The 2 Gb TP9100 RAID enclosure can be configured with any of the following topologies:

- “Simplex Single-port RAID Topology” on page 53
- “Duplex Single-Port RAID Topology” on page 54
- “Simplex Dual-Port RAID Topology” on page 55
- “Duplex Dual-Port RAID Topology” on page 56
- “Dual-port Duplex Two-Host RAID Configuration” on page 57
- “Dual-Port Duplex RAID Configuration” on page 58

Simplex Single-port RAID Topology

Figure 3-1 illustrates a simplex single-port RAID configuration that uses a single host. This configuration:

Supports transfer speeds up to 200 MB/s

Does not support failover capabilities




Figure 3-1 Simplex Single-port RAID Topology

Duplex Single-Port RAID Topology

Figure 3-2 illustrates a duplex single-port RAID configuration. This configuration:

- Supports transfer speeds up to 400 MB/s
- Supports failover capabilities
- Supports SGI FailSafe high-availability solutions




Figure 3-2 Duplex Single-port RAID Topology

Simplex Dual-Port RAID Topology

Figure 3-3 illustrates a simplex dual-port RAID configuration using two hosts. This configuration:

- Supports transfer speeds up to 400 MB/s
- Supports failover capabilities
- Supports SGI FailSafe high-availability solutions

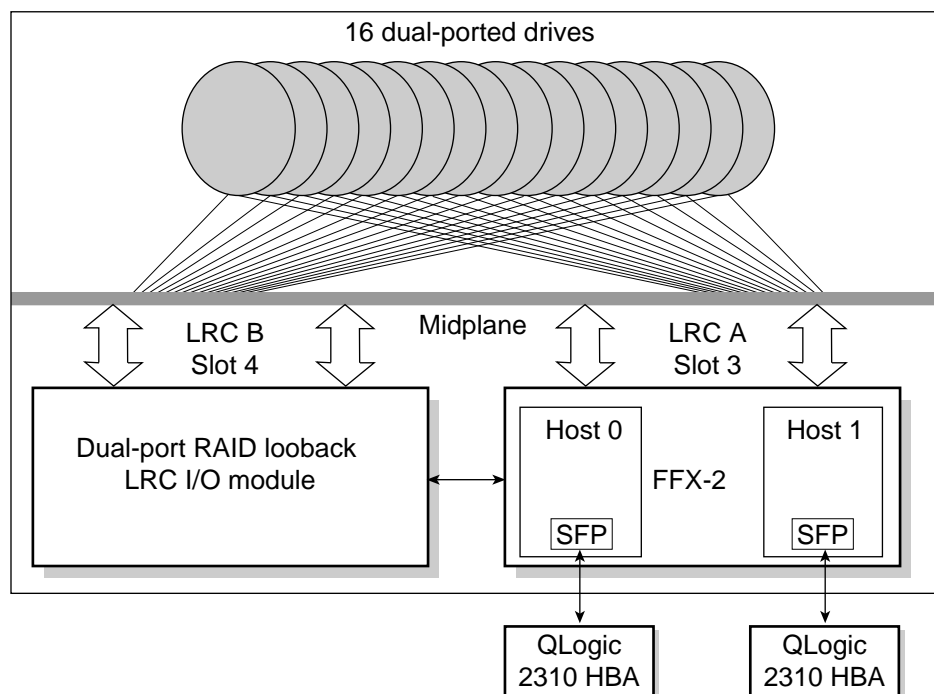


Figure 3-3 Simplex Dual-port Dual-host RAID Topology

Duplex Dual-Port RAID Topology

Figure 3-4 illustrates a duplex dual-port RAID configuration using two hosts and two controllers. This configuration:

- Supports transfer speeds up to 400 MB/s
- Supports failover capabilities
- Supports SGI FailSafe high-availability solutions




Figure 3-4 Duplex Dual-port RAID Configuration



Caution: If two independent systems 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 one system only.

Dual-port Duplex Two-Host RAID Configuration

Figure 3-5 illustrates a dual-port, duplex, dual-path RAID configuration that uses two hosts. This configuration:

- Supports transfer speeds up to 400 MB/s
- Supports failover capabilities
- Supports SGI FailSafe high-availability solutions




Figure 3-5 Dual-port Dual-path Attached Duplex RAID Topology



Caution: If two independent systems 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 one system only.

Dual-Port Duplex RAID Configuration

Figure 3-6 illustrates a dual-port quad-path attached duplex RAID configuration. This configuration supports the following features:

- Transfer speeds up to 400 MB/s
- Failover capabilities
- SGI FailSafe high-availability solution




Figure 3-6 Dual-port Quad-path Duplex RAID Topology



Caution: If two independent systems 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 one system only.

Using the RAID Controller

This chapter explains the operation of the RAID controller in the following sections:

- “Software Tools for the Controller” on page 59
- “RAID Levels” on page 60
- “CAP Strategy for Selecting a RAID Level” on page 61
- “Disk Topologies” on page 65
- “System Drives” on page 68
- “Drive State Reporting” on page 69
- “Automatic Rebuild” on page 71

Software Tools for the Controller

Two software components allow you to manage the RAID controllers: the RAID controller firmware and the RAID management software.

RAID firmware has the following characteristics:

- Resides on RAID controller (FFX-2) in the LRC I/O module.
- Controls the low-level hardware functions.
- Controls RAID functionality.
- Can be upgraded or “flashed” in the field by trained service personnel.

RAID management software (TPM) has the following characteristics:

- Resides on the host system.
- Use in-band management to interface firmware.
- Provides a graphical user interface (GUI).

RAID Levels

RAID stands for “redundant array of inexpensive disks.” In a RAID storage system, multiple disk drives are grouped into **arrays**. Each array is configured as a single system drive consisting of one or more disk drives.

Correct installation of the disk array and the controller requires a proper understanding of RAID technology and concepts. The controllers implement several versions of the Berkeley RAID technology, as summarized in Table 4-1.

Note: Although JBOD (“just a bunch of disks”) is not strictly a RAID level, it is included at various points in this discussion for comparison to RAID levels. It is sometimes referred to as RAID 7.

Table 4-1 Supported RAID Levels

RAID Level	Description	Minimum Drives	Maximum Drives	Fault-tolerant?
0	Block striping is provided, which yields higher performance than is possible with individual disk drives. No redundancy is provided.	2	16	No
1	Disk drives are paired and mirrored. All data is duplicated 100% on an equivalent disk drive.	2	2	Yes
3	Data is striped across several physical disk drives. Parity protection is used for data redundancy. This level provides a larger bandwidth for applications that process large files.	3	16	Yes
5	Data and parity information is striped across all physical disk drives. Parity protection is used for data redundancy.	3	16	Yes
0+1 (6)	Combination of RAID levels 0 and 1. Data is striped across several physical disk drives. This level provides redundancy through mirroring.	4	16	Yes
JBOD (7)	Each disk drive is operated independently like a normal disk drive, or multiple disk drives can be spanned and seen as a single large drive. This level does not provide data redundancy.	1	1	No

You must select an appropriate RAID level when you define or create system drives. This decision is based on how you prioritize the following:

- Disk capacity utilization (number of disk drives)
- Data redundancy (fault tolerance)
- Disk performance

The controllers make the RAID implementation and the disk drives' physical configuration transparent to the host operating system. This transparency means that the host operating logical drivers and software utilities are unchanged, regardless of the RAID level selected.

Although a system drive may have only one RAID level, RAID levels can be mixed within a drive pack (LUN), as illustrated in Figure 4-1.




Figure 4-1 Example of RAID Levels within a Drive Pack (LUN)

In Figure 4-1, the smaller system drive (B0) is assigned a RAID 5 level of operation, while the larger system drive (B1) is assigned a RAID 0+1 level of operation.

Remember that different RAID levels exhibit different performance characteristics for a particular application or environment. The controller affords complete versatility in this regard by allowing multiple RAID levels to be assigned to a drive pack.

Drives are fault-tolerant when you use a RAID level providing redundancy. In the simplex configuration, however, if the controller or host bus adapter fails, the data is not accessible until the failure is corrected.

CAP Strategy for Selecting a RAID Level

Capacity, availability, and performance are three benefits, collectively known as CAP, that should characterize your expectations of the disk array subsystem.

It is impossible to configure an array optimizing all of these characteristics; that is a limitation of the technology. For example, maximum capacity and maximum availability cannot exist in a single array. Some of the disk drives must be used for redundancy, which reduces capacity. Similarly, configuring a single array for both maximum availability and maximum performance is not an option.

The best approach is to prioritize requirements. Decide which benefit is most important for the operating environment. The controller in the 2 Gb TP9100 storage system is versatile enough to offer any of these preferences, either singly or in the most favorable combination possible.

