

**SGI™ Total Performance 9100
Storage System Owner's Guide**

007-4068-003

CONTRIBUTORS

Written by Carolyn Curtis and Mick Taylor

Illustrated by Dan Young, Carolyn Curtis, and Mick Taylor

Production by Glen Traefald

Engineering contributions by Marty Castilla, Tim Piper, Bill Andrews, Greg Schaem, Sammy Wilborn, and Jaideep Deshmukh

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

© 2000, Silicon Graphics, Inc. All Rights Reserved

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.

The contents of this document may not be copied or duplicated in any manner, in whole or in part, without the prior written permission of Silicon Graphics, Inc.

For regulatory information, see Appendix B.

Silicon Graphics and the Silicon Graphics logo are registered trademarks and SGI and IRIS InSight are trademarks of Silicon Graphics, Inc.

Xyratex is a registered trademark of Xyratex, Inc.

Mylex is a registered trademark of Mylex Corporation.

Intel is a registered trademark of Intel Corporation.

Windows and Windows NT are registered trademarks of Microsoft Corporation.

Record of Revision

Version	Description
001	November 1999 Original Printing
002	February 2000
003	June 2000

Contents

List of Figures	ix
List of Tables	xi
About This Guide.xiii
Audiencexiii
Structure of This Documentxiii
Related Publicationsxiv
Conventions Used in This Guide xv
Product Support xv
Reader Commentsxvi
1. Storage System Features	1
Storage System Formats	1
Storage System Enclosure	2
SGI TP 9100 Storage System Rack	5
Expansion Enclosure Cabling	8
SGI TP 9100 Storage System Tower.	9
Modules in the Storage System	11
Power Supply/Cooling Module	12
Enclosure System Interface/Operator (ESI/Ops) Panel Module	13
RAID Controller Module.	17
FC-AL Loop Resiliency Circuit (LRC) I/O Modules	19
Disk Drive Modules	23
Storage System Bay Numbering	26
Opening and Closing the Rack Door	28
2. Connecting to a Host and Powering On and Off.	29
Connecting to a Host	29
Grounding Issues	30

Connecting the Power Cord and Powering On the SGI TP 9100 Tower	31
Checking Grounding for the Tower	32
Powering On the SGI TP 9100 Tower	33
Checking AC Power and Storage System Status for the Tower	34
Cabling and Powering On the Rack	35
Checking Grounding for the Rack	37
Powering On the Rack	37
Checking AC Power and System Status for the Rack	39
Powering Off	40
Powering Off the SGI TP 9100 Rack	40
Powering Off the SGI TP 9100 Tower or a Single Enclosure	40
3. Features of the RAID Controller Module	43
SCSI Enclosure Services (SES) and Disk Drive Control	43
Configuration on Disk (COD).	44
Drive Roaming	45
Mylex Online RAID Expansion (MORE).	46
Disk Drive Requirements and Restrictions for MORE Options	46
Caveats for MORE Options	47
Add Logical Drive Option	48
Expand Logical Drive Option	48
Data Caching	49
Write Cache Enabled (Write-Back Cache Mode)	49
Write Cache Disabled (Write-Through or Conservative Cache Mode)	50
4. Using the RAID Controller	51
Software Tools for the Controller.	51
Global Array Manager (GAM)	52
Web Array Manager (WAM).	52
RAID Levels.	53
CAP Strategy for Selecting a RAID Level	55
Configuring for Maximum Capacity	56
Configuring for Maximum Availability.	57
Configuring for Maximum Performance	59

Loop Configuration	59
System Drives	64
System Drive Properties	64
System Drive Affinity and Programmable LUN Mapping	64
Drive State Reporting	65
Automatic Rebuild	66
5. Troubleshooting	69
Solving Initial Startup Problems	70
Using Storage System LEDs for Troubleshooting	71
ESI/Ops Panel LEDs	72
Power Supply/Cooling Module LED and Problem Solution	76
FC LRC I/O Module LEDs and Problem Solutions	77
RAID Controller LEDs and Problem Solutions	78
Disk Drive LEDs and Problem Solutions	79
Using the Alarm for Troubleshooting	80
Storage System Temperature Conditions, Problems, and Solutions	81
Thermal Control	81
Thermal Alarm	82
Thermal Shutdown	82
Using Test Mode	83
6. Installing and Replacing Disk Drive Modules	85
Adding a Disk Drive Module	86
Replacing a Disk Drive Module	96
LUN Integrity and Disk Drive Module Failure	96
Replacing the Disk Drive Module	97
A. Technical Specifications	99
Storage System Physical Specifications	99
Environmental Requirements	101
Power Requirements	102
RAID Controller Module Specifications	103
Disk Drive Module Specifications	104
SGI Cables for the SGI TP 9100 Storage System	104

- B. Regulatory Information 105**
 - Class A Warnings and Specifications for the SGI TP 9100 Rack 105
 - FCC Warning 105
 - Attention 106
 - European Union Statement 106
 - International Special Committee on Radio Interference (CISPR). 106
 - Canadian Department of Communications Statement 106
 - Attention 107
 - VCCI Class 1 Statement 107
 - Class A Warning for Taiwan 107
 - Class B Warnings and Specifications: SGI TP 9100 Tower and Enclosures 107
 - Manufacturer’s Declaration of Conformity 108
 - FCC Class B Conformance 108
 - Industry Canada Notice (Canada Only) 110
 - CE Notice 110
 - VCCI Notice (Japan Only) 110
 - Korean Regulatory Notice 110
 - NOM 024 Information (Mexico Only) 111
 - EN 55022 Compliance (Czech Republic) 111
 - Polish Notice (Poland Only) 112
- Index. 113**

List of Figures

Figure 1-1	Enclosure (Rackmount), Front View	2
Figure 1-2	Rackmount Enclosure (single RAID, or Base), Rear View	3
Figure 1-3	Rackmount Enclosure (dual RAID, or Base), Rear View	4
Figure 1-4	SGI TP 9100 Rack, Front View, Example.	5
Figure 1-5	SGI TP 9100 Rack, Rear View, Example	6
Figure 1-6	Expansion Enclosure Cabling	8
Figure 1-7	Tower (Deskside), Front View	9
Figure 1-8	Tower, Rear View	10
Figure 1-9	Power Supply/Cooling Module	12
Figure 1-10	ESI/Ops Panel Module in Enclosure.	13
Figure 1-11	ESI/Ops Panel Indicators and Switches: Rackmount Enclosure	14
Figure 1-12	ESI/Ops Panel Indicators and Switches: Tower	14
Figure 1-13	RAID Controller Module Locations	17
Figure 1-14	RAID Controller Panel Features: Rackmount Enclosure	18
Figure 1-15	RAID Controller Panel Features: Tower	19
Figure 1-16	I/O Module Location, Rackmountable Enclosure	20
Figure 1-17	I/O Module Panel Features, Rackmount Enclosure	20
Figure 1-18	I/O Module Panel Features and Location, Tower	22
Figure 1-19	Disk Drive Module.	23
Figure 1-20	Required Disk Drive Module Locations	25
Figure 1-21	Module Locations and Bay Numbering: Rackmount Enclosure	26
Figure 1-22	Module Locations and Bay Numbering: Tower.	27
Figure 1-23	Opening the Rack Rear Door: Latch	28
Figure 2-1	Bifurcated Power Cord for the Tower	31
Figure 2-2	ESI/Ops Panel, Tower: Powering On	33
Figure 2-3	ESI/Ops Panel Indicators and Switches: Tower	34
Figure 2-4	Rack Power Cabling	36

Figure 2-5 ESI/Ops Panel, Rackmount: Powering On 38

Figure 2-6 ESI/Ops Panel Indicators and Switches: Rackmount Enclosure . . 39

Figure 4-1 RAID Levels Within a Drive Pack (LUN): Example 55

Figure 4-2 I/O Modules, Channels, and Loops: Tower 60

Figure 4-3 I/O Modules, Channels, and Loops: Rackmount Enclosure
(Front View)61

Figure 4-4 Single-Controller/Point-to-Point/Single-Host Example 62

Figure 4-5 Dual-Active Controller/Point-to-Point/Dual-Host Example . . . 63

Figure 5-1 ESI/Ops Panel Indicators and Switches: Rackmount Enclosure . . 72

Figure 5-2 ESI/Ops Panel Indicators and Switches: Tower 73

Figure 5-3 Power Supply/Cooling Module LED. 76

Figure 5-4 FC LRC I/O Module LEDs: Rackmount Enclosure 77

Figure 5-5 FC LRC I/O Module LEDs: Tower 77

Figure 5-6 RAID Controller Panel LEDs: Rackmount Enclosure 78

Figure 5-7 RAID Controller Panel LEDs: Tower 78

Figure 5-8 Disk Drive Module LEDs 79

Figure 6-1 Drive Installation Order 86

Figure 6-2 Unlocking the Disk Drive Module (Disengaging the
Antitamper Lock) 87

Figure 6-3 Releasing the Module Handle 88

Figure 6-4 Inserting the Disk Drive Module: Rackmount Enclosure 89

Figure 6-5 Inserting the Disk Drive Module: Tower 90

Figure 6-6 Drive Seated: Rackmount Enclosure 91

Figure 6-7 Drive Seated: Tower 92

Figure 6-8 Engaging the Camming Lever: Rackmount Enclosure 93

Figure 6-9 Engaging the Camming Lever: Tower 94

Figure 6-10 Locking the Disk Drive Module (Setting the Antitamper Lock) . . 95

List of Tables

Table 1-1	Storage System LED Indications	15
Table 4-1	Supported RAID Levels	53
Table 4-2	RAID Level Maximum Capacity	56
Table 4-3	Array Operating Conditions	57
Table 4-4	RAID Levels and Availability	58
Table 4-5	RAID Levels and Performance	59
Table 4-6	Physical Disk Drive States	65
Table 5-1	ESI/Ops Panel LEDs and Troubleshooting	74
Table 5-2	Disk Drive LED Function	80
Table A-1	Physical Specifications, Rackmount Enclosures and Rack	99
Table A-2	Physical Specifications, Tower	100
Table A-3	Module Weights	100
Table A-4	Ambient Temperature and Humidity Requirements	101
Table A-5	Environmental Requirements	101
Table A-6	Minimum Power Requirements	102
Table A-7	Rack PDU Power Specifications	102
Table A-8	RAID Module Specifications	103
Table A-9	Drive Carrier Module Specifications (1.6-Inch 18 GB Drive)	104
Table A-10	SGI Fibre Channel Fabric Cabling Options for SGI TP 9100 Storage System	104

About This Guide

This guide explains how to operate and maintain the SGI Total Performance 9100 fibre channel storage system. 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 SGI Total Performance (TP) 9100 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 owners and users of the SGI TP 9100 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 Features,” 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 Module,” describes SCSI Enclosure Services, configuration on disk, drive roaming, Mylex Online RAID Expansion, 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 Disk Drive 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 and B 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, have available the latest versions of the owner’s guide for the server. Also have available the owner’s guides 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, choose Help > Online Books > SGI EndUser or SGI Admin, and select the applicable owner’s or hardware guide.
- Technical Publications Library: if you have access to the Internet, enter the following URL in your Web browser location window:
<http://techpubs.sgi.com/library/>

Conventions Used in This Guide

The following conventions are used throughout this guide:

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.
user input	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 tell us. Be sure to include the title and document number of the manual with your comments. (Online, the document number is located in the front matter of the manual. In printed manuals, the document number can be found on the back cover.)

You can contact us in any of the following ways:

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

We value your comments and will respond to them promptly.

Storage System Features

This chapter gives information on the basic features of the SGI Total Performance (TP) 9100 fibre channel storage system, in these sections:

- “Storage System Formats” on page 1
- “Modules in the Storage System” on page 11
- “Storage System Bay Numbering” on page 26
- “Opening and Closing the Rack Door” on page 28

Storage System Formats

The storage system is available in the formats described in these subsections:

- “Storage System Enclosure” on page 2
- “SGI TP 9100 Storage System Rack” on page 5
- “Expansion Enclosure Cabling” on page 8
- “SGI TP 9100 Storage System Tower” on page 9

The storage system can be connected to one or two fibre channel boards (host bus adapters, or HBAs) in a supported SGI server separately or in combination (loop).