The three benefits are further explained in these subsections:

- “Configuring for Maximum Capacity” on page 62
- “Configuring for Maximum Availability” on page 63
- “Configuring for Maximum Performance” on page 65

Configuring for Maximum Capacity

Table 4-2 shows the relationship between RAID levels and effective capacities offered for the quantity X disk drives of N capacity. As an example, it provides computed capacities for six 2-GB disk drives.

Table 4-2 RAID Level Maximum Capacity

RAID Level	Effective Capacity	Example: Capacity in GB
0	$X*N$	$6*2 = 12$
1	$(X*N)/2$	$6*2/2 = 6$
3	$(X-1)*N$	$(6-1)*2 = 10$
5	$(X-1)*N$	$(6-1)*2 = 10$
0+1	$(X*N)/2$	$(6*2)/2 = 6$
JBOD	$X*N$	$6*2 = 12$

The greatest capacities are provided by RAID 0 and JBOD, with the entire capacity of all disk drives being used. Unfortunately, with these two solutions, there is no fault

tolerance. RAID 3 and RAID 5 give the next best capacity, followed by RAID 1 and RAID 0+1.

Configuring for Maximum Availability

Table 4-3 presents definitions of array operating conditions.

Table 4-3 Array Operating Conditions

Array Condition	Meaning
Normal (online)	The array is operating in a fault-tolerant mode, and can sustain a disk drive failure without data loss.
Critical	The array is functioning and all data is available, but the array cannot sustain a second disk drive failure without potential data loss.
Degraded	The array is functioning and all data is available, but the array cannot sustain a second disk drive failure without potential data loss. Additionally, a reconstruction or rebuild operation is occurring, reducing the performance of the array. The rebuild operation takes the array from a critical condition to a normal condition.
Offline	The array is not functioning. If the array is configured with a redundant RAID level, two or more of its member disk drives are not online. If the array is configured as a RAID 0 or JBOD, one or more of its member disk drives are not online.
Not fault-tolerant	No fault-tolerant RAID levels have been configured for any of the disk drives in the array.

You can achieve an additional measure of fault tolerance (or improved availability) with a hot spare, or standby disk drive. This disk drive is powered on but idle during normal array operation. If a failure occurs on a disk drive in a fault-tolerant set, the hot spare takes over for the failed disk drive, and the array continues to function in a fully fault-tolerant mode after it completes its automatic rebuild cycle. Thus the array can suffer a second disk drive failure after rebuild and continue to function before any disk drives are replaced.

Controller Cache and Availability

The RAID controller has a write cache of 512 MB. This physical memory is used to increase the performance of data retrieval and storage operations.

The controller can report to the operating system that a write is complete as soon as the controller receives the data. Enabling write cache (write-back cache) improves performance, but exposes the data to loss if a system crash or power failure occurs before the data in the cache is written to disk. To prevent data loss, use an uninterruptable power supply (UPS).

In systems using dual-active RAID controllers, data is copied to the cache of the partner controller before the write complete is reported to the host initiator. During the time the data is being written to the partner controller, the system is exposed to possible data loss if a system crash or power failure occurs. Again, a UPS is recommended to preserve data integrity.



Caution: No UPS has been tested, qualified, or approved by SGI.

RAID Levels and Availability

Table 4-4 summarizes RAID levels offered by the RAID controller and the advantages (and disadvantages) of the RAID levels as they apply to availability.

Table 4-4 RAID Levels and Availability

RAID Level	Fault Tolerance Type	Availability
0	None	Data is striped across a set of multiple disk drives. If a disk drive in the set ceases to function, all data contained on the set of disk drives is lost. (This configuration is not recommended if fault tolerance is needed.)
1	Mirrored	Data is written to one disk drive, and then the same data is written to another disk drive. If either disk drive fails, the other one in the pair is automatically used to store and retrieve the data.
3 and 5	Striped	Data and parity are striped across a set of at least three disk drives. If any fail, the data (or parity) information from the failed disk drive is computed from the information on the remaining disk drives.

Table 4-4 RAID Levels and Availability **(continued)**

RAID Level	Fault Tolerance Type	Availability
0+1	Mirrored and striped	Data is striped across multiple disk drives, and written to a mirrored set of disk drives.
JBOD	None	This configuration offers no redundancy and is not recommended for applications requiring fault tolerance.

Configuring for Maximum Performance

Table 4-5 presents the relative performance advantages of each RAID level.

Table 4-5 RAID Levels and Performance

RAID Level	Access Profile Characteristics
0	Excellent for all types of I/O activity
1	Excellent for write-intensive applications
3	Excellent for sequential or random reads and sequential writes
5	Excellent for sequential or random reads and sequential writes
0+1	Excellent for write-intensive applications
JBOD	Mimics normal, individual disk drive performance characteristics

Disk Topologies

After you have determined the RAID level to use, determine the loop configuration. Note the following:

- The largest RAID group that can be created is 15+1 (16 drives).
- For a tower, the maximum SGI supported configuration is 16 drives total, those in the system itself; no expansion to another enclosure or tower is possible.
- For a RAID enclosure and two expansion enclosures, the maximum release 5 configuration is 32 drives. A maximum of 16 system drives can be created (see “System Drives” on page 68 for more information).

The disk drive modules are dual-ported. A RAID controller sees 16 to 32 drives on each loop (A and B), because it finds both ports of each drive. Via the I/O modules, it alternates allocation of the drives between channels, so that the drive addresses are available for failover. At startup, half the drives are on channel 0 via their A port and the other half are on channel 1 via their B port; each I/O module controls a separate loop of half the drives. Figure 4-2 diagrams this arrangement for the tower.

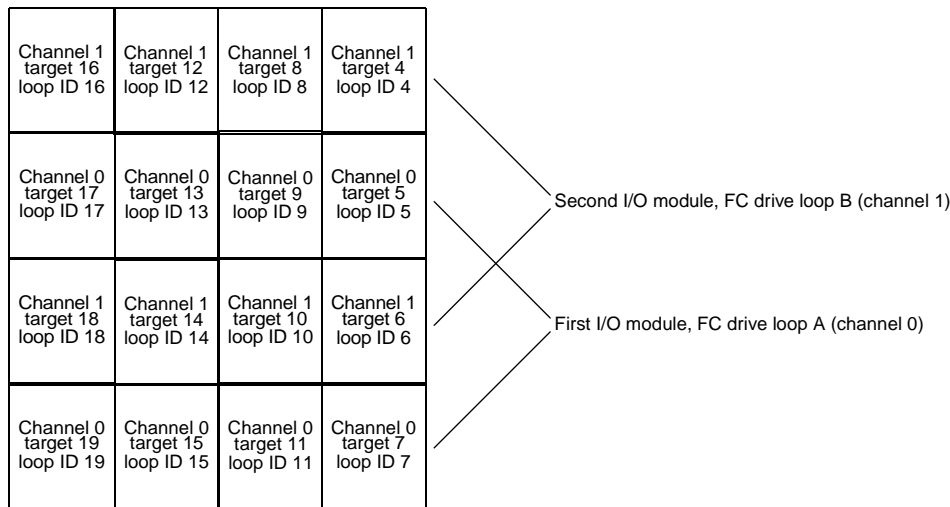


Figure 4-2 Tower I/O Modules, Channels, and Loops

Figure 4-3 diagrams disk addressing for a rackmount RAID system with a full complement of disks in three enclosures.