Storage System Enclosure

The basic unit of the SGI TP 9100 storage system is the *enclosure*, which contains a minimum of 5 and maximum of 12 disk drives and the component modules that handle RAID control, I/O, power and cooling, and operations. Figure 1-1 shows the front of the enclosure, with the full complement of 12 disk drive modules.




Figure 1-1 Enclosure (Rackmount), Front View

The enclosure is available with or without single or dual RAID modules. An enclosure with single or dual RAID modules is a *RAID (base) enclosure*. An enclosure without a RAID module is an *expansion enclosure*; The expansion enclosure must be cabled to a RAID enclosure and provides additional disk modules. The RAID controller can address up to 32 disk drives; thus, two expansion enclosures can be cabled to it.

Figure 1-2 shows the back of a rackmountable RAID enclosure with a single RAID controller (RAID controller 0) and points out the component modules.




Figure 1-2 Rackmount Enclosure (single RAID, or Base), Rear View

Figure 1-3 shows the back of a rackmountable RAID enclosure with dual RAID controllers (RAID controller 0 and RAID controller 1) and points out the component modules.




Figure 1-3 Rackmount Enclosure (dual RAID, or Base), Rear View

Dummy modules occupy unoccupied empty bays and must be in place for proper operation and to ensure correct airflow. The component modules are discussed in detail in “Modules in the Storage System” on page 11.

The RAID enclosure can be housed in a standalone tower or in an SGI TP 9100 rack (or in a standard 19-inch rack). Each enclosure occupies 4U (4 EIA units) of rack space.

SGI TP 9100 Storage System Rack

Figure 1-4 shows the front of an SGI TP 9100 rack with two enclosures.




Figure 1-4 SGI TP 9100 Rack, Front View, Example

Figure 1-5 is a rear view of an SGI TP 9100 rack with three RAID enclosures.




Figure 1-5 SGI TP 9100 Rack, Rear View, Example

For instructions on opening the rear door of the rack, see “Opening and Closing the Rack Door” on page 28.

The SGI TP 9100 rack ships with at least two enclosures: either two RAID enclosures or one RAID enclosure and one expansion enclosure. At the front of the rack, filler panels cover unoccupied bays; these must be in place for proper operation and to ensure correct airflow.

The SGI 9100 rack is 38U high, with ten bays. The power distribution units (PDUs) at the sides of the rack provide power for the bays. 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 bays and sockets accommodate the following:

- Socket 1 at the top of each PDU is for the 2U bay at the top of the rack, which accommodates the SGI Fibre Channel Hub or one or more fibre channel switches.
If this bay is unoccupied, a 2U filler panel provides correct airflow.
- Sockets 2 through 10 on each PDU are for the nine 4U bays, which accommodate SGI TP 9100 enclosures.

All nine bays can be occupied by independent RAID enclosures, or by any combination of a RAID enclosure and one or two expansion enclosures cabled to each RAID enclosure. For unoccupied bays, a 4U filler panel provides proper operation and correct airflow.

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



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.

Expansion Enclosure Cabling

Figure 1-6 shows cabling between a RAID (base) enclosure and two expansion enclosures. Enclosures are cabled at the factory, or by an SGI System Support Engineer (SSE).




Figure 1-6 Expansion Enclosure Cabling

SGI TP 9100 Storage System Tower

The tower (deskside) version of the storage system houses one RAID enclosure. Figure 1-7 shows the front of the tower.




Figure 1-7 Tower (Deskside), Front View

Figure 1-8 is a rear view of the tower.




Figure 1-8 Tower, Rear View

Dummy modules occupy empty bays and must be in place for proper operation and to ensure correct airflow. The component modules are discussed in detail in “Modules in the Storage System” on page 11.

The enclosure in the tower can be adapted for rackmounting; contact your service provider if you wish this done.

Modules in the Storage System

The enclosure chassis consists of two segmented aluminum assemblies with a midplane PCB between them. Each chassis assembly contains 12 bays; a bay is defined as the space required to house a single 3.5-inch disk drive (1.6 inches high, or 1 inch high with a foam filler) in its carrier module. Larger modules—the power supplies—are accommodated in multiple bay spaces. The midplane PCB provides logic level signal and low-voltage power distribution paths.

Component modules in the storage system are

- Two power supply/cooling modules
- Enclosure system interface/operator panel (ESI/ops)
- One or two RAID controller modules with FC-AL input/output
- Two loop resiliency circuit (LRC) I/O modules
- FC-AL disk drive carrier modules:
 - RAID enclosure: 5 to 12
 - expansion enclosure: 2 to 12

The modules are described in separate sections in this chapter:

- “Power Supply/Cooling Module” on page 12
- “Enclosure System Interface/Operator (ESI/Ops) Panel Module” on page 13
- “RAID Controller Module” on page 17
- “FC-AL Loop Resiliency Circuit (LRC) I/O Modules” on page 19
- “Disk Drive Modules” on page 23

This chapter also includes a section explaining chassis bay numbering: “Storage System Bay Numbering” on page 26.

Power Supply/Cooling Module

Two 550-watt power supply/cooling modules are mounted in the rear of the storage system. Figure 1-9 shows a module. Each module contains two fans as well as the power supply and its associated electronics. The fans derive power from the chassis midplane, not from the power supply itself.




Figure 1-9 Power Supply/Cooling Module

Power supply voltage operating ranges are nominally 115V or 230V, selected automatically.

An LED on the front panel of the power supply/cooling module indicates the status of the power supply and the fan.

The modules operate together; if one fails, the other maintains the power supply and cooling while the faulty unit is replaced. The faulty module, however, still provides proper airflow for the storage system.

Power cords are included with the storage systems; they are described in “Connecting the Power Cord and Powering On the SGI TP 9100 Tower” on page 31 and “Powering Off” on page 40 in Chapter 2. Specifications for the cables and for the power supply are in Table A-7 on page 102 in Appendix A.

Enclosure System Interface/Operator (ESI/Ops) Panel Module

The ESI/ops panel module provides the storage system with a microcontroller for monitoring and controlling all elements of the storage system. Each element (power, cooling, temperature, device status) is interfaced to the processor using an I²C (I square C) bus.

Figure 1-10 shows the location of the ESI/ops panel in a rackmountable enclosure.




Figure 1-10 ESI/Ops Panel Module in Enclosure

Figure 1-8 shows its location in the tower.

Figure 1-11 shows details of the panel in a rackmounted enclosure.




Figure 1-11 ESI/Ops Panel Indicators and Switches: Rackmount Enclosure

Figure 1-12 shows details of the panel in a tower.




Figure 1-12 ESI/Ops Panel Indicators and Switches: Tower

The ESI/ops panel shows a consolidated status for all modules. Table 1-1 summarizes the function of the LEDs on the ESI/ops panel, as well as those on other modules in the storage system.

Table 1-1 Storage System LED Indications

LED	Meaning
Green, steady	Positive condition
Green, flashing	Disk drives: data activity
Alternating green and amber, flashing	Noncritical condition
Amber, steady	Fault

In addition to the indicators on the ESI/ops panel, each module type has its own status LED(s). Using the LEDs for troubleshooting is explained in “Using Storage System LEDs for Troubleshooting” on page 71 in Chapter 5.

Other features of the ESI/ops panel are as follows:

- Key switch: activates run or service mode, as explained in “Key Switch for Starting the Storage System.”
- Standby pushbutton: activates standby mode when the key switch is in the Service position; see “Powering On the Rack” on page 37 and “Powering On the SGI TP 9100 Tower” on page 33 in Chapter 2.
- Alarm mute pushbutton: mutes the alarm; see “Audible Alarm” on page 16.

The ESI/ops panel module firmware includes SCSI Enclosure Services (SES), which monitors the modules in the storage system and controls the ESI/ops panel LEDs. The ESI/ops panel requires two disk drives in specific drive bays to serve as conduits for information from the storage system to the ESI/ops panel module. SES is also required for the RAID controller module(s); see “SCSI Enclosure Services (SES) and Disk Drive Control” on page 43 in Chapter 3 for more information.

Key Switch for Starting the Storage System

For general use, the storage system is started by turning the key switch on the far right (top) of the ESI/ops panel to the **RUN** position; the storage system ships with the key in this position. For specific instructions on starting and stopping the storage system, see

“Powering On the Rack” on page 37 and “Powering On the SGI TP 9100 Tower” on page 33 in Chapter 2.

Audible Alarm

The audible alarm on the ESI/ops panel sounds when a fault state is present. Pressing the alarm mute pushbutton reduces the volume of the alarm, but leaves a beep at approximately 10-second intervals to indicate that a fault state is still present. The mute pushbutton is beneath the indicators on the ESI/ops panel (see Figure 1-11 and Figure 1-12).

“Using the Alarm for Troubleshooting” on page 80 in Chapter 5 explains alarm function.

RAID Controller Module

Each storage system contains one or two RAID controllers, which bring RAID functionality and high-speed fibre data transfer performance for connection to a fibre channel host bus adapter (HBA) installed in a host system. Figure 1-13 shows an enclosure with dual RAID controllers. The enclosure can also accommodate a single RAID controller. Each RAID controller has a GBIC port that can accommodate SGI supported GBIC modules.




Figure 1-13 RAID Controller Module Locations

The first controller (RAID controller 0) must be installed in bay 3/2 and the second in bay 2/2 (see “Storage System Bay Numbering” on page 26 for details on bay numbering). When the controllers are configured in a dual-active configuration, they provide continuous access to data, even if one of them fails. An installation with a single controller is known as a simplex system; an installation with two controllers is known as a dual-active system.

The intelligent caching controller supports industry-standard RAID levels (0, 1, 3, 5, and 0+1) for multiple-drive arrays, and JBOD for single-drive control functionality. The controller also supports RAID 10, 30, and 50 in spanned disk arrays. (“RAID Levels” on page 53 in Chapter 4 has information on supported RAID levels.)

The RAID controller is shipped with the cache backup battery disconnected. Before the RAID controller is installed in the enclosure, connect the battery, then use the GAM or WAM utility to fast charge the battery.

LEDs on the controller panel give information on controller activity; see Figure 1-14 (tower) and Figure 1-15 (rackmount enclosure). The fault/controller not ready LED is amber; the other three LEDs are green.




Figure 1-14 RAID Controller Panel Features: Rackmount Enclosure

The RAID controller’s activity LED flashes every time the Global Array Management (GAM) server issues a poll. The default interval is 7 seconds; to reset the interval, see the documentation for the GAM software.

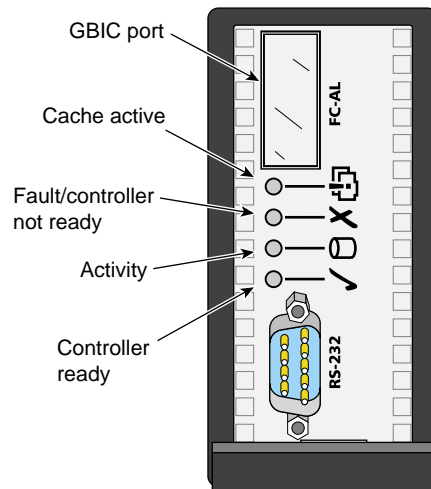


Figure 1-15 RAID Controller Panel Features: Tower

A fibre channel device, such as a host bus adapter in a server, is cabled to the FC-AL (GBIC) port on the RAID controller.

For more information on the controller, see Chapter 3 and Chapter 4.

FC-AL Loop Resiliency Circuit (LRC) I/O Modules

The storage system is shipped with two FC-AL loop resiliency circuit (LRC) I/O modules, which make dual (redundant) data paths possible. Each module provides connection between the RAID controller and one loop of up to 12 drives.

Cables between I/O modules must have a male DB9 connector on one end to plug into the mating DB9 connector on the left side of the I/O module and an HDSSC connector on the other end to plug into the GBIC connector on the right side of the I/O module.

Figure 1-16 shows the location of the I/O modules in the enclosure.




Figure 1-16 I/O Module Location, Rackmountable Enclosure

Figure 1-17 shows features of the I/O module panel in the enclosure.




Figure 1-17 I/O Module Panel Features, Rackmount Enclosure

Each port on the I/O module panel has an LED indicating its status.

Each I/O module has an ID selector switch for the system ID, which is set to 0 on both modules at the factory. Disk drive IDs are based on this ID.

- For the tower, this ID is always 0 on both I/O modules.
- In the rack:
 - the RAID (base) enclosure is set to 0 on both its I/O modules
 - the first expansion enclosure cabled to a base enclosure is set to 1 on both its I/O modules
 - the second expansion enclosure cabled to a base enclosure is set to 2 on both its I/O modules

At initial startup, each I/O module controls a separate loop of six drives; see “Loop Configuration” on page 59 in Chapter 4 for details.

Figure 1-18 shows features and location of the I/O module panel in a tower.




Figure 1-18 I/O Module Panel Features and Location, Tower

Disk Drive Modules

A disk drive module (FC-AL) consists of a hard disk mounted in a carrier. Each drive bay in the front of the storage system houses a single 1.6-inch high, 3.5-inch disk drive in its carrier, or a 1-inch high drive, which has a foam filler. Figure 1-19 shows the disk drive module.




Figure 1-19 Disk Drive Module

The label on the front of the drive has the drive capacity (for example, **FCAL 9.17GB**), the serial number, and worldwide name. The serial number of the drive is also on a label on the top of the carrier.

The SGI TP 9100 storage system uses disk technology that allows you to replace a disk while the system continues to run. This “hot swapping” must be done only under controlled circumstances; “Replacing a Disk Drive Module” on page 96 in Chapter 6 has more information.

The drives in certain bays (left top and left bottom bays of the tower; bottom far left and far right of the rackmount enclosure) are required for storage system management; see

“Required Drives” on page 25 and “Storage System Bay Numbering” on page 26 for specific information on their location.

The drives are dual-ported and can be configured as two separate loops; see “Loop Configuration” on page 59 in Chapter 4.

Disk Drive Carrier

The extruded aluminum carrier (canister) provides thermal conduction and radio frequency and electromagnetic induction protection, and affords the drive maximum physical protection.

The drive module cap has an ergonomic handle that provides the following functions:

- Camming of the carrier into and out of the drive bays
- Positive spring loading of the drive/backplane connector

Disk Drive LEDs

Each drive carrier has two LEDs, an upper (green) and lower (amber). In normal operation, the green LED is on and flickers as the drive operates. The amber LED illuminates when the drive is faulty.

Disk Drive Antitamper Lock

The drive carrier has an antitamper lock that disables the normal latch action of the handle; it is fitted in the drive carrier handle, as shown in Figure 1-19. The lock is set through the hole in the lower part of the handle trim. An indicator shows the setting:

- When the drive is locked, a red indicator is visible in the center rectangular aperture in the handle.
- When the drive is unlocked, a black indicator is visible.

The antitamper lock setting is changed with a key through the small round cutout in the lower part of the handle trim piece. (A key is included with each tower.) Changing the setting is explained in “Adding a Disk Drive Module” on page 86 and in “Replacing a Disk Drive Module” on page 96 in Chapter 6.

Dummy Disk Drives (Drive Fillers)

Drive fillers or dummy drives (front dummy fascias) are provided for all unused drive bays and are required to maintain a balanced airflow.

Required Drives

Figure 1-20 shows the location of the drives required by the ESI/ops panel and the RAID controller(s), namely drive bays 1/3 and 4/3. "Storage System Bay Numbering" on page 26 has full details of bay numbering.

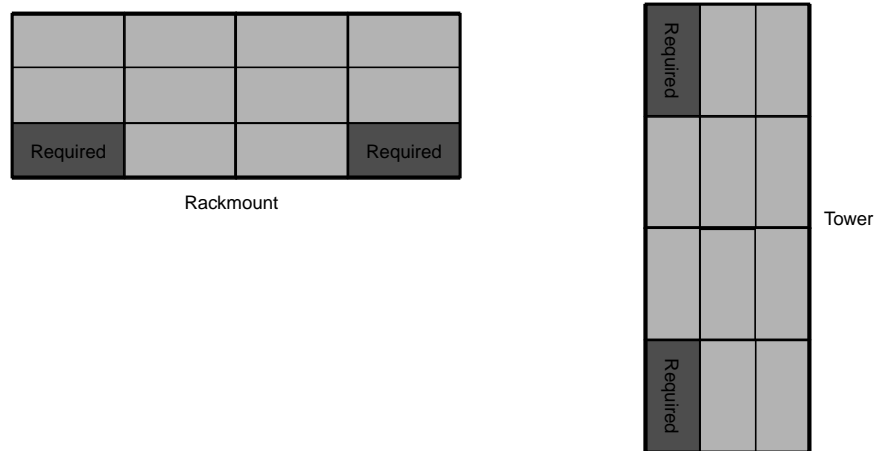


Figure 1-20 Required Disk Drive Module Locations

Storage System Bay Numbering

Figure 1-21 shows the correct positions of the modules and the enclosure bay numbering convention for the rackmount enclosure. A bay is defined as the space required to house a single 3.5-inch disk drive in its carrier module.

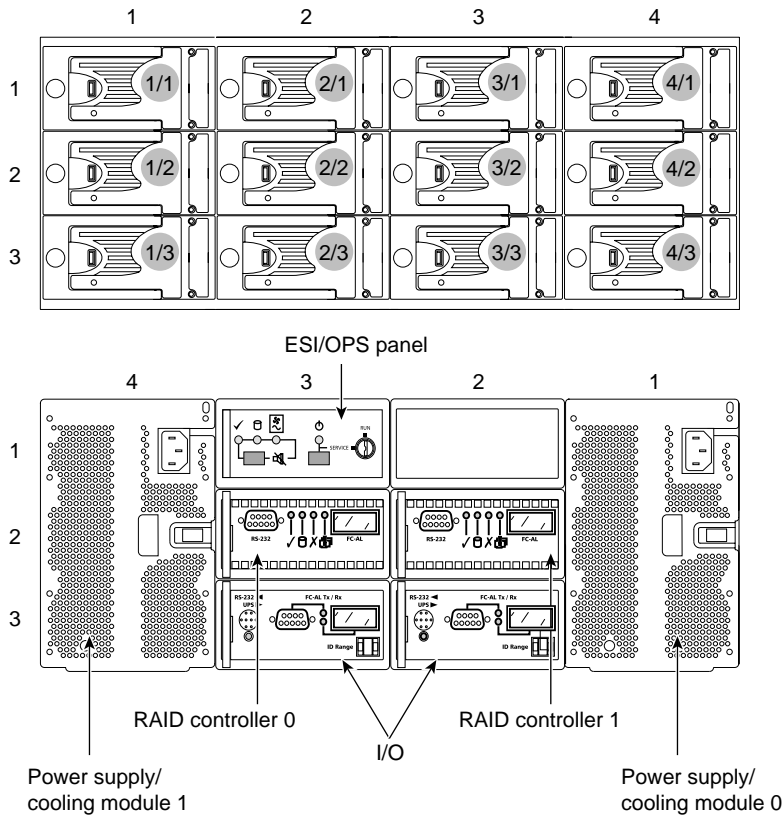


Figure 1-21 Module Locations and Bay Numbering: Rackmount Enclosure

The rackmount enclosure is 4 x 3, that is, 4 bays wide by 3 bays high.

- The disk drive bays, located in front, are numbered 1 to 4 from left to right and 1 to 3 from top to bottom. Drives in bays 1/3 and 4/3 are required for storage system management; these bays must always be occupied.
- The rear bays are numbered 1 to 4 from right to left and 1 to 3 from top to bottom.

Module locations are identified by combining the column and row numbers (top and side numbers in Figure 1-22). For example, the ESI/ops panel on the rackmountable enclosure is in rear bay 3/1 (third column from right, top row).

Figure 1-22 shows the correct positions of the modules and the enclosure bay numbering convention for the tower.

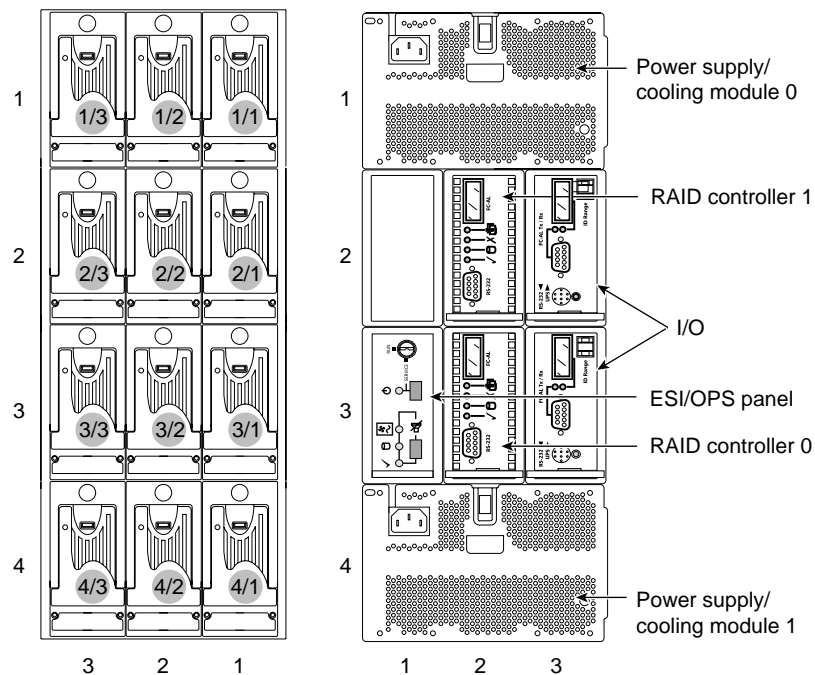


Figure 1-22 Module Locations and Bay Numbering: Tower

The tower is 3 x 4, that is, 3 bays wide by 4 bays high.

- The disk drive bays, located in front, are numbered 1 to 3 from right to left and 1 to 4 from top to bottom. Drives in bays 1/3 and 4/3 are required for storage system management; these bays must always be occupied.
- The rear bays are numbered 1 to 3 from left to right and 1 to 4 from top to bottom.

Module locations are identified by combining the column and row numbers (top and side numbers in Figure 1-22). For example, the ESI/ops panel on the tower is in rear bay 3/1 (leftmost column, third row).

Opening and Closing the Rack Door

To open the rack rear door, follow these steps:

1. Locate the latch on the rack rear door.
2. Push up the top part of the latch, as shown in the second panel of Figure 1-23.




Figure 1-23 Opening the Rack Rear Door: Latch

3. Press the button, as shown in the third panel of Figure 1-23. 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 1-23. The door opens.

To close the door, lift the locking brace at the bottom.

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 29
- “Grounding Issues” on page 30
- “Connecting the Power Cord and Powering On the SGI TP 9100 Tower” on page 31
- “Cabling and Powering On the Rack” on page 35
- “Powering Off” on page 40

Note: For instructions on opening the rear door of the rack, see “Opening and Closing the Rack Door” on page 28 in Chapter 1.

Connecting to a Host

Physical attachment to the host is from the FC-AL DB-9 port on the RAID controller module to an FC-AL port in a host.

Besides cabling directly to an HBA in a host, you can connect the storage system to an SGI Fibre Channel Hub using fibre channel copper cabling with DB-9 connectors at each end, or to an SGI fibre channel 8-port or 16-port switch (using a copper cable and a copper GBIC). See Table A-10 in Appendix A for information on these cables.

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

For more information on configurations, see “Loop Configuration” on page 59 in Chapter 4.

Grounding Issues

Each chassis—storage or host—must be well grounded through its power connector. All chassis connected with XIO copper cables must share the same transformer, must be grounded through the same earthing rod, and must be on the same branch circuit.

If you have any doubts about the quality of the ground connection, consult with a qualified electrician. Using optical cabling between the fibre enclosure(s) and the host XIO connection eliminates any problems related to common grounding.

Caution: Any difference in ground potential greater than 500 millivolts (0.5 volts) between two chassis connected with copper XIO cables can cause severe equipment damage and can create hazardous conditions.

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.

Copper cabling is advised over optical cabling for the following situations:

- The fibre channel switch is inside the fibre channel storage rack.
- The fibre channel switch is next to storage rack and uses the same ground.
- The switch is next to CPU and uses the same ground.



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 Cord and Powering On the SGI TP 9100 Tower

The tower requires 115-220 volts (autoranging), and is shipped with the following:

- One bifurcated power cord, shown in Figure 2-1; this cable gives you the advantage of a single point of disconnect.




Figure 2-1 Bifurcated Power Cord for the Tower

The bifurcated power cord has dual IEC 320 receptacles and an IEC plug inline cord extension.