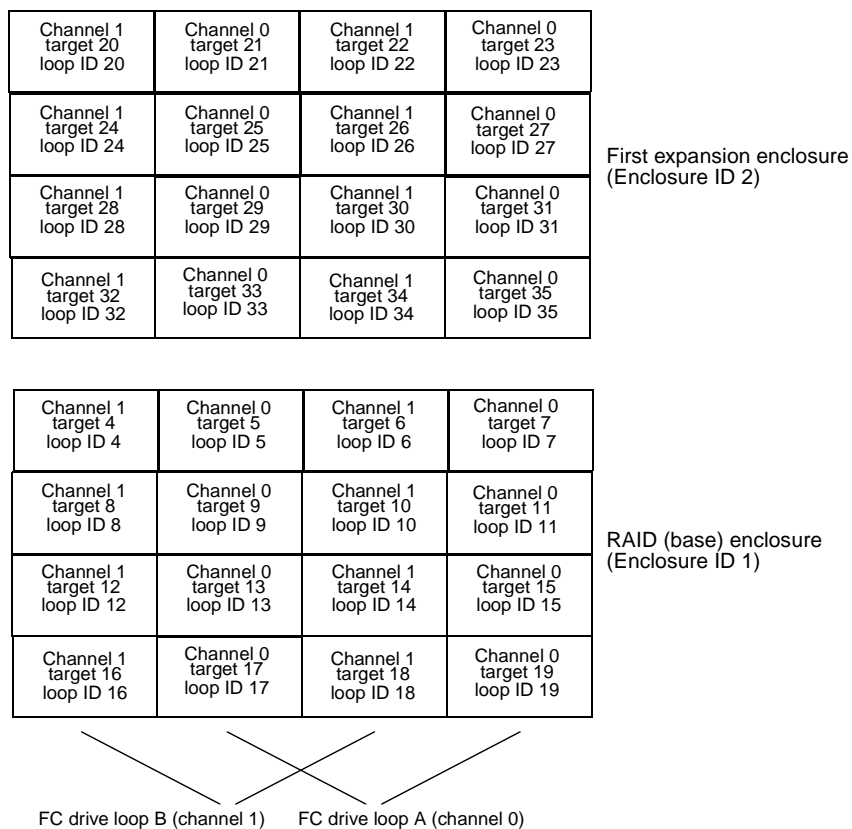


Figure 4-3 Rackmount Enclosure I/O Modules, Channels, and Loops (Front View)

However, you can use TPM to reassign target drives in accordance with your CAP strategy to channels 0 and 1. Check and confirm if the controller parameters need to be

modified for the intended application; see the documentation for the management software included with the storage system for information on controller parameters.

Note: Changes to the controller parameter settings take effect after the controller is rebooted.

System Drives

System drives are the logical devices 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. This section discusses system drives in these subsections:

- “System Drive Properties” on page 68
- “System Drive Affinity and Programmable LUN Mapping” on page 69

System Drive Properties

System drives have the following properties:

- The minimum size of a system drive is 8 MB; the maximum size is 2 TB.
- Up to 16 system drives can be created.
- Each system drive has a RAID level that 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); see “Data Caching” on page 50 for an explanation of this feature.
- Each system drive has its own LUN affinity. This capability is further discussed in “System Drive Affinity and Programmable LUN Mapping” on page 69.
- More than one system drive can be defined on a single drive pack (LUN).

System Drive Affinity and Programmable LUN Mapping

System drive affinity and programmable LUN mapping are configuration features that work together to define how the host accesses the available storage space.

System drive affinity allows system drives to be assigned to any combination of controller and host ports as follows:

- Configurations with one RAID controller that has two host ports (through a switch, for example) can use system drive affinity to define affinity of each system drive to one or both host ports.
- System drives that are not owned by a controller/host port are not accessible.

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



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 one system only.

Programmable LUN mapping lets you assign any LUN ID (even multiple LUN IDs) to any system drive on each port, or configure system drive assignments without specifying the LUN, defaulting to the current mapping algorithm. System drives with the “all” affinity are mapped to a LUN ID on every controller/host port.

Drive State Reporting

The RAID controller sends information about the status of each physical disk drive to the array management software. The controller records the operational state of each drive and a list of available target ID addresses. The controller determines which drives are present and what target IDs are available. Then, it determines the status of the drives that are present. If the disk drive is present, the location of the disk drive is considered configured and the operational state of the disk drive is then determined. If the controller determines the disk drive at the available target ID location is absent, the location of the

disk drive is considered unconfigured and the operational state is marked unconfigured, offline, or dead.

If a configured disk drive is removed or fails, and a new disk drive replaces the failed disk drive at the same location, the new disk drive is set to online spare. This allows the automatic rebuild operation to function with replaced drives. When a disk drive is inserted into the system, the controller recognizes that the drive has been replaced. If a configured disk drive fails and the controller loses power or is reset, the disk drive remains offline.

Unconfigured disk drives can be removed and the device state will remain unconfigured. New disk drives added to the system are considered unconfigured until used in a new configuration.

Unconfigured disk drive fault lights (LEDs) are disabled and any insertion, removal, or errors related to these unconfigured devices do not result in fault light activity or error message generation.

If the RAID controller is running firmware version 7.0 or later, COD information is written to all configured drives. Unconfigured drives are not updated; their COD information is set to all zeros.

Table 4-6 describes possible physical disk drive states. This information applies only to physical disk drives, not to system drives.

Table 4-6 Physical Disk Drive States

State	Description
Online optimal	The disk drive is powered on, has been defined as a member of a drive pack, and is operating properly.
Online spare	The disk drive is powered on, is able to operate properly, and has been defined as a standby or hot spare.
Offline failed or Unconfigured offline	<p>The disk drive is one of the following:</p> <ul style="list-style-type: none"> • Not present • Present, but not powered on • A newly inserted replacement drive • Marked as offline by the controller due to operational failure.

Table 4-6 Physical Disk Drive States (**continued**)

State	Description
Online rebuild	The disk drive is in the process of being rebuilt. (In a RAID 1 or 0+1 array, data is being copied from the mirrored disk drive to the replacement disk drive. In a RAID 3 or 5 array, data is being regenerated by the exclusive OR (XOR) algorithm and written to the replacement disk drive.)
Unconfigured	This location is unconfigured.
Environmental	An environmental device is present at this address.

For more information, see the TPM documentation and online help.

Automatic Rebuild

The RAID controller provides automatic rebuild capabilities in the event of a physical disk drive failure. The controller performs a rebuild operation automatically when a disk drive fails and the following conditions are true:

- All system drives that are dependent on the failed disk drive are configured as a redundant array; RAID 1, RAID 3, RAID 5, or RAID 0+1;
- The Automatic rebuild management controller parameter is enabled;
- The Operational fault management controller parameter is enabled; and
- A replacement drive with a capacity that is at least as large as the consumed capacity of the failed drive is present in the system.

Note: If a replacement drive of the exact size is not available, the controller selects the smallest replacement drive found with a capacity that is at least as large as the consumed capacity of the failed drive. The consumed capacity is the capacity assigned to the configured system drive(s). If the consumed capacity of the failed disk drive is a percentage of the total capacity, a larger physical disk drive can be rebuilt with a much smaller physical disk drive.

During the automatic rebuild process, system activity continues as normal. However, system performance may be slightly degraded.

Note: The priority of rebuild activity can be adjusted using the controller parameters to adjust the Rebuild and check consistency rate.

In order to use the automatic rebuild feature, you must maintain an online spare disk drive in the system. The number of online spare disk drives in a system is limited only by the maximum number of disk drives available on each drive channel. SGI recommends creating an online spare disk drive as part of the original configuration, or soon after creating the original configuration. If the online spare disk drive is created after a disk drive failure has occurred, the automatic rebuild does not start until the controllers have been reset.

A disk drive may be labeled as an online spare using the “**create hot spare option**” of the TPM configuration utility.

The RAID controllers also support the ability to perform a hot swap disk drive replacement while the system is online. A disk drive can be disconnected, removed, and replaced with a different disk drive without taking the system offline.

Caution: System drives associated with a failed or removed disk drive become critical. Failure or removal of another disk drive may result in data loss.

The automatic rebuild feature is dependent upon having an online spare disk drive available or hot swapping the failed disk drive with a replacement drive. If these conditions are not met, the automatic rebuild features does not operate transparently, or without user intervention. Automatic rebuild will not start if an online spare is configured after a disk drive has failed.

Note: A “ghost drive” is created when a disk drive fails, power is removed from the system, the disk drive is replaced or a spare drive is added to the system, and power is returned to the system. Automatic rebuild does not occur in this situation. Additionally, the system does not recognize the replacement/spare disk drive and creates a ghost drive in the same location as the failed disk drive. If the replacement/spare disk drive was inserted into the same slot as the failed drive, the ghost drive appears in the first available empty slot, beginning with channel 0, target 0. The ghost drive represents a deleted, dead drive that still exists in the configuration and the replacement/spare disk drive has a drive state of unconfigured. In order for the rebuild to occur, the replacement/spare disk drive’s state must change from unconfigured to online spare.

The rebuild procedure begins after a REBUILD has been started or power has been cycled to the controllers. Cycling the power also removes the “ghost drive” from the configuration.

Troubleshooting

The 2 Gb TP9100 storage system includes a processor and associated monitoring and control logic that allows it to diagnose problems within the storage system's power, cooling, and drive systems.

SES (SCSI enclosure services) communications are used between the storage system and the RAID controllers. Status information on power, cooling, and thermal conditions is communicated to the controllers and is displayed in the management software interface.

The enclosure services processor is housed in the ESI/ops panel module. The sensors for power, cooling, and thermal conditions are housed within the power supply/cooling modules. Each module in the storage system is monitored independently.

Note: For instructions on opening the rear door of the rack, see “Opening and Closing the Rear Rack Door” on page 31.

This chapter contains the following sections:

- “RAID Guidelines” on page 76
- “Solving Initial Startup Problems” on page 76
- “Using Storage System LEDs for Troubleshooting” on page 78
- “Using the Alarm for Troubleshooting” on page 91
- “Solving Storage System Temperature Issues” on page 92
- “Using Test Mode” on page 93
- “Care and Cleaning of Optical Cables” on page 94

RAID Guidelines

RAID stands for “redundant array of independent disks”. In a RAID system multiple disk drives are grouped into arrays. Each array is configured as system drives consisting of one or more disk drives. A small, but important set of guidelines should be followed when connecting devices and configuring them to work with a controller.

Follow these guidelines when configuring a RAID system:

- Distribute the disk drives equally among all the drive channels on the controller. This results in better performance. The TP9100 has two drive channels.
- A drive pack can contain a maximum of 16 drives.
- A drive pack can contain drives that are on any drive channel.
- If configuring an online spare disk drive, ensure that the spare disk drive capacity is greater than or equal to the capacity of the largest disk drive in all redundant drive packs.
- When replacing a failed disk drive, ensure that the replacement disk drive capacity is greater than or equal to the capacity of the failed disk drive in the affected drive pack.

Solving Initial Startup Problems

If cords are missing or damaged, plugs are incorrect, or cables are too short, contact your supplier for a replacement.

If the alarm sounds when you power on the storage system, one of the following conditions exists:

- A fan is slowing down. See “Power Supply/Cooling Module LEDs” on page 84 for further checks to perform.
- Voltage is out of range. The tower requires 115/220 Volts (autoranging), and the rack requires 200-240 Volts (autoranging).
- There is an overtemperature or thermal overrun condition. See “Solving Storage System Temperature Issues” on page 92.

- There is a storage system fault. See “ESI/Ops Panel LEDs and Switches” on page 79.
- There are mixed single-port and dual-ports modules within an enclosure. Only one type of module may be installed in an enclosure.

If the SGI server does not recognize the storage system, check the following:

- Ensure that the device driver for the host bus adapter board has been installed. If the HBA was installed at the factory, this software is in place; if not, check the HBA and the server documentation for information on the device driver.
- Ensure the FC-AL interface cables from the LRC I/O module to the Fibre Channel board in the host computer are installed correctly.
- Check the selector switches on the ops panels of the storage system as follows:
 - On a tower or a RAID enclosure the ops panel should be set to address 1.
 - On the first expansion enclosure attached to a RAID system, the ops panel should be set to address 2.
 - On the first enclosure in a JBOD system, the ops panel should be set to address 1. Other enclosures that are daisy chained to the first enclosure should be addressed sequentially (2-7).
- Ensure that the LEDs on all installed drive carrier modules are green. Note that the drive LEDs flash during drive spinup.
- Check that all drive carrier modules are correctly installed.
- If an amber disk drive module LED drive fault is on, there is a drive fault. See Table 5-7 on page 90.

If the SGI server connected to the storage system is reporting multiple Hard Error (SCS_DATA_UNDERRUN) errors in the `/var/adm/SYSLOG`, the cabling connected to the controller reporting the errors requires cleaning or replacement. For more information on cleaning cables, refer to “Care and Cleaning of Optical Cables” on page 94.

Using Storage System LEDs for Troubleshooting

This section summarizes LED functions and gives instructions for solving storage system problems in these subsections:

- “ESI/Ops Panel LEDs and Switches” on page 79
- “Power Supply/Cooling Module LEDs” on page 84
- “RAID LRC I/O Module LEDs” on page 85
- “RAID Loopback LRC I/O Module LEDs” on page 89
- “Drive Carrier Module LEDs” on page 90

ESI/Ops Panel LEDs and Switches

Figure 5-1 shows details of the ESI/ops panel.




Figure 5-1 ESI/Ops Panel Indicators and Switches

Table 5-1 summarizes functions of the LEDs on the ESI/ops panel

Table 5-1 ESI/Ops Panel LEDs

LED	Description	Corrective Action
Power on	This LED illuminates green when power is applied to the enclosure.	N/A
Invalid address	This LED flashes amber when the enclosure is set to an invalid address mode.	Change the enclosure address thumb wheel to the proper setting. If the problem persists, contact your service provider.

Table 5-1 ESI/Ops Panel LEDs (**continued**)

LED	Description	Corrective Action
System/ESI fault	This LED illuminates amber and the audible alarm sounds when the ESI processor detects an internal problem. This LED flashes when an over- or under-temperature condition exists.	Contact your service provider.
PSU/cooling/ temperature fault	This LED illuminates amber if an over- or under-temperature condition exists. This LED flashes if there is an ESI communications failure.	Check for proper airflow clearances and remove any obstructions. If the problem persists, lower the ambient temperature. In case of ESI communications failure, contact your service provider.
Hub mode	This LED illuminates green when the host side switch is enabled (RAID only).	N/A
2-Gb link speed	This LED illuminates green when 2-Gb link speed is detected.	N/A

The Ops panel switch settings for a JBOD enclosure are listed in Table 5-2. The ops panel switch settings for a RAID enclosure are listed in Table 5-3 on page 82. The switches are read only during the power-on cycle.

Table 5-2 Ops Panel Configuration Switch Settings for JBOD

Switch Number	Function	Function When Off	Function When On
1 (On ^a)	Loop select single (1x16) or dual (2x8)	LRC operates as 2 loops of 8 drives (2x8). Refer to drive addressing mode 2.	LRC operates as 1 loop of 16 drives (1x16 loop mode)
2 (On)	Loop terminate mode	If no signal is present on external FC port, the loop is left open	If no signal is present on external FC port, then the loop is closed
3 (Off) ^b	N/A	N/A	N/A
4 (Off)	N/A	N/A	N/A

Table 5-2 Ops Panel Configuration Switch Settings for JBOD (continued)

Switch Number	Function	Function When Off		Function When On
		Sw 5	Sw 6	Function
5 and 6	RAID host hub speed select Used only for RAID configurations.	Sw 5	Sw 6	Function
		Off	Off	Force 1 Gb/s
		On	Off	Force 2 Gb/s
		Off	On	Reserved
		On	On	Auto loop speed detect based on LRC port signals Note: This feature is not supported
7 and 8	Drive loop speed select	Sw 7	Sw 8	Function
		Off	Off	Force 1 Gb/s
		On	Off	Force 2 Gb/s
		Off	On	Speed selected by EEPROM bit
		On	On	Auto loop speed detect based on LRC port signals Note: This feature is not supported
9 and 10	Drive addressing mode	Sw 9	Sw 10	Function
		On	On	Mode 0 - Single loop, base 16, offset of 4, 7 address ranges ^c
		Off	On	Mode 1 - Single loop, base 20, 6 address ranges
		On ^d	Off	Mode 2 - JBOD, dual loop, base 8, 15 address ranges
		Off	Off	Mode 3 (Not used)
11 (On)	Soft select	Selects switch values stored in EEPROM		Selects switch values from hardware switches Note: Soft select switch must be set to On
12 (Off)	Not used	Not used		Not used

a. Bolded entries indicate default switch settings for a 1X16 JBOD. Set switches 1 and 10 to **Off** for 2X8 JBOD.

b. Switches 3, 5, and 6 are used in RAID configurations

c. Mode 0 (switches 9 and 10 set to **On**) is the SGI default factory setting

d. Selecting mode 2 forces 2x8 dual loop selection

Table 5-3 Ops Panel Configuration Switch Settings for RAID

Switch Number	Function	Function when Off		Function when On
1 (On)	Loop select single (1x16) or dual (2x8)	LRC operates as two loops of 8 drives		LRC operates as 1 loop of 16 drives (1x16 loop mode)
2 (On)	Loop terminate mode	If no signal is present on the external FC port, then the loop is left open		If no signal is present on the external FC port then the loop is closed internally
3 (Off)	Hub mode select (RAID only)	Hub ports connect independently		RAID host FC ports are linked together internally
4 (Off)	Not used.			
5 and 6	RAID host hub speed select Note: Set switches 5 and 6 to Off to force 1 Gb/s if connecting RAID controllers to 1-Gb/s HBAs or switches	Sw 5	Sw 6	Function
		Off	Off	Force 1 Gb/s
		On	Off	Force 2 Gb/s
		Off	On	Reserved
On	On	Auto loop speed detect based on LRC port signals Note: This feature is not supported		
7 and 8	Drive loop speed select	Sw 7	Sw 8	Function
		Off	Off	Force 1 Gb/s
		On	Off	Force 2 Gb/s
		Off	On	Speed selected by EEPROM bit
On	On	Auto loop speed detect based on LRC port signals Note: This feature is not supported		

Table 5-3 Ops Panel Configuration Switch Settings for RAID (continued)

Switch Number	Function	Function when Off		Function when On
		Sw 9	Sw 10	Function
9 and 10	Drive addressing mode	On	On	Mode 0 - Single loop, base 16, offset of 4, 7 address ranges ^b
		Off	On	Mode 1 - Single loop, base 20, 6 address ranges
		On	Off	Mode 2 - JBOD, dual loop, base 8, 15 address ranges ^c
		Off	Off	Mode 3 (Not used)
11 (On)	Soft select	Selects switch values stored in EEPROM		Selects switch values from hardware switches Note: Soft select switch must be set to On
12 (Off)	Not used	Not used		Not used

a. Bolded entries indicate SGI's default switch settings for RAID

b. Mode 0 (switches 9 and 10 set to **On**) is the SGI default factory setting

c. Mode 2 (2x8) is not supported in a RAID configuration

Note the following:

- If all LEDs on the ESI/ops panel flash simultaneously, see "Using Test Mode" on page 93.
- If test mode has been enabled (see "Using Test Mode" on page 93), the amber and green drive bay LEDs flash for any non-muted fault condition.