- Two country-specific power cords:
 - Attach one at the free end of the bifurcated power cord to connect to a power source.or
 - Attach one to each power supply/cooling module power socket instead of using the bifurcated power cord, and connect the other end of each cord to a power source.
- Two IEC 320 power cords for cabling to the rack power distribution units (PDUs)
These are included in case the tower is converted to a rackmountable enclosure.

When you attach the power cord to a power source, the tower powers on automatically.

Caution: Use power cords supplied with the storage system or ones that match the specification quoted in Table A-7 on page 102 in Appendix A.

Country-specific power cords are available from SGI that are appropriate to local standards.

Checking Grounding for the Tower

To ensure that a safe grounding system is provided for a tower, follow these steps:

1. Make sure that the power cord is connected to the storage system, but 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.

2. Check for continuity between the earth pin of the power cord and any exposed metal surface of the storage system.

Powering On the SGI TP 9100 Tower

Follow these steps to power on the storage system:

1. Ensure that all modules are firmly seated in their correct bays and that blank plates are fitted in any empty bays.
2. Make sure 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.
3. For an unattended installation where the storage system is required to power on whenever AC power is present, the **RUN/SERVICE** key switch should be in the **RUN** position. The storage system ships from the factory with this setting; see Figure 2-2.




Figure 2-2 ESI/Ops Panel, Tower: Powering On

If the **RUN/SERVICE** key switch is in the **RUN** position, the storage system powers on when the power cord is plugged into a power source.

If the **RUN/SERVICE** key switch is in the **RUN** position and power is interrupted, the storage system returns to the power-on condition when power is restored.

In run mode, the standby pushbutton is disabled.

Note: Setting the **RUN/SERVICE** key switch on the ESI/ops panel to the **SERVICE** position puts the storage system in service mode; this mode is used by SGI System Support Engineers only. In service mode, if power is lost for any reason, the storage system defaults to standby mode when power is restored. Press the standby pushbutton to power on the storage system.

Checking AC Power and Storage System Status for the Tower

The standby LED on the ESI/ops panel (see Figure 2-3) turns green if AC power is present.




Figure 2-3 ESI/Ops Panel Indicators and Switches: Tower

At power-on, check ESI/ops panel LEDs for system status. Under normal conditions, these LEDs should all be illuminated constant green. If a problem is detected, the ESI processor in the operator panel changes the relevant LED from green to amber. See

“Solving Initial Startup Problems” on page 70 and “Using Storage System LEDs for Troubleshooting” on page 71 in Chapter 5.

Other modules in the storage system also have LEDs, which are described in “Using Storage System LEDs for Troubleshooting” on page 71 in Chapter 5.

Cabling and Powering On the Rack

The rack requires 220 volts 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 2-4 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 sets up the rack and cables its power. The information in this section is provided for your information and for safety reasons.

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 System Support Engineers 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 2-4 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 30.
2. For the grounding check, make sure 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. Make sure 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. Make sure 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. Make sure 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. For an unattended installation where the storage system is required to power on whenever AC power is present, the **RUN/SERVICE** key switch on each ESI/ops panel of each enclosure you are operating should be in the **RUN** position. The enclosures ship from the factory with this setting; see Figure 2-5.




Figure 2-5 ESI/Ops Panel, Rackmount: Powering On

If the **RUN/SERVICE** key switch is in the **RUN** position, the enclosure powers on when power reaches the enclosure.

If the **RUN/SERVICE** key switch is in the **RUN** position and power is interrupted, the enclosure returns to the power-on condition when power is restored.

In run mode, the standby pushbutton is disabled.

Note: Setting the **RUN/SERVICE** key switch on the ESI/ops panel to the **SERVICE** position puts the storage system in service mode; this mode is used by SGI System Support Engineers only. In service mode, if power is lost for any reason, the storage system defaults to standby mode when power is restored. Press the standby pushbutton to power on the storage system.

4. If necessary, connect each PDU power cord to a power source. The PDU power cords can be routed through an aperture at the top or the bottom of the rack. See Figure 2-4.



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.

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

5. Press the rack breaker switch at the bottom of each PDU so that the word **ON** shows.

Checking AC Power and System Status for the Rack

When you power on the system, the standby LED on each ESI/ops panel (see Figure 2-6) in each enclosure you are operating should turn green. If it does not, check that the power supply/cooling modules in the enclosure are correctly cabled to the rack PDUs.




Figure 2-6 ESI/Ops Panel Indicators and Switches: Rackmount Enclosure

At power-on, check ESI/ops panel LEDs for system status. Under normal conditions, these LEDs should all be illuminated constant green. If a problem is detected, the ESI processor in the operator panel changes the relevant LED from green to amber. See “Solving Initial Startup Problems” on page 70 and “Using Storage System LEDs for Troubleshooting” on page 71 in Chapter 5.

Other modules in the storage system also have LEDs, which are described in “Using Storage System LEDs for Troubleshooting” on page 71 in Chapter 5.

Powering Off

Powering off the SGI TP 9100 rack and the tower are explained in separate subsections. You power off a single enclosure in a rack the same way as you do a tower.

Powering Off the SGI TP 9100 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 three enclosures and leave others powered on. Figure 2-4 points out their location.

To power off the entire rack, follow these steps:

1. Make sure users are off the affected systems.
2. Push down the main breaker switch at the bottom of each PDU so that the word **OFF** shows.
3. If appropriate, disconnect the PDU power cords from the power sources.

Powering Off the SGI TP 9100 Tower or a Single Enclosure

The storage system toggles between standby mode (all drives are powered off) and run (operation) mode whenever the standby pushbutton is pressed.

After making sure users are off the affected systems, power off the tower or enclosure in one of two ways:

- Turn the key switch to the **SERVICE** position and then press the standby pushbutton for 2 seconds.

To power on the tower again from this setting, press the standby pushbutton for 2 seconds again.

- Remove power at the power source: unplug the tower, or detach the enclosure's power cord from the rack PDU.

Note the following:

- If the key switch on the ESI/ops panel is in the **RUN** position when power is restored to the storage system, the system powers on.
- If the key switch is in the **SERVICE** position when power is restored to the storage system, the system does not power on.

Features of the RAID Controller Module

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

- “SCSI Enclosure Services (SES) and Disk Drive Control” on page 43
- “Configuration on Disk (COD)” on page 44
- “Drive Roaming” on page 45
- “Mylex Online RAID Expansion (MORE)” on page 46
- “Data Caching” on page 49

SCSI Enclosure Services (SES) and Disk Drive Control

The RAID controllers use SCSI-3 SCSI Enclosure Services (SES) commands to manage the physical storage system. SES 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 1/3 and 4/3; these drives must be present for the system to function. See Figure 1-22 on page 27 in Chapter 1 for diagrams of their location.

SES access is 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 1/3 and 4/3, and are passed through to the SES device.

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

- Disk drive insertion status
- Power supply status
- Cooling element status
- Storage system temperature

The LEDs on the ESI/ops panel show this status.

Configuration on Disk (COD)

Configuration on disk (COD) retains the latest version of the saved configuration at a reserved location on each disk drive in the storage system: the same information is written to a reserved location of the NVRAM of the controller(s).

The configuration information includes the configuration ID number and the configuration sequence number. Both numbers are reset to 0 when the configuration is changed. The configuration sequence number increments whenever there is a physical disk drive state change. This feature allows the configuration to be maintained through a controller replacement.

COD plays a significant role during the power-up sequence after a controller is replaced. The replacement controller progresses through a series of checks and validation steps to determine the validity of any configuration currently present in NVRAM, followed by additional checks and validation steps to determine the validity of the COD information on all disk drives in the storage system. The final configuration is determined by one of the following criteria:

- The majority of COD information with the same configuration ID
- The COD available on the disk drive (when there is one disk drive and one controller in the storage system)
- The first configuration encountered (when there is an equal number of valid configurations)

At startup, the RAID controller checks if the storage system configuration matches that stored on disk (see “Configuration on Disk (COD)”); if it does not match, the storage system takes steps to bring it into conformity with the stored configuration.

Drive Roaming

Drive roaming allows disk drives to be moved to other channel/target ID locations while the storage system is powered off. Drive roaming allows for easier disassembly and assembly of systems; it also enhances performance by optimizing channel usage.

Drive roaming uses the Configuration on Disk (COD) information stored on the physical disk drive and in the storage system drive configuration structure. When the storage system reboots, a table is generated for each disk drive’s current location and its location before the storage system was powered off. The controller uses this information to remap the physical disk drives into the proper location in the system, or logical, drive.

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.

Drive roaming is designed for use within one system environment, for example, a single storage system or a cluster of systems sharing a single-controller or a dual-active controller configuration.

Note: If the COD information on a replacement disk drive is questionable or invalid, the disk drive is labeled dead. Once a drive is marked as standby, the COD is updated with the correct information. If you use the Global Array Manager (GAM) or Web Array Manager (WAM) software to change that drive’s state, the invalid COD is overwritten automatically (see “Global Array Manager (GAM)” and “Web Array Manager (WAM)” on page 52 for brief descriptions of GAM and WAM). If the drive with the invalid COD is inserted into the same slot as one that died and was removed, it starts a rebuild. If it is placed in any other nonconfigured slot, it comes up dead and must be marked as standby.

Mylex Online RAID Expansion (MORE)

The Mylex Online RAID Expansion operation adds capacity to an existing LUN while the controller is online with the host. For example, a storage system using a LUN with five disk drives can add another disk drive to create a LUN with six drives. The MORE operation can be performed on all RAID levels except JBOD. The procedures mentioned in this section are performed using the GAM or WAM GUI. The GAM or WAM online help and other GAM or WAM documentation describe the processes fully. "Global Array Manager (GAM)" on page 52 in Chapter 4 is an introduction to GAM and WAM.

The controller continues to service host I/O requests during the LUN expansion process, which includes restriping data from the old (smaller) LUN to the new (expanded) LUN.

The MORE operation provides two options for configuring the added capacity: Add Logical Drive (Add System Drive) option and Expand Logical Drive (Expand or Enlarge System Drive) option. These options and the requirements for disks to be added are explained in separate subsections:

- "Disk Drive Requirements and Restrictions for MORE Options" on page 46
- "Caveats for MORE Options" on page 47
- "Add Logical Drive Option" on page 48
- "Expand Logical Drive Option" on page 48

Disk Drive Requirements and Restrictions for MORE Options

For either MORE option, the disk drives must meet these criteria:

- One to six disk drives can be added to a LUN at one time. The maximum number of physical disk drives in the new storage system drives must not exceed eight.
- The minimum number of disk drives in the source LUN is two; the maximum is seven.
- The disk drive(s) being added must be in standby mode.
- The disk drive(s) being added must not already be part of a LUN.
- The capacity of each of the added disk drives must be greater than or equal to that of the smallest disk drive in the LUN.
- The system (logical) drive to be expanded must be online.

Caveats for MORE Options

For either MORE option (Add Logical Drive or Expand Logical Drive), the following caveats apply:

- Add capacity input parameters and execution parameters are stored in the controller's NVRAM. If a power failure occurs during the process, the controller firmware automatically restarts the process.
- In the event of a disk drive failure, the process continues to completion in critical mode. After migration is completed in critical mode, if automatic rebuild mode is enabled and if a spare disk drive is available, a rebuild starts immediately after migration is complete, restoring the critical system drive to an online state.

However, if the disk drive fails during the initialization of the newly created system drive (Add System Drive option), the migration operation stops immediately and is considered complete. At this time the rebuild operation begins, if applicable.

The migration process is aborted if a disk drive fails when capacity is added to a RAID 0 set or when a two disk drives fail.

- If a MORE process must be terminated (for example, two disk drives fail and cannot be recovered), the controller automatically aborts the migration process.
- Initialize, Rebuild, Consistency Check, and MORE are mutually exclusive operations. If one of the first three is already running, MORE fails to start.
- No configuration update commands issued from the host are allowed during the MORE process.
- Data is not lost if the RAID controller becomes disabled (for example, it is powered off by mistake or a power supply fails). The process resumes after power-on. Full recovery occurs after a power failure and the operation continues with no data loss.
- Write-back cache is disabled during the MORE process but resumes at the end of the operation. (That is, write cache is enabled during the MORE process, but is disabled at the end of the operation.)
- The following parameters do not change as a result of a MORE operation:
 - existing system drive write policy
 - LUN affinity/LUN mapping
 - cache size
 - all other controller parameters

- Removing and replacing a controller during migration is fatal. A hardware error that renders the controller unusable or the NVRAM unreadable is also fatal during migration.
- An add capacity operation does not start if the target LUN is in critical mode.