Power Supply/Cooling Module LEDs

Figure 5-2 shows the meanings of the LEDs on the power supply/cooling module.




Figure 5-2 Power Supply/Cooling Module LED

If the green “PSU good” LED is not lit during operation, or if the power/cooling LED on the ESI/ops panel is amber and the alarm is sounding, contact your service provider.

RAID LRC I/O Module LEDs

Figure 5-3 shows the LEDs on the dual-port RAID LRC I/O module.

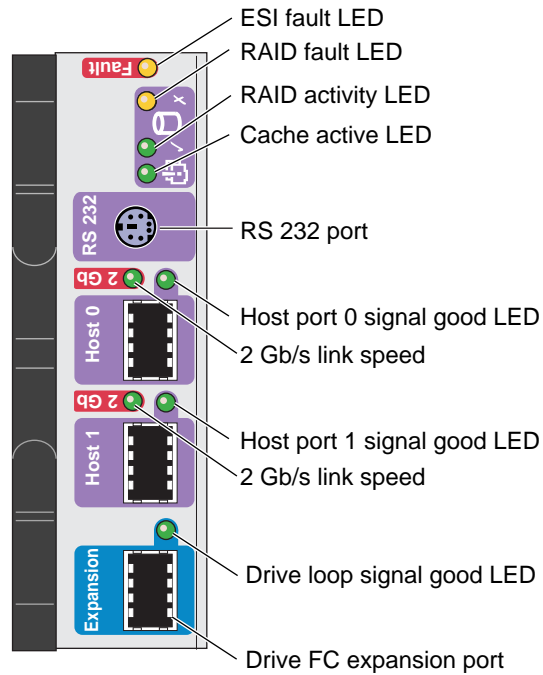


Figure 5-3 Dual-port RAID LRC I/O Module LEDs

Table 5-4 explains what the LEDs in Figure 5-3 indicate.

Table 5-4 Dual-port RAID LRC I/O Module LEDs

LED	Description	Corrective Action
ESI fault	This LED illuminates amber and the audible alarm sounds when the ESI processor detects an internal problem.	Check for mixed single-port and dual-port modules within an enclosure. Also check the drive carrier modules and PSU/cooling modules for faults. If the problem persists, contact your service provider.
RAID fault	This LED illuminates amber when a problem with the RAID controller is detected.	Contact your service provider.
RAID activity	This LED flashes green when the RAID controller is active.	N/A
Cache active	This LED flashes green when data is read into the cache.	N/A
Host port signal good (1 and 2)	This LED illuminates green when the port is connected to a host.	Check both ends of the cable and ensure that they are properly seated. If the problem persists, contact your service provider.
Drive loop signal good	This LED illuminates green when the expansion port is connected to an expansion enclosure.	Check both ends of the cable and ensure that they are properly seated. If the problem persists, contact your service provider.

Figure 5-4 shows the LEDs on the single-port RAID LRC I/O module.




Figure 5-4 Single-port RAID LRC I/O Module LEDs

Table 5-5 explains what the LEDs in Figure 5-4 indicate.

Table 5-5 Single-port RAID LRC I/O Module LEDs

LED	Description	Corrective Action
ESI fault	This LED illuminates amber and the audible alarm sounds when the ESI processor detects an internal problem.	Check the drive carrier modules and PSU/cooling modules. If the problem persists, contact your service provider.
RAID fault	This LED illuminates amber when a problem with the RAID controller is detected.	Contact your service provider.

Table 5-5 Single-port RAID LRC I/O Module LEDs **(continued)**

LED	Description	Corrective Action
RAID activity	This LED flashes green when the RAID controller is active.	N/A
Cache active	This LED flashes green when data is read into the cache.	N/A
Host port signal good	This LED illuminates green when the port is connected to a host.	Check both ends of the cable and ensure that they are properly seated. If the problem persists, contact your service provider.
Drive loop signal good	This LED illuminates green when the expansion port is connected to an expansion enclosure.	Check both ends of the cable and ensure that they are properly seated. If the problem persists, contact your service provider.

RAID Loopback LRC I/O Module LEDs

The LEDs on the rear of the RAID loopback LRC I/O module function similarly to those on the RAID LRC I/O modules. See “RAID LRC I/O Module LEDs” on page 85 for more information.

JBOD LRC I/O Module LEDs

Figure 5-5 shows the JBOD LRC I/O module LEDs.




Figure 5-5 JBOD LRC I/O Module LEDs

Table 5-6 explain what the LEDs in Figure 5-5 indicate.

Table 5-6 JBOD LRC I/O Module LEDs

LED	Description	Corrective Action
ESI fault	This LED illuminates amber and the audible alarm sounds when the ESI processor detects an internal problem.	Check the drive carrier modules and PSU/cooling modules. If the problem persists, contact your service provider.
FC-AL signal present	These LEDs illuminate green when the port is connected to an FC-AL.	Check the cable connections. If the problem persists, contact your service provider.

Drive Carrier Module LEDs

Each disk drive module has two LEDs, an upper (green) and a lower (amber), as shown in Figure 5-6.



Figure 5-6 Drive Carrier Module LEDs

Table 5-7 explains what the LEDs in Figure 5-6 indicate.

Table 5-7 Disk Drive LED Function

Green LED	Amber LED	State	Remedy
Off	Off	Disk drive not connected; the drive is not fully seated.	Check that the drive is fully seated
On	Off	Disk drive power is on, but the drive is not active.	N/A

Table 5-7 Disk Drive LED Function (**continued**)

Green LED	Amber LED	State	Remedy
Blinking	Off	Disk drive is active. (LED might be off during power-on.)	N/A
Flashing at 2-second intervals	On	Disk drive fault (SES function).	Contact your service provider for a replacement drive and follow instructions in Chapter 6.
N/A	Flashing at half-second intervals	Disk drive identify (SES function).	N/A

In addition, the amber drive LED on the ESI/ops panel alternates between on and off every 10 seconds when a drive fault is present.

Using the Alarm for Troubleshooting

The ESI/ops panel includes an audible alarm that indicates when a fault state is present. The following conditions activate the audible alarm:

- RAID controller fault
- Fan slows down
- Voltage out of range
- Over-temperature
- Storage system fault

You can mute the audible alarm by pressing the alarm mute button for about a second, until you hear a double beep. The mute button is beneath the indicators on the ESI/ops panel (see Figure 5-1 on page 79).

When the alarm is muted, it continues to sound with short intermittent beeps to indicate that a problem still exists. It is silenced when all problems are cleared.

Note: If a new fault condition is detected, the alarm mute is disabled.

Solving Storage System Temperature Issues

This section explains storage system temperature conditions and problems in these subsections:

- “Thermal Control” on page 92
- “Thermal Alarm” on page 93

Thermal Control

The storage system uses extensive thermal monitoring and ensures that component temperatures are kept low and acoustic noise is minimized. Airflow is from front to rear of the storage system. Dummy modules for unoccupied bays in enclosures and blanking panels for unoccupied bays in the rack must be in place for proper operation.

If the ambient air is cool (below 25 °C or 77 °F) and you can hear that the fans have sped up by their noise level and tone, then some restriction on airflow might be raising the storage system’s internal temperature. The first stage in the thermal control process is for the fans to automatically increase in speed when a thermal threshold is reached. This might be a normal reaction to higher ambient temperatures in the local environment. The thermal threshold changes according to the number of drives and power supplies fitted.