Add Logical Drive Option

The Add Logical Drive (Add System Drive) option adds a new system drive with the increased capacity, rather than by increasing the size of the existing system drive(s). The system drive's data is striped across a larger number of physical disk drives than before the expansion operation.

The Add Logical Drive option is rejected if the maximum of eight system drives is specified. When the initialization is complete, you must create and format the new system drive partition in accordance with Windows NT requirements.

Using the Add Logical Drive option, a LUN consisting of more than one system drive processes a MORE request on all system drives with one migration request, processing the system drives one at a time until all have been processed. Allowing one migration service to process all system drives on the same set creates one new extra capacity region when the migration is finished. The new system drive is added at the end of the LUN (selected drive pack).

Expand Logical Drive Option

The Expand Logical Drive (Expand or Enlarge System Drive) option appends the added capacity to the system drive specified in the MORE operation request. This option requires that only one system drive be defined using the set of physical disk drives; if more than one system drive exists on the same set of disk drives, the MORE operation is rejected.

Creating and formatting the added capacity requires backing up all data on the system drive to be enlarged (if appropriate), deleting the current partition, and creating the new, larger system drive.

After the MORE operation, you must delete, re-create, and reformat the system drive partition according to the operating system in use. The data can be saved back onto the system drive following the operating system format process.

Using the Expand Logical Drive option, the system drive specified is enlarged and its data is striped across a larger number of physical disk drives. For this option, the system drive specified must be the only system drive on the pack of disk drives being migrated.

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 49
- “Write Cache Disabled (Write-Through or Conservative Cache Mode)” on page 50

Write caching is set independently for each system drive in GAM or WAM.

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 if

- Only one controller is present and this controller fails
- 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 GAM or WAM for a dual-active controller configuration, and a controller failure has occurred.
- A MORE operation is initiated. During a MORE operation, the storage system automatically enters conservative cache mode. After MORE completes, the system drives are restored to their original settings. "Mylex Online RAID Expansion (MORE)" on page 46 has information about MORE.
- 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. 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. 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.

Using the RAID Controller

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

- “Software Tools for the Controller” on page 51
- “RAID Levels” on page 53
- “CAP Strategy for Selecting a RAID Level” on page 55
- “Loop Configuration” on page 59
- “System Drives” on page 64
- “Drive State Reporting” on page 65
- “Automatic Rebuild” on page 66

Software Tools for the Controller

Server-client software allows you to manage the RAID controllers. The software tools are as follows:

- Windows NT and Windows 2000 Global Array Manager (GAM) server driver and a Java-based Web Array Manager (WAM) client
- IRIX Global Array Manager (GAM) server driver and a Java-based Web Array Manager (WAM) client
- Linux Global Array Manager (GAM) server driver and a Java-based Web Array Manager (WAM) client

Because the client is Java-based, it can run on a variety of platforms.

Global Array Manager (GAM)

The Global Array Manager Client software (GAM Client) configures, manages, and monitors RAID arrays connected to the controllers, even across remote servers. For more information on GAM components, requirements, and operation, refer to the documentation included on the Mylex RAID Management Software CD-ROM.

The Global Array Manager “wizard” for setting up and configuring new logical drives and disk arrays is called RAID Assist. In its simplest form, RAID Assist provides an Auto Configuration option that immediately configures all available drives into an optimal RAID 5 configuration. RAID Assist’s Assisted Configuration sets up a new array according to predefined parameters, and asks questions to gather the information necessary to build the array. If the configuration needs to go beyond these options, the Manual Configuration option allows additional control over logical drive setup parameters.

For detailed information on using GAM for configuration, see the Mylex RAID Management Software CD-ROM or the *Global Array Manager Client Software Installation and User Guide*.

Web Array Manager (WAM)

The Web Array Manager Client software (WAM Client) is used to:

- Monitor, manage, maintain, and configure SGI RAID Controllers and the physical and logical drives that are connected to these controllers, even across remote servers
- Monitor and manage SCSI Host Bus Adapters and the physical devices that are connected to these adapters, even across remote servers

The Web Array Manager has three components:

- Global Array Manager Server component
- Web Array Manager Client component
- Java Runtime Environment

Configuration functions are easily performed using WAM’s Express (automatic) and Custom (manual) configuration tools. Just select a few preferences, and the Express configuration tool automatically does the rest. Use the Custom configuration tool for more control over drive group setup or individual configuration parameters.

The Global Array Manager (GAM) Server component collects and disseminates information on disk array or HBA subsystem status and resource utilization. The WAM Client component organizes this information through an intuitive graphical display. Errors and events are recorded in a log, and if a problem is serious enough to warrant immediate attention, operators can be alerted via popup windows, pagers, email, or an application.

You can use the WAM Client to manage or perform maintenance on individual disk arrays and drives (with the appropriate authentication). This includes forcing physical devices offline (also known as “killing” a drive), rebuilding drives, creating hot spares, and initiating consistency checks on arrays that are configured for redundancy.

For detailed information on using WAM client for configuration, see the *Mylex Web Array Manager™ Client Software Installation Guide and User Manual* or the *Web Array Manager Client Software Installation and User Guide*.

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 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	8	No
1	Disk drives are paired and mirrored. All data is duplicated 100% on an equivalent disk drive.	2	2	Yes

Table 4-1 (continued) Supported RAID Levels

RAID Level	Description	Minimum Drives	Maximum Drives	Fault-Tolerant
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	8	Yes
5	Data and parity information is striped across all physical disk drives. Parity protection is used for data redundancy.	3	8	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	8	Yes
JBOD (7)	“Just a bunch of disks.” 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
30 and 50	Allow storage capacity to be increased by a factor of 4 from RAID 3 and RAID 5 configurations by combining four RAID 3 or RAID 5 drive groups into a superdrive group across which data is striped. Each of the four drive groups must contain the same number of physical devices, use the same stripe block size, and be the same RAID level. The appended 0 denotes that the RAID level is using striping (RAID 0) across drive groups.	3	8	Yes

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.

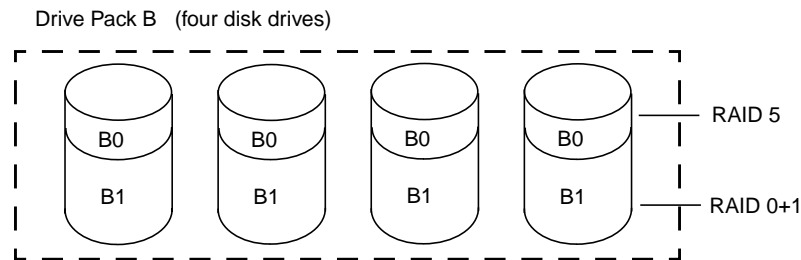


Figure 4-1 RAID Levels Within a Drive Pack (LUN): Example

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 SGI TP 9100 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 56
- “Configuring for Maximum Availability” on page 57
- “Configuring for Maximum Performance” on page 59

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 disk drives of size 2 GB each.

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 operational 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 taking place, 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 128 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 uninterruptible 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 (continued) RAID Levels and Availability

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

Loop Configuration

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

- The largest RAID 3 or RAID 5 group that can be created is 7+1 drives.
- For a tower, the maximum configuration is 12 drives total, those in the system itself; no expansion to another enclosure or tower is possible. A maximum of 12 drives can be attached to a drive loop.
- For a RAID enclosure and two expansion enclosures, the maximum configuration is 32 drives. Although the enclosures can hold 36 drives, the controller sees a maximum of 32. A maximum of 32 drives can be attached to a drive loop.

- A maximum of 8 system drives can be created (see “System Drives” on page 64 for information on these).

The disk drive modules are dual-ported. A RAID controller sees 12 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.

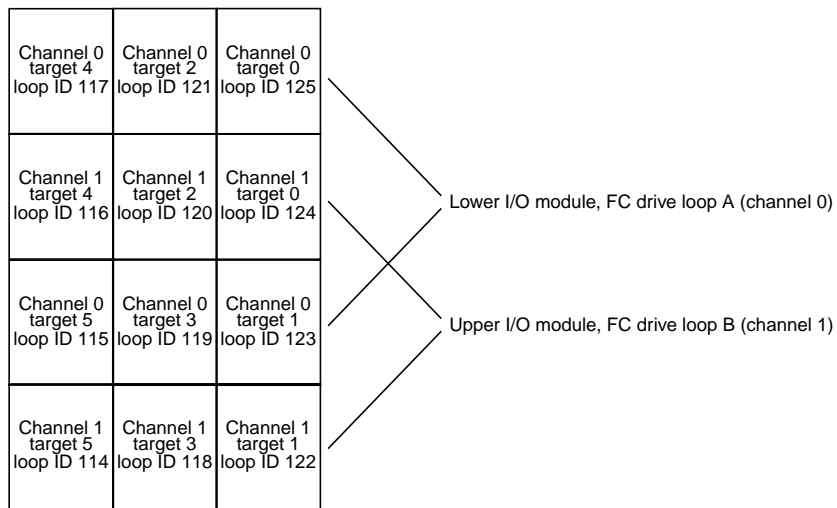


Figure 4-2 I/O Modules, Channels, and Loops: Tower

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

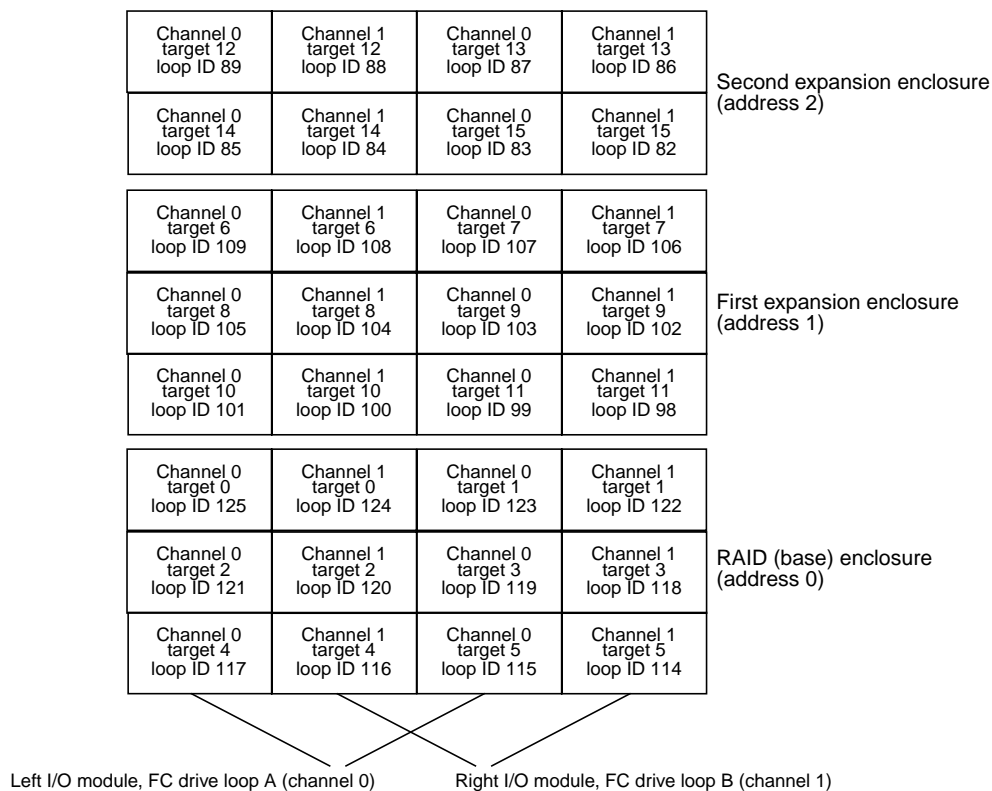


Figure 4-3 I/O Modules, Channels, and Loops: Rackmount Enclosure (Front View)

However, you can use GAM or WAM 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.

Figure 4-4 diagrams an example relationship of a single RAID controller, the I/O modules, and loops. In this configuration, all drives are on both channels for availability. In a configuration with two loops, if one loop fails, access is through the alternative path, which might affect performance.

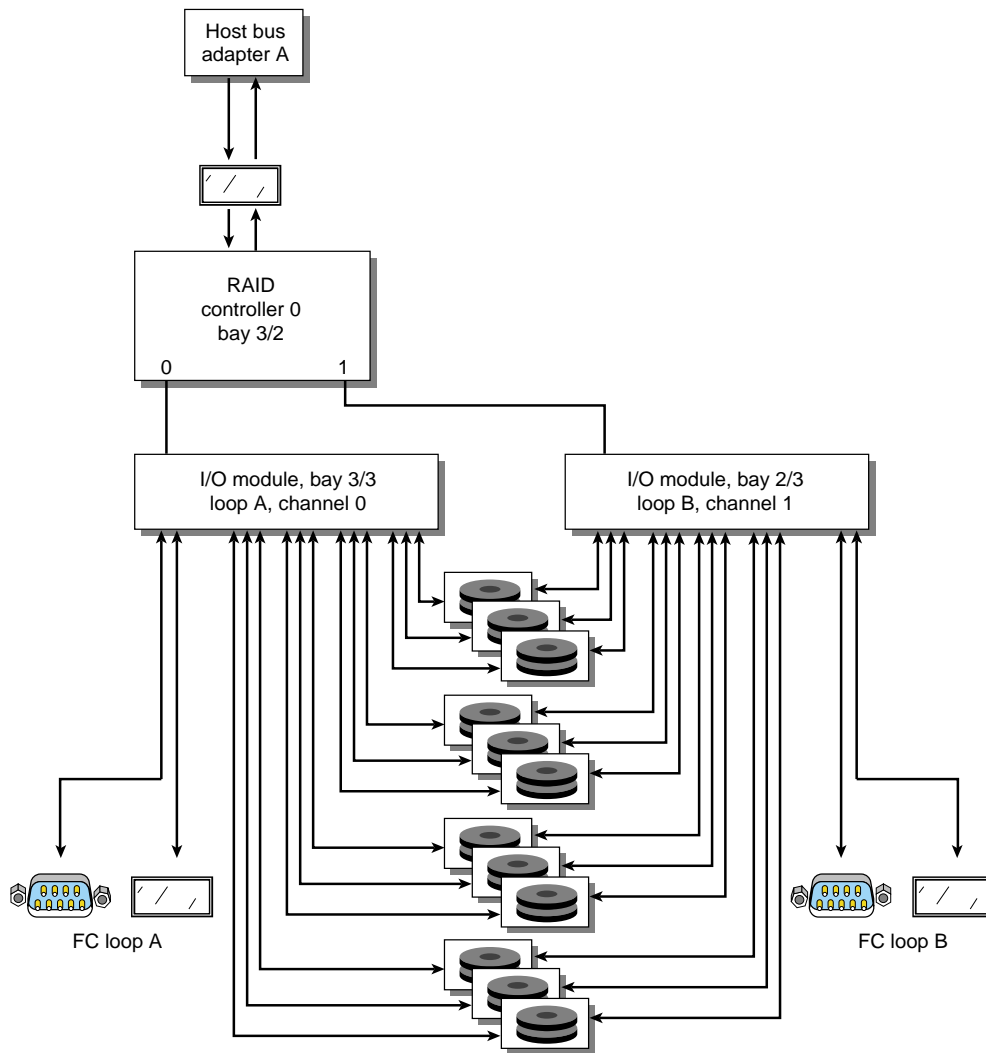


Figure 4-4 Single-Controller/Point-to-Point/Single-Host Example

Figure 4-5 shows two host bus adapters (HBA), dual-active RAID controllers, and dual I/O modules and loops. In this configuration, all drives are on both channels of each RAID controller for availability. In such a configuration, if one controller or loop fails, failover access is through the alternate HBA, controller, or loop.

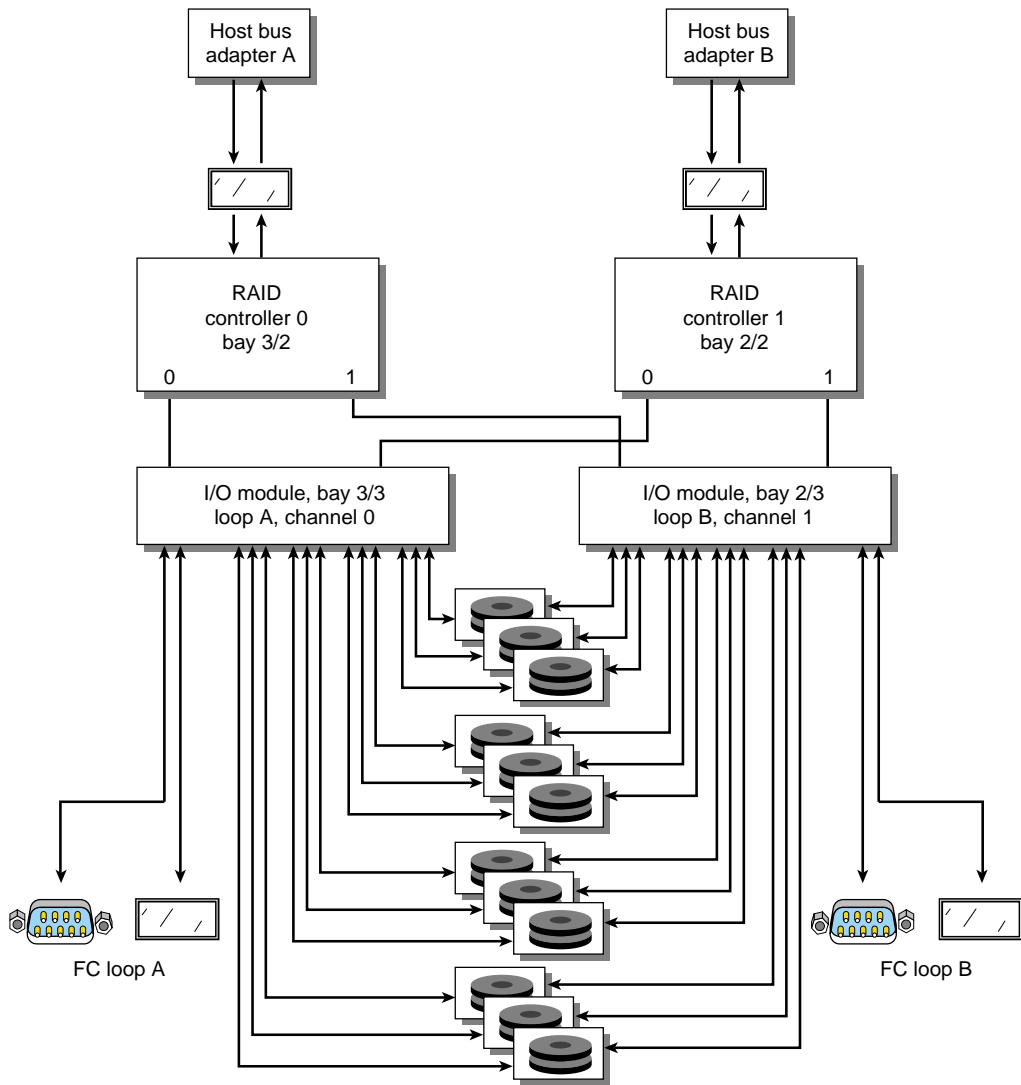


Figure 4-5 Dual-Active Controller/Point-to-Point/Dual-Host Example

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 64
- “System Drive Affinity and Programmable LUN Mapping” on page 64

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 eight 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 49 in Chapter 3 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 64.
- 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:

- 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: When setting up for multipath failover, use the GAM or WAM software to set the Topology to MULTI-PORT.

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 controller provides information to the configuration tool or utility to help it maintain operational information for each physical disk drive. The controller keeps a record of each physical disk drive’s operational state and information on available target ID addresses. The controller first determines the presence or absence of the physical disk drive or available target ID location and then determines the operational state of the disk drives present.

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	The disk drive is powered on, has been defined as a member of a drive pack, and is operating properly.
Standby	The disk drive is powered on, is able to operate properly, has not been defined as part of any drive pack, and has been defined as a standby or hot spare.
Dead	The disk drive is not present, is present but not powered on, is a newly inserted replacement disk drive, or fails to operate properly and is labeled Dead by the controller (whether or not it has been defined as a member of a drive pack).
Rebuild	The disk drive is in the process of being rebuilt, that is, data is being copied from the mirrored disk drive to the replacement disk drive in a RAID 1 or RAID 0+1 (RAID 6) array, or data is being regenerated via the exclusive OR (XOR) redundancy algorithm and written to the replacement disk drive in a RAID 3 or 5 array.

Table 4-6 (continued) Physical Disk Drive States

State	Description
Unconfigured	This location is unconfigured.
Environmental	An environmental device is present at this address.

For more information, see the GAM or WAM documentation and online help.

Automatic Rebuild

The 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 both of the following conditions are true:

- A standby or hot spare disk drive of identical or larger size is attached to the same controller.
- 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.

Note: The controller always attempts first to locate a replacement (standby or hot spare) disk drive that is exactly the same size as the failed disk drive. If none is found, the controller attempts to locate a replacement disk drive that is at least the same size as the failed disk drive.

During the automatic rebuild process, storage system activity continues as normal; however, storage system performance may degrade slightly.

Note: The priority of rebuild activity can be adjusted using the GAM or WAM controller parameters to adjust the rebuild rate.

To use the automatic rebuild feature, you must maintain a hot spare or standby disk drive in the storage system. The number of hot spare disk drives in a storage system is limited only by the maximum number of disk drives available on each disk channel.

You can create a hot spare disk drive in one of two ways:

- A disk drive can be labeled as a spare or standby using the GAM or WAM disk information dialog.
- The configuration is created or changed using the GAM or WAM RAID Assist dialog.

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.

Troubleshooting

The SGI TP 9100 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 GAM or WAM management utilities.

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 Rack Door" on page 28 in Chapter 1.

This chapter consists of the following sections:

- "Solving Initial Startup Problems" on page 70
- "Using Storage System LEDs for Troubleshooting" on page 71
- "Using the Alarm for Troubleshooting" on page 80
- "Storage System Temperature Conditions, Problems, and Solutions" on page 81
- "Using Test Mode" on page 83

Note: If a fault is indicated on the ESI/ops panel, refer first to Table 5-1 in "Using Storage System LEDs for Troubleshooting" on page 71 and then to the section referenced within that table.

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 RAID module fault LED is also on, there is a RAID controller fault; check that the RAID module is correctly seated.

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

- A fan is slowing down; see “ESI/Ops Panel Power Supply/Cooling Module LED: Faults and Remedies” on page 75 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 “ESI/Ops Panel Power Supply/Cooling Module LED: Faults and Remedies” on page 75.
- There is a storage system fault; see “ESI/Ops Panel System LED: Faults and Remedies” on page 75.

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

- Make sure 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.
- Make sure the FC-AL interface cables from the RAID module to the fibre channel board in the host computer are fitted correctly.
- Check the selector switches on the I/O modules in the storage system:
 - On a tower or a RAID enclosure, both I/O modules should be set to 0.
 - On the first expansion enclosure attached to a RAID system, both I/O modules should be set to 1.
 - On the second expansion enclosure attached to a RAID system, both I/O modules should be set to 2.
- See if the LEDs on all installed drive carrier modules are green. Note that the drive LEDs flash during drive spinup.
- Check that all disk drive modules are correctly installed.

- Check that the RAID controller has completed its startup and the green controller ready LED is illuminated; see “RAID Controller LEDs and Problem Solutions” on page 78.

If an amber disk drive module LED drive fault is on, there is a drive fault: see Table 5-2 on page 80.

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” on page 72
- “Power Supply/Cooling Module LED and Problem Solution” on page 76
- “FC LRC I/O Module LEDs and Problem Solutions” on page 77
- “RAID Controller LEDs and Problem Solutions” on page 78
- “Disk Drive LEDs and Problem Solutions” on page 79

ESI/Ops Panel LEDs

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




Figure 5-1 ESI/Ops Panel Indicators and Switches: Rackmount Enclosure

Figure 5-2 shows details of the ESI/ops panel in the tower.