If fans are speeding up, follow these steps:

1. Check that there is clear, uninterrupted airflow at the front and rear of the storage system.
2. Check for restrictions due to dust buildup; clean as appropriate.
3. Check for excessive recirculation of heated air from the rear of the storage system to the front.
4. Check that all blank plates and dummy disk drives are in place.
5. Reduce the ambient temperature.

Thermal Alarm

The four types of thermal alarms and the associated corrective actions are described in Table 5-8.

Table 5-8 Thermal Alarms

Alarm Type	Indicators	Solutions
High temp warning begins at 54° C (129° F)	Audible alarm sounds Ops panel system fault LED flashes Fans run at higher speed than normal SES temperature status is non-critical PSU Fault led is lit	If possible, power down the enclosure; then check the following: <ul style="list-style-type: none"> • Ensure that the local ambient temperature meets the specifications outlined in “Environmental Requirements” on page 105 • Ensure that the proper clearances are provided at the front and rear of the rack. • Ensure that the airflow through the rack is not obstructed. If you are unable to determine the cause of the alarm, please contact your service provider.
High temp failure begins at 58° C (136° F)	Audible alarm sounds Ops panel system fault LED flashes Fans run at higher speed than normal SES temperature status is critical PSU Fault led is lit	
Low temp warning begins at 10° C (50° F)	Audible alarm sounds Ops panel system fault is lit SES temperature status is non-critical	
Low temp failure begins at 0° C (32° F)	Audible alarm sounds Ops panel system fault is lit SES temperature status is critical	

Using Test Mode

When no faults are present in the storage system, you can run test mode to check the LEDs and the audible alarm on the ESI/ops panel. In this mode, the amber and green LEDs on each of the drive carrier modules and the ESI/ops panel flash on and off in sequence; the alarm beeps twice when test mode is entered and exited.

To activate test mode, press the alarm mute button until you hear a double beep. The LEDs then flash until the storage system is reset, either when you press the alarm mute button again or if an actual fault occurs.

Care and Cleaning of Optical Cables



Warning: Never look into the end of a fiber optic cable to confirm that light is being emitted (or for any other reason). Most fiber optic laser wavelengths (1300 nm and 1550 nm) are invisible to the eye and cause permanent eye damage. Shorter wavelength lasers (for example, 780 nm) are visible and can cause significant eye damage. Use only an optical power meter to verify light output.



Warning: Never look into the end of a fiber optic cable on a powered device with any type of magnifying device, such as a microscope, eye loupe, or magnifying glass. Such activity can cause a permanent burn on the retina of the eye. Optical signal cannot be determined by looking into the fiber end.

Fiber optic cable connectors must be kept clean to ensure long life and to minimize transmission loss at the connection points. When the cables are not in use, replace the caps to prevent deposits and films from adhering to the fiber. A single dust particle caught between two connectors will cause significant signal loss. In addition to causing signal loss, dust particles can scratch the polished fiber end, resulting in permanent damage. Do not touch the connector end or the ferrules; your fingers will leave an oily deposit on the fiber. Do not allow uncapped connectors to rest on the floor.

If a fiber connector becomes visibly dirty or exhibits high signal loss, carefully clean the entire ferrule and end face with special lint-free pads and isopropyl alcohol. The end face in a bulkhead adapter on test equipment can also be cleaned with special lint-free swabs and isopropyl alcohol. In extreme cases, a test unit may need to be returned to the factory for a more thorough cleaning.

Never use cotton, paper, or solvents to clean fiber optic connectors; these materials may leave behind particles or residues. Instead, use a fiber optic cleaning kit especially made for cleaning optical connectors, and follow the directions. Some kits come with canned air to blow any dust out of the bulkhead adapters. Be cautious, as canned air can damage the fiber if not used properly. Always follow the directions that come with the cleaning kit.

Installing and Replacing Drive Carrier Modules

This chapter explains how to install a new drive carrier or replace an existing one in the following sections:

- “Adding a Drive Carrier Module” on page 95
- “Replacing a Drive Carrier Module” on page 98

Note: The RAID controller supports hot-swap disk drive replacement while the storage system is online: depending on the RAID level, a disk drive can be disconnected, removed, or replaced with another disk drive without taking the storage system offline.



Caution: Observe all ESD precautions when handling modules and components. Avoid contact with backplane components and module connectors. Failure to observe ESD precautions could damage the equipment.

Adding a Drive Carrier Module

Note the following:

- All disk drive bays must be filled with either a drive carrier module or a dummy drive; no bay should be left completely empty.
- The drives in bays 1/1 and 4/4 are required for enclosure management; these bays must always be occupied.

To add a new disk drive module to the storage system, follow these steps:

1. Ensure that you have enough drive carrier modules and dummy modules to occupy all bays.
2. Carefully open the bag containing the drive carrier module.



Warning: The disk drive handle might have become unlatched in shipment and might spring open when you open the bag. As you open the bag, keep it a safe distance from your face.

3. Place the drive carrier module on an antistatic work surface and ensure that the anti-tamper lock is disengaged (unlocked). A disk drive module cannot be installed if its anti-tamper lock is activated outside the enclosure.

Drives are shipped with their locks set in the unlocked position. However, if a drive is locked, insert the key (included with the disk drive) into the socket in the lower part of the handle trim and turn it 90 degrees counterclockwise until the indicator visible in the center aperture of the handle shows black. See Figure 6-1.




Figure 6-1 Unlocking the Drive Carrier Module

4. Open the handle of the replacement carrier by pressing the latch handle towards the right (see Figure 6-2).




Figure 6-2 Opening the Module Handle

5. Orient the module so that the hinge of the handle is on the right. Then slide the disk carrier module into the chassis until it is stopped by the camming lever on the right of the module (see Figure 6-3).




Figure 6-3 Inserting the Disk Drive Module in a Rackmount Enclosure

6. Swing the drive handle shut and press it to seat the drive carrier module. The camming lever on the right of the module will engage with a slot in the chassis. Continue to push firmly until the handle fully engages with the module cap. You should hear a click as the latch engages and holds the handle closed.
7. Repeat steps 2 through 6 for all drive modules to be installed.
8. When you have finished installing the drive carrier module(s), activate the anti-tamper lock(s). Insert the key and turn it 90 degrees clockwise. The indicator in the drive carrier module turns red when the drive is locked. See Figure 6-4.




Figure 6-4 Locking the Drive Carrier Module

9. Fit any empty drive bays with dummy drive carrier modules. The drive handle and camming mechanisms operate the same way that those in a standard drive carrier module do.

Replacing a Drive Carrier Module

This section explains how to replace a defective drive carrier module in the following sections:

- “LUN Integrity and Drive Carrier Module Failure” on page 99
- “Replacing the Disk Drive Module” on page 100

LUN Integrity and Drive Carrier Module Failure

When a disk drive fails in a RAID 5, 3, 1, or 0+1 LUN, the amber LEDs on all disks in the LUN (except the failed one) alternate on and off every 1.2 seconds until the fault condition is cleared. The amber LED on the failed disk remains lit.

Note: Before replacing a drive carrier module, use the storage system software to check the disk status.

For a RAID 5, 3, 1, or 0+1 LUN, you can replace the disk module without powering off the array or interrupting user applications. If the array contains a hot spare on standby, the controller automatically rebuilds the failed module on the hot spare. A hot spare is a special LUN that acts as a global disk spare that can be accessed by any RAID 5, 3, 1, or 0+1 LUN. A hot spare is unowned until it becomes part of a LUN when one of the LUN's disk modules fails.

A RAID 0 array must be taken offline to be replaced if a single disk module fails. Also, if a second disk drive fails in a RAID 5, 3, or 1 LUN, the system drive is marked offline—regardless of whether a second hot spare is available—and the host cannot access data from that system drive.

In these cases, the LUN's data integrity is compromised and it becomes unowned (not accessible by the controller). After you replace the failed disk modules (one at a time), you must delete and then re-create the affected LUN(s). If the data on the failed disks was backed up, restore it to the new disks.

Note: If a disk fails in a LUN and the storage system puts the hot spare into the LUN, use the software included with the storage system to check disk module status, and replace the failed disk as soon as possible. The replacement becomes the new hot spare; this arrangement (drive roaming) differs from that of other RAID systems. Therefore, it is important to keep track of the location of the hot spare.

Replacing the Disk Drive Module

If an LED indicates that a disk drive is defective, follow these steps to remove the faulty drive:

1. Make sure enough disk drives and dummy drives are available to occupy all bays.
2. Ensure that users are logged off of the affected systems; back up data if necessary.

Note: Replace disk drive modules one at a time.

3. If the drive module is locked, insert the key into the anti-tamper lock and turn it 90 degrees counterclockwise. The indicator in the drive carrier module turns black when the drive is unlocked.




Figure 6-5 Unlocking the Disk Drive Module

4. Ensure that the faulty drive has spun down.



Caution: Damage can occur to a drive if it is removed while still spinning.

5. Open the handle by pressing the latch on the module handle towards the right. Then gently slide the module out of the enclosure approximately 25 mm (1 inch) and wait 30 seconds. See Figure 6-6.




Figure 6-6 Removing the Drive Carrier Module

6. Withdraw the module from the drive bay. Replace it immediately; follow instructions in “Adding a Drive Carrier Module” on page 95.
7. If you are replacing a module in a LUN that uses a hot spare, note the location of the replacement module; it is the new hot spare.

Technical Specifications

This appendix contains the following sections:

- “Storage System Physical Specifications” on page 103
- “Environmental Requirements” on page 105
- “Power Requirements” on page 106
- “LRC I/O Module Specifications” on page 107
- “Disk Drive Module Specifications” on page 108
- “SGI Cables for the 2 Gb TP9100 Storage System” on page 108

Storage System Physical Specifications

Table A-1 provides the dimensions for the SGI 2 Gb TP9100 enclosure, tower, and rack.

Dimension	Rackmount Enclosure	Tower Enclosure	Rack
Height	13.4 cm (5.3 in.)	50.1 cm (19.7 in.)	Operating: 180 cm (5.94 ft) Shipping: 210 cm (6.93 ft)
Width	44.6 cm (17.5 in.)	23 cm (9 in.) including feet	Operating: 60 cm (1.98 ft) Shipping: 120 cm (3.96 ft)
Depth	50 cm (19.7 in.)	52.3 cm (20.6 in.)	Operating: 81 cm (2.64 ft) Shipping: 120cm (6.93 ft)

Table A-2 shows the weights of various component modules.

Table A-2 Weights

Component	Weight
Enclosure, fully populated	Rackmount: 32.3 kg (71 lb) Tower: 42.3 kg (93.0 lb)
Enclosure, empty	Rackmount: 17.9 kg (39.4 lb) Tower: 12 kg (26.4 lb)
Power supply/cooling module	3.6 kg (7.9 lb)
Disk carrier module with 36-GB drive	0.88 kg (1.9 lb)
LRC I/O module	1.2 kg (2.6 lb)
Tower conversion kit	10 kg (22 lb)

Table A-3 shows the power requirements and specifications of the 2 Gb TP9100.

Table A-3 Power Specifications

Specification	Value
Voltage range for Rack	200-240 VAC
Voltage range for Tower	100-120/220-240 VAC
Voltage range selection	Automatic
Frequency	50-60 Hz
Power factor	>0.98
Harmonics	Meets EN61000-3-2
Power cord:	
Cord type	SV or SVT, 18 WG minimum, 3 conductor
Plug	250 V, 10 A
Socket	IEC 320 C-14, 250 V, 15 A

Environmental Requirements

Table A-4 provides temperature and humidity requirements for both the rack and tower storage systems.

Table A-4 Ambient Temperature and Humidity Requirements

Factor	Temperature	Relative Humidity	Maximum Wet Bulb
Operating temperature	5 °C to 40 °C (41 °F to 104 °F)	20% to 80% noncondensing	23 °C (73 °F)
Non-operating temperature	0 °C to 50 °C (32 °F to 122 °F)	8% to 80% noncondensing	27 °C (80 °F)
Storage temperature	1 °C to 60 °C (34 °F to 140 °F)	8% to 80% noncondensing	29 °C (84 °F)
Shipping temperature	-40 °C to +60 °C (-40 °F to 140 °F)	5% to 100% nonprecipitating	29 °C (84 °F)

Table A-5 gives other environmental specifications for both the rack and tower storage systems.

Table A-5 Environmental Specifications

Environmental Factor	Requirement
Altitude, operating	0 to 3047 m (0 to 10,000 ft)
Altitude, non-operating	-305 to 12,192 m (-1000 to 40,000 ft)
Shock, operating	Vertical axis 5 g peak 1/2 sine, 10 ms
Shock, non-operating	30 g 10 ms 1/2 sine
Vibration, operating	0.21 grms 5-500 Hz random
Vibration, Non-operating	1.04 grms 2-200 Hz random
Acoustics	Less than 6.0 B LwA operating at 20 °C
Safety and approvals	CE, UL, cUL
EMC	EN55022 (CISPR22-A), EN55024 (CISPR24), FCC-A

Power Requirements

Table A-6 provides minimum storage system power requirements.

Table A-6 Minimum Power Requirements

Factor	Requirement
Voltage	Tower: 100 to 120 or 220 to 240 VAC Rack: 200 to 240 VAC
Frequency	50 to 60 Hz
Maximum power consumption	700 VA
Typical power consumption	400 VA or less
Inrush current (25 °C (77 °F) cold start 1 PSU)	100 A maximum peak for 4ms, 25 A thereafter at maximum voltage

Table A-7 provides additional information for the power distribution units (PDUs) in the rack.

Table A-7 Rack PDU Power Specifications

Factor	Requirement or Specification
Ratings	200 to 240 VAC, 24 A, 50 to 60 Hz
Over-voltage category	II
Maximum load per PDU	24 A
Maximum load per bank of outlet sockets on each circuit breaker	10 A
Plug	NEMA L6-30

LRC I/O Module Specifications

Table A-8 provides specifications for the LRC I/O module.

Table A-8 LRC I/O Module Specifications

Specification	Value
Connectors	2 x SFP module LC optical, maximum cable length 300 m 1x SFP expansion port, maximum copper cable length 1 m
External FC-AL signal cables	SGI dual-port HBAs: 25 m (82 ft) Storage area network (SAN) and SGI single-port HBAs: maximum 100 m (328 ft) optical (see Table A-10 for information on cables)
Drive interface	2 x FC-AL loops, connected internally to FCAL LRC I/O
Power dissipation	3 A @ 3.3 V 2 A @ 5 V 2 A @ 12 V
RAID levels	0, 1, 3, 5, and 0+1 (RAID level 6) JBOD (RAID level 7)
LED Indicators	Drive loop signal good - green Host port 1 and 2 signal good - green Cache active - green RAID active - green RAID fault - amber ESI/LRC module fault LED - amber
Memory	512 MB maximum
Cache	Selectable write-through or write-back Read always enabled
Battery	NiCad cache battery protects 512 MB data for up to 72 hours

Disk Drive Module Specifications

Consult your supplier for details of disk drives supported for use with the RAID storage system. Table A-9 provides specifications for a typical drive carrier module.

Table A-9 Drive Carrier Module Specifications (1.6-inch 36-GB Drive)

Factor	Requirement
Dimensions	Height 2.91 cm (1.1 in.)
	Width 10.65 cm (4.2 in.)
	Depth 20.7 cm (8.1 in.)
Weight	.88 kg (1.9 lb) with 36-GB drive
Operating temperature	5 °C to 40 °C (41 °F to 104 °F) when installed
Power dissipation	22 W maximum

SFI Cables for the 2 Gb TP9100 Storage System

Table A-10 lists SFI cable options that can be connected to the 2 Gb TP9100 product.

Table A-10 SFI Fibre Channel Fabric Cabling Options for the 2 Gb TP9100 Storage System

Cable	Length	Marketing Code	Part Number
1 m copper SFP to copper SFP cable	1 m (3.3 ft)	TP912G-CASCADE	018-1081-001
FC optical cable (62.5 µm)	3 m (9.8 ft)	X-F-OPT-3M	018-0656-001
	10 m (32.8 ft)	X-F-OPT-10M	018-0656-101
	25 m (82 ft)	X-F-OPT-25M	018-0656-201
	100 m (328 ft)	X-F-OPT-100M	018-0656-301
	300 m (980 ft) ^a	X-F-OPT-300M	018-0656-401

a. This cable is not authorized for use with SFI Fibre Channel switches.

Regulatory Information

The SGI 2 Gb Total Performance 9100 (2 Gb TP9100) conforms to Class A specifications.

Note: This equipment is for use with Information Technology Equipment only.

FCC Warning

This equipment has been tested and found compliant with the limits for a Class A digital device, pursuant to Part 15 of the FCC rules. These limits are designed to provide reasonable protection against harmful interference when the equipment is operated in a commercial environment. This equipment generates, uses, and can radiate radio frequency energy and, if not installed and used in accordance with the instruction manual, may cause harmful interference to radio communications. Operation of this equipment in a residential area is likely to cause harmful interference in which case the user will be required to correct the interference at his own expense.

Attention

This product requires the use of external shielded specified by the manufacturer or optical cables in order to maintain compliance pursuant to Part 15 of the FCC Rules.