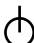


Figure 5-2 ESI/Ops Panel Indicators and Switches: Tower

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

Table 5-1 ESI/Ops Panel LEDs and Troubleshooting

LED	Off	Green Steady	Green or Amber Flashing	Amber Steady	Remedy
 System status	Power is off, or standby mode is selected.	All ESI functions OK.	System configuration error.	ESI / PC function failure.	See "ESI/Ops Panel System LED: Faults and Remedies" on page 75.
 Drive status	Power is off, or standby mode is selected.	All drives OK.	Drive fault (alternating green and amber). Note: ESI polling causes flash at regular interval.	N/A	Replace the drive as explained in Chapter 6, "Installing and Replacing Disk Drive Modules."
 Power/ Cooling	Power is off, or standby mode is selected.	All power and fans OK.	Not used.	Power supply/cooling module or fan fault.	See "ESI/Ops Panel Power Supply/Cooling Module LED: Faults and Remedies" on page 75. Contact service provider for replacement.
 Standby	No AC power is supplied.	Power is good, or standby status is selected.	Not used.	Storage system is in standby mode due to ESI thermal shutdown.	See "Thermal Alarm" on page 82, "Thermal Shutdown" on page 82.

Note the following:

- If all LEDs on the ESI/ops panel flash simultaneously, see "Using Test Mode" on page 83.
- If test mode has been enabled (see "Using Test Mode" on page 83), the amber and green drive bay LEDs flash for any non-muted fault condition.
- The LED for the uninterruptible power supply (UPS) is covered with a sticker. This LED is lit whether or not a UPS is present. Use of a UPS is not supported in this release. Leave the sticker in place.

ESI/Ops Panel System LED: Faults and Remedies

If the system status LED on the ESI/ops panel is amber and the alarm is sounding, the ESI processor has detected an internal fault, such as failure of an internal communications path. Follow these steps:

1. Check if the LEDs on the power supply/cooling modules are amber. If a PSU error is present, there might be a communications problem with that power supply/cooling module. Contact your service representative for a replacement.
2. Check if the disk drive module LEDs are amber. If none are, there might be either an ESI processor problem or a chassis midplane problem. Contact your service provider.

See also “Thermal Shutdown” on page 82.

ESI/Ops Panel Power Supply/Cooling Module LED: Faults and Remedies

If the alarm is sounding, the power supply/cooling LED on the ESI/ops panel is amber, and the LED on one or both power supply/cooling modules is amber, the cause is one of the following:

- Power fault
- Fan failure
- Thermal condition that is causing the power supply/cooling module to overheat

To troubleshoot the problem, follow these steps:

1. Check the power connection to each power supply/cooling module.
2. If possible, reduce the ambient temperature. If the problem persists, contact your service provider.

If one of the power supply/cooling modules is missing or defective, all of the following occur:

- The alarm sounds
- The system LED on the ESI/ops panel flashes
- The power supply/cooling module LED remains illuminated

If these warnings occur, check the LEDs on the power supply/cooling modules. If one or both are amber, contact your service representative.

Power Supply/Cooling Module LED and Problem Solution

Figure 5-3 shows the location of the LED on the power supply/cooling module.




Figure 5-3 Power Supply/Cooling Module LED

If this LED is amber, or if the power/cooling LED on the ESI/ops panel is amber and the alarm is sounding, contact your service provider.

FC LRC I/O Module LEDs and Problem Solutions

Two LEDs between the DB-9 connectors on the I/O panel glow steady green to indicate a good FC-AL signal on the cables attached to the adjacent connectors. Figure 5-4 shows the LEDs on the I/O module panel in a rackmount enclosure.




Figure 5-4 FC LRC I/O Module LEDs: Rackmount Enclosure

Figure 5-5 shows the LEDs on the I/O module panel in a tower.




Figure 5-5 FC LRC I/O Module LEDs: Tower

If the LED for a connector is amber, check the connection. If the problem persists, contact your service provider for a replacement.

RAID Controller LEDs and Problem Solutions

Figure 5-6 shows the LEDs on the RAID controller panel in a rackmount enclosure; Figure 5-7 shows them in a tower.




Figure 5-6 RAID Controller Panel LEDs: Rackmount Enclosure




Figure 5-7 RAID Controller Panel LEDs: Tower

The RAID controller's activity LED flashes every time the Global Array Management (GAM) server issues a poll. The default interval is 7 seconds; to reset the interval, see the documentation for the GAM software.

If the fault/controller not ready LED is on, check that the module is properly seated. If you cannot find the fault, contact your service provider.

Disk Drive LEDs and Problem Solutions

Each disk drive module has two LEDs, an upper (green) and a lower (amber), as shown in Figure 5-8.




Figure 5-8 Disk Drive Module LEDs

The LEDs on the disk drive modules provide the information summarized in Table 5-2.

Table 5-2 Disk Drive LED Function

Green LED	Amber LED	State	Remedy
Off	Off	No disk drive fitted; that is, 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
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
- Drive fault
- Fan slows down
- Voltage out of range
- Overtemperature
- Thermal overrun
- Storage system fault

You can mute the audible alarm by pressing the alarm mute pushbutton for about a second, until you hear a double beep. The mute pushbutton is beneath the indicators on the ESI/ops panel (see Figure 5-1 and Figure 5-2).

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.

Storage System Temperature Conditions, Problems, and Solutions

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

- “Thermal Control” on page 81
- “Thermal Alarm” on page 82
- “Thermal Shutdown” on page 82

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 thermal alarm is a combination of all these:

- The alarm sounds.
- The power supply/cooling LED on the ESI/ops panel is amber.
- The LED on one or both power supply/cooling modules is amber.
- The air temperature of the air exiting the PSU is above 55 °C (131 °F).

This alarm sounds when the internal temperature (measured in the airflow through the storage system) exceeds a threshold. Follow these steps:

1. Check that local ambient environment temperature is below the upper temperature specification of 40 °C (104 °F).
2. Check that there is clear, uninterrupted airflow at the front and rear of the storage system.
3. Check for restrictions due to dust buildup; clean as appropriate.
4. Check for excessive recirculation of heated air from the rear to the front.

Thermal Shutdown

The storage system automatically shuts down to protect itself and the disk drives from damage. Thermal shutdown occurs:

- After the thermal alarm condition is present; shutdown takes place at a higher threshold than the thermal alarm
- When all fans fail

- When only one fan is operating and the storage system internal temperature reaches 40° C (104° F)

The storage system shuts down ten seconds after any of these conditions is present.

When a thermal shutdown is imminent, all amber LEDs on the ESI/ops panel and on all disk drives flash and the alarm sounds continuously and cannot be muted.

When thermal shutdown conditions are present, follow these steps (if possible):

1. Check for airflow restrictions.
2. Check for power supply/cooling module faults as detailed in “ESI/Ops Panel Power Supply/Cooling Module LED: Faults and Remedies” on page 75.
3. Check for excessive ambient temperatures (over 40 °C (104 °F))

If the overheating problem is not solved, thermal shutdown occurs. The standby LED on the ESI/ops panel and the power supply status LED are amber and the storage system powers itself off.

Follow these steps:

1. Correct the source of the overheating.
2. Allow the storage system to cool down.
3. Unplug the power cord from the storage system and leave it unplugged for at least 30 seconds to reset the shutdown condition.
4. Reconnect the power cord and restart the storage system following the normal procedure (see “Powering On the Rack” on page 37 and “Powering On the SGI TP 9100 Tower” on page 33 in Chapter 2).
5. Check for cooling faults that persist, particularly fan failure. If a fan has failed, contact your service provider for a replacement.

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

Installing and Replacing Disk Drive Modules

This chapter explains how to replace a faulty disk drive module or install a new one, in the following sections:

- “Adding a Disk Drive Module” on page 86
- “Replacing a Disk Drive Module” on page 96

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.

Adding a Disk Drive Module

Note the following:

- All disk drive bays must be filled with either a disk drive module or a dummy drive; no bay should be left completely empty.
- The drives in bays 1/3 and 4/3 are required for enclosure management; these bays must always be occupied.
- Install drives in the order shown in Figure 6-1. Generally, after the two required bays are occupied, drive bays are filled starting at the bottom of the storage system.

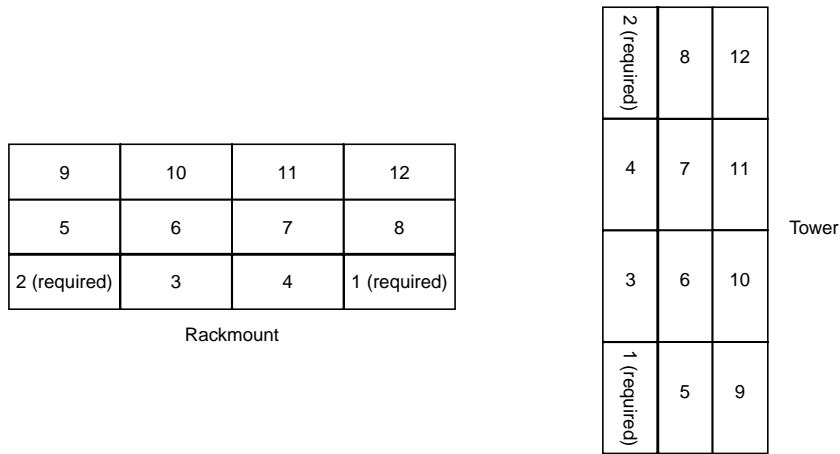


Figure 6-1 Drive Installation Order

Follow these steps to add a new disk drive module to the storage system:

1. Make sure that enough disk drives and dummy drives have been ordered to occupy all bays.
2. For each new drive, carefully open the bag containing the drive.



Warning: The disk drive handle might have come 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 disk drive modules on an antistatic work surface, and make sure that their antitamper locks are disengaged (unlocked). A disk drive module cannot be installed if its antitamper lock is activated outside the enclosure.

Drives are shipped with the 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 clockwise until the indicator visible in the center aperture of the handle shows black. See Figure 6-2.




Figure 6-2 Unlocking the Disk Drive Module (Disengaging the Antitamper Lock)

4. Release the disk drive module handle by pressing on the small indentation on the left of the module front; see Figure 6-3.




Figure 6-3 Releasing the Module Handle

5. Insert the disk drive module into the chassis. Figure 6-4 shows the procedure in a rackmount enclosure; Figure 6-5 shows it in a tower.




Figure 6-4 Inserting the Disk Drive Module: Rackmount Enclosure

Note: The drive modules are not keyed. Make sure you are inserting the drive module in the correct orientation.




Figure 6-5 Inserting the Disk Drive Module: Tower

Note: The drive modules are not keyed. Make sure you are inserting the drive module in the correct orientation.

6. Push the disk drive module gently into the chassis until it is stopped by the camming lever on the right of the module. Figure 6-6 shows the procedure in a rackmount enclosure; Figure 6-7 shows it in a tower.




Figure 6-6 Drive Seated: Rackmount Enclosure




Figure 6-7 Drive Seated: Tower

7. Press on the drive handle to seat the drive. The camming lever on the right of the module engages with a slot in the chassis. Figure 6-8 shows this operation in a rackmount enclosure; Figure 6-9 shows it in a tower.




Figure 6-8 Engaging the Camming Lever: Rackmount Enclosure




Figure 6-9 Engaging the Camming Lever: Tower

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.

8. Repeat steps 2 through 7 for all drive modules to be installed.

9. When you have finished installing all drives, activate the drive antitamper lock on each one:
 - Carefully insert the lock key provided into the cutout in the handle trim. See Figure 6-10.




Figure 6-10 Locking the Disk Drive Module (Setting the Antitamper Lock)

- Rotate the key 90 degrees counterclockwise until the indicator in the handle shows red.
 - Remove the key and store it.
10. Fit all empty drive bays with dummy drives. The drive handle and camming mechanism operate the same way as for disk drives.

Replacing a Disk Drive Module

This section explains how to replace a defective disk drive module, in the following sections

- “LUN Integrity and Disk Drive Module Failure” on page 96
- “Replacing the Disk Drive Module” on page 97

LUN Integrity and Disk Drive Module Failure

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

Note: Before replacing a disk module, use the software included with the storage system to check disk module 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 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 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. Make sure users are off the affected systems; back up data if necessary.

Note: Replace disk drive modules one at a time,

3. Using the key, turn the antitamper lock to the off position:
 - Carefully insert the lock key provided into the cutout in the handle trim, as shown in Figure 6-2.
 - Rotate the key 90 degrees clockwise until the indicator in the handle shows black.
 - Remove the key and store it.
4. Make sure the faulty drive has spun down.

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

5. Release the handle by pressing on the small indentation area on the left of the module front.
6. Gently withdraw the module approximately 25 mm (one inch) and wait 30 seconds.
7. Withdraw the module from the drive bay. Replace the drive immediately; follow instructions in "Adding a Disk Drive Module" on page 86.
8. 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 99
- “Environmental Requirements” on page 101
- “Power Requirements” on page 102
- “RAID Controller Module Specifications” on page 103
- “SGI Cables for the SGI TP 9100 Storage System” on page 104

Storage System Physical Specifications

Table A-1 gives the physical specifications for the rackmount enclosures and the rack.

Table A-1 Physical Specifications, Rackmount Enclosures and Rack

Dimension	RAID (Base) Enclosure	Expansion Enclosure	Rack
Height	17.7 cm (6.95 in.)	17.7 cm (6.95 in.)	Operating: 180.0 cm (5.94 ft) Shipping: 210.0 cm (6.93 ft)
Width	44.6 cm (17.5 in.)	44.6 cm (1.75 in.)	Operating: 60.0 cm (1.98 ft) Shipping: 120.0 cm (3.96 ft)
Depth	58.6 cm (23.0 in.)	58.6 cm (23.0 in.)	Operating: 80.0 cm (2.64 ft) Shipping: 120.0 cm (3.96 ft)
Weight, approximate	Maximum: about 44 kg (97 lb); Empty: 13 kg (28.6 lb)	Maximum: about 43 kg (94.3 lb)	Operating, maximum: 575 kg (1265 lb) Shipping, maximum: 650 kg (1430 lb) Empty (with PDUs): 200 kg (440 lb)

Exact weight of the units is affected by the type and number of disk drive modules installed. Exact weight of the rack is affected by the enclosures installed.

The height of 1U is 1.75 inches (4.445 cm).

Table A-2 gives the physical specifications for the tower.

Table A-2 Physical Specifications, Tower

Dimension	Tower
Height	57 cm (22.4 in.), including wheels
Width	30.2 cm (11.8 in.) at base 183 cm (71.9 in.) at top
Depth	59.2 cm (23.2 in.)
Weight, maximum configuration	About 53.5 kg (about 117.7 lb)

Exact weight of the tower is affected by the type and number of disk drive modules installed.

Table A-3 gives weights of various component modules.

Table A-3 Module Weights

Module	Weight
Power supply/cooling module	4 kg (8.8 lb)
ESI/Ops panel module	Under 1 kg (under 2.2 lb)
RAID controller module	About 0.8 kg (about 1.76 lb)

Environmental Requirements

Table A-4 gives temperature and humidity requirements for both storage system configurations.

Table A-4 Ambient Temperature and Humidity Requirements

Factor	Temperature	Relative Humidity	Maximum Wet Bulb
Operating temperature	10 °C to 40 °C (50 °F to 104 °F)	20% to 80% noncondensing	23 °C (73.4 °F)
Nonoperating temperature	0 °C to 50 °C (32 °F to 122 °F)	8% to 80% noncondensing	27 °C (80.6 °F)
Shipping	-40 °C to +60 °C (-40 °F to 140 °F)	5% to 100% nonprecipitating	29 °C (84.2 °F)

Table A-5 gives other environmental requirements for the storage system.

Table A-5 Environmental Requirements

Factor	Requirement
Altitude	0 to 2133 m (6996.2 feet or 1.3 miles)
Operational shock	Vertical axis 5 g peak 1/2 sine, 10 ms
Nonoperational shock	20 g 20 ms square wave
Operational vibration	Random vibration power spectrum available on request
Acoustics	Freestanding enclosure declared 'A' weighted sound power level equal to or less than 68 decibels

Power Requirements

Table A-6 gives minimum storage system power requirements.

Table A-6 Minimum Power Requirements

Factor	Requirement
Voltage	Tower: 100 to 120 or 200 to 240 VAC Rack: 200 to 240 VAC
Frequency	47 to 63 Hz
Maximum power consumption	800 VA
Typical power consumption	400 VA or less
Inrush current (25 °C (77 °F) cold start 1 PSU)	40/80 A @ 110/230 VAC

Table A-7 gives 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
Overvoltage category	II
Maximum load per PDU	24 A
Maximum load per bank of outlet sockets on each circuit breaker	8 A
Plug	NEMA L6-30

RAID Controller Module Specifications

Table A-8 gives specifications for the RAID controller module.

Table A-8 RAID Module Specifications

Factor	Requirement
Controller type	Mylex Corp. DACFFX
CPU	Two Intel 960RD RISC 32-bit microprocessors
Host interface connectors	One GBIC port (accommodates any standard GBIC module), up to 100 MB/sec
External FC-AL signal cables	SGI dual-port HBAs: 25 m (82 ft) Storage area network (SAN) and SGI single-port HBAs: maximum 12 m (39.36 ft) copper or 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
Mode of operation	Simplex: single controller
Current limit	1.5 A
RAID levels	0, 1, 3, 5, and 0+1 (RAID level 6) 10, 30, and 50 JBOD (RAID level 7)
Configuration and management	Via FC-AL using Global Array Manager (GAM) or Web Array Manager (WAM) utility
Indicators	LEDs for Controller Ready, Activity, Failed, and Cache
Memory	Default 128 MB; approved memory options available
Cache	Selectable write-through or write-back Read always enabled
Battery	NiCd cache battery protects 128 MB data for up to 72 hours

Disk Drive Module Specifications

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

Table A-9 Drive Carrier Module Specifications (1.6-Inch 18 GB Drive)

Factor	Requirement
Weight	1.0 kg (2.2 lb) or 1.4 kg (3 lb)
Operating temperature	10 °C to 40 °C (50 °F to 104 °F) when installed
Power dissipation	24 watts maximum

SGI Cables for the SGI TP 9100 Storage System

Table A-10 lists SGI cable options that can be connected to the SGI TP 9100 product.

Table A-10 SGI Fibre Channel Fabric Cabling Options for SGI TP 9100 Storage System

Cable	Length	Marketing Code	Part Number
1 meter copper HSSDC to HSSDC cable	1 m (39.3 in.)	X-FS-COP-1M	N/A
3 meter copper HSSDC to HSSDC cable	3 m (9.8 ft)	X-FS-COP-3M	N/A
12 meter copper HSSDC to HSSDC cable	12 m (39.4 ft)	X-FS-COP-12M	N/A
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 SGI fibre channel switches.

Regulatory Information

The SGI Total Performance (TP) 9100 rack conforms to Class B specifications. The SGI TP 9100 rackmountable enclosures and tower storage systems conform to Class B specifications. The information is given separately, in these sections:

- “Class A Warnings and Specifications for the SGI TP 9100 Rack” on page 105
- “Class B Warnings and Specifications: SGI TP 9100 Tower and Enclosures” on page 107

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

Class A Warnings and Specifications for the SGI TP 9100 Rack

This section contains Class A warnings and specifications. It applies only to the SGI Total Performance 9100 rack. It does not apply to the SGI TP 9100 tower or to enclosures used outside the rack.

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 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; Germany’s BZT Class A limits for 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。

Class B Warnings and Specifications: SGI TP 9100 Tower and Enclosures

This section contains Class B warnings and specifications. It applies only to the SGI TP 9100 tower or to enclosures used outside the rack. It does not apply to the SGI Total Performance 9100 rack.

The SGI TP 9100 rackmountable enclosures and tower storage system conform to several national and international specifications and European Directives listed in "Manufacturer's Declaration of Conformity" on page 108. The CE insignia displayed on each device is an indication of conformity to the European requirements.

Caution: The SGI TP 9100 storage system has several governmental and third-party approvals, licenses, permits. Do not modify this product in any way that is not expressly approved by SGI. If you do, you may lose these approvals and your governmental agency authority to operate this device.

Manufacturer's Declaration of Conformity

A "Manufacturer's Declaration of Conformity" is available on the World Wide Web. Look on your system (regulatory) label on the rear of your workstation to determine your CMN (model) number, which you need to identify your Declaration of Conformity, and to find the Series Number of your processor.

To locate the information on the World Wide Web, enter the following in your Web browser location window:

<http://www.sgi.com/compliance/index.html>

FCC Class B Conformance

This equipment complies with Part 15 of the FCC Rules. Operation is subject to the following two conditions: (1) This device may not cause harmful interference, and (2) this device must accept any interference received, including interference that may cause undesired operation.

This equipment has been tested and found to comply with the limits for a Class B digital device, pursuant to Part 15 of the FCC Rules. These limits are designed to provide reasonable protection against harmful interference in a residential installation. This equipment generates, uses and can radiate radio frequency energy and, if not installed and used in accordance with the instructions, may cause harmful interference to radio communications. However, there is no guarantee that interference will not occur in a particular installation. If this equipment does cause harmful interference to radio or television reception, which can be determined by turning the equipment off and on, the user is encouraged to try to correct the interference by one or more of the following measures:

- Reorient or relocate the receiving antenna.
- Increase the separation between the equipment and receiver.

- Connect the equipment into an outlet on a circuit different from that to which the receiver is connected.
- Consult the dealer or an experienced radio/TV technician for help.

Electromagnetic Emissions

This device complies with the Class B limits of Part 15 of the FCC Rules. Operation is subject to the following two conditions:

- This device may not cause harmful interference.
- This device must accept any interference received, including interference that may cause undesired operation.

Also, this device complies with Class B electromagnetic emissions limits of C.I.S.P.R. Publication 22, Limits and Methods of Measurement of Radio Interference Characteristics of Information Technology Equipment.

Shielded Cables

The SGI TP 9100 rackmountable enclosures and tower storage system are FCC-compliant under test conditions that include the use of shielded cables between the storage system and a host. Shielded cables reduce the possibility of interference with radio, television, and other devices. If you use any cables that are not from SGI, make sure they are shielded. Telephone cables do not need to be shielded.

Electrostatic Discharge

SGI designs and tests its products to be immune to the effects of electrostatic discharge (ESD). ESD is a source of electromagnetic interference and can cause problems ranging from data errors and lockups to permanent component damage.

While operating the SGI TP 9100 enclosure or tower storage system, it is important that you keep all the covers, doors, and dummy (filler) panels in place. The shielded cables that came with the product should be installed correctly.

Industry Canada Notice (Canada Only)

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

CE Notice

Marking by the "CE" symbol indicates compliance of the device to directives of the European Community. A "Declaration of Conformity" in accordance with the above standards has been made and is available from SGI upon request.



VCCI Notice (Japan Only)

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B情報技術装置です。この装置は、家庭環境で使用することを目的としていますが、この装置がラジオやテレビジョン受信機に近接して使用されると、受信障害を引き起こすことがあります。取扱説明書に従って正しい取扱いをして下さい。

Korean Regulatory Notice

이 기기는 가정용으로 전자파적합등록을 한 기기로서 주거지역에서는 물론 모든 지역에서 사용할 수 있습니다.

NOM 024 Information (Mexico Only)

La información siguiente se proporciona en el dispositivo o en dispositivos descritos en este documento, en cumplimiento con los requisitos de la Norma Oficial Mexicana (NOM 024):

Exportador:	Silicon Graphics, Inc.
Importador:	Silicon Graphics, Inc.
Embarcar a:	Av. Vasco de Quiroga No. 3000 Col. Santa Fe C.P. 01210 México, D.F. México
Tensión alimentación:	100-120/200-240 V~
Frecuencia:	50/60 Hz
Consumo de corriente:	6.5/3.5 A

EN 55022 Compliance (Czech Republic)

Pokud není na typovém štítku počítače uvedeno, že spadá do třídy A podle EN 55022, spadá automaticky do třídy B podle EN 55022. Pro zařízení zařazená do třídy A (ochranné pásmo 30m) podle EN 55022 platí následující. Dojde-li rušení telekomunikačních nebo jiných zařízení, je uživatel povinen provést taková opatření, aby rušení odstranil.

Polish Notice (Poland Only)

Współpracujące ze sobą urządzenia (komputer, monitor, drukarka) powinny być zasilane z tego samego źródła.

Instalacja elektryczna pomieszczenia powinna zawierać w przewodzie fazowym rezerwową ochronę przed zwarciami, w postaci bezpiecznika o wartości znamionowej nie większej niż 10A (amperów).

W celu całkowitego wyłączenia urządzenia z sieci zasilania, należy wyjąć wtyczkę kabla zasilającego z gniazdka, które powinno znajdować się w pobliżu urządzenia i być łatwo dostępne.

Pozostałe Instrukcje Bezpieczeństwa

- Nie należy używać wtyczek adapterowych lub usuwać kołka obwodu