European Union Statement

This device complies with the European Directives listed on the “Declaration of Conformity” which is included with each product. The CE mark insignia displayed on the device is an indication of conformity to the aforementioned European requirements.



International Special Committee on Radio Interference (CISPR)

This equipment has been tested to and is in compliance with the Class A limits per CISPR publication 22, Limits and Methods of Measurement of Radio Interference Characteristics of Information Technology Equipment; and Japan's VCCI Class 1 limits.

Canadian Department of Communications Statement

This digital apparatus does not exceed the Class A limits for radio noise emissions from digital apparatus as set out in the Radio Interference Regulations of the Canadian Department of Communications.

Attention

Cet appareil numérique n'émet pas de perturbations radioélectriques dépassant les normes applicables aux appareils numériques de Classe A prescrites dans le Règlement sur les interférences radioélectriques établi par le Ministère des Communications du Canada.

VCCI Class 1 Statement

この装置は、情報処理装置等電波障害自主規制協議会（VCCI）の基準に基づくクラス A 情報技術装置です。この装置を家庭環境で使用すると電波妨害を引き起こすことがあります。この場合には使用者が適切な対策を講ずるよう要求されることがあります。

Class A Warning for Taiwan

警告使用者：

這是甲類的資訊產品，在居住的環境中使用時，可能會造成射頻干擾，在這種情況下，使用者會被要求採取某些適當的對策。

Index

A

affinity. See LUN affinity and system drive.
airflow, 92
alarm
 and troubleshooting, 91
 at power-on, 76
 muting, 91
 thermal, 93
 with LED on power supply/cooling module, 84
automatic rebuild, 71
availability, configuring for maximum, 63-65

B

bay numbering
 rackmount enclosure, 23
 tower, 24
breaker in rack, 43

C

cable
 fibre channel, 108
 power, 37-42
 to host, 36
caching
 disabled, 51
 enabled, 50, 64
 write-back, 50, 64
 write-through, 51

CAP, 61-65
capacity, availability, and performance. See CAP.
capacity, configuring for maximum, 62-63
Carrier module, 19
 dummy, 21
 replacement procedure, 95
chassis grounding, 37
Class A, 109-111
COD, 48-49
components
 enclosure, 8-22
 weight, 104
configuration, 65-68
 RAID, 65
 selecting RAID level, 61-65
configuration on disk, 48-49
controller parameters, take effect after reboot, 68
conventions, xv
Cooling
 overview, 12
current limit, 36
customer service, xv

D

data caching, 50-51
 See also caching.
deskside. See tower.
device driver, 77

Disk drive carrier module, 19
 replacement procedure, 95

disk drive module
 adding, 95-98
 antitamper lock
 disengaging, 96
 dual-ported, 66
 LEDs
 and troubleshooting, 90
 at power-on, 77
 replacing, 98-101
 required, 47
 specifications, 108
 states, 69-71
 total addressed by RAID controller, 7

Disk topologies
 RAID, 52

Disk topology, RAID, 52

documentation, other, xiv

door of rack, opening and closing, 31

drive carrier module
 antitamper lock
 disengaging, 100

drive roaming, 49
 and hot spare, 99

drive state reporting, 69-71

Dummy disk drive carrier module, 21

E

enclosure, 6-25
 components, 8-22
 expansion, 7
 system ID, 77
 height, 103
 power off procedure, 46
 RAID, 7
 in rack, 26
 system ID, 77

environmental
 device, 71
 drive state, 71
 requirements, 105

ESI
 and ESI/ops panel LEDs, 47-48

ESI/ops panel module
 and SES, 75
 LEDs
 and SES, 47-48
 and troubleshooting, 79-80

expansion enclosure. See enclosure, expansion.

F

fan
 increased noise level, 92
 slowing, 76, 91
 speeding up, 92

FRUs
 disk carrier module, 95

G

ghost drive, 72

grounding, 37

grounding, checking
 rack, 42

H

HBA, cabling to, 36

height
 enclosure, 103
 rack, 103

host
 cabling to, 36
 does not recognize storage system, 77

hot spare, 4, 49, 70, 99
 and availability, 63
 and drive roaming, 99
hot swap disk drive replacement, 95
hub
 cabling to, 36
 in rack, 29
humidity requirements, 105

I

ID

 selector switch
 and troubleshooting, 77
 system
 expansion enclosure, 77
 RAID module, 77

I/O module

 and loops, 66

J

JBOD, 60

 and availability, 65
 and capacity, 62
 and performance, 65

JBOD controller module

 LEDs
 and troubleshooting, 89

L

LED

 and troubleshooting, 78-91
 checking at power-on, 76
 disk drive module
 and troubleshooting, 90

ESI/ops panel module

 and ESI, 47-48
 and troubleshooting, 79-80

JBOD controller module

 and troubleshooting, 89

power supply/cooling module

 and troubleshooting, 84

RAID controller module

 and troubleshooting, 85

loop

 configuration, 65-68

LRC I/O module

 RAID, 14
 specifications, 107

LUN

 affinity and system drive, 69
 integrity and disk drive module failure, 99
 mapping, 69

M

 manuals, other, xiv

O

 online drive state, 70

Ops panel

 overview, 11
 switches and indicators, 11

P

 performance, configuring for maximum, 65

 physical specifications, 103

port

 RAID controller module
 to host, 36

- power, 46
 - checking, 39, 43
 - cord, 37-42
 - requirements, 106
 - PDU, 106
 - rack, 106
 - voltage requirement
 - rack, 40
 - tower, 37
- Power supplies. See PSU/cooling module.
- power supply/cooling module
 - LED
 - and troubleshooting, 84
- powering
 - off
 - enclosure in rack, 46
 - rack, 45
 - tower, 46
 - on
 - alarm sounds, 76
 - checking system status, 39, 43
 - problems, 76
 - rack, 40-44
 - tower, 37-40
- programmable LUN mapping, 69
- PSU/cooling module
 - overview, 12
- R**
- rack, 26-31
 - breaker, 43
 - cabling, 40-42
 - height, 103
 - in U, 26
 - power requirements, 106
 - powering on, 40-44
 - rear door, opening and closing, 31
- rackmount enclosure
 - bay numbering, 23
- RAID
 - disk topologies, 52
 - LRC I/O module, 14
- RAID controller module
 - and loops, 66
 - drives addressed, 7
 - LEDs
 - and troubleshooting, 85
 - write cache size, 64
- RAID enclosure, 7
 - in rack, 26
- RAID level
 - RAID 0, 60
 - and availability, 64
 - and capacity, 62
 - and disk failure, 99
 - and performance, 65
 - RAID 0+1, 60
 - and availability, 65
 - and capacity, 62
 - and disk failure, 99
 - and performance, 65
 - RAID 1, 60
 - and availability, 64
 - and capacity, 62
 - and disk failure, 99
 - and performance, 65
 - RAID 3, 60
 - and availability, 64
 - and capacity, 62
 - and disk failure, 99
 - and performance, 65
 - RAID 5, 60
 - and availability, 64
 - and capacity, 62
 - and disk failure, 99
 - and performance, 65
 - RAID 6. See RAID 0+1.

RAID 7, 60
 See also JBOD.
 strategy for selecting, 61-65
 supported, 60-61

rebuild
 automatic, 71

regulatory information, 109-111

S

SCS_DATA_OVERRUN error, 77

server does not recognize storage system, 77

service, xv

SES, 51
 and ESI/ops panel module, 75

SGI Fibre Channel Hub
 in rack, 29

SGI switch
 cabling to, 36
 in rack, 29

SGI, contacting, xv

Simplex RAID configuration, 52

slot numbering
 tower, 24

specifications
 disk drive module, 108
 LRC I/O module specifications, 107
 storage system physical, 103

starting storage system. See powering on.

stopping storage system. See powering off.

support, xv

switch
 cabling to, 36
 in rack, 29

system drive, 68-69
 and data caching, 68
 and LUN affinity, 69

maximum, 68
 size, 68

T

temperature
 requirements, 105

test mode, 93

thermal
 alarm, 93
 control, 92

Topology, RAID, 52

Tower
 power off procedure, 46

tower, 32-34
 adapting for rackmounting, 34
 bay numbering, 24
 cabling, 37-40
 powering on, 37-40

TP9100
 features, 1

U

unconfigured drive state (unconfigured location), 71

V

voltage, 106
 out of range, 76
 requirement
 rack, 40
 tower, 37

W

write cache

 disable, 51

 enable, 50, 64

 size, RAID controller, 64

write-back caching, 50, 64

 and system drive, 68

write-through caching, 51

 and system drive, 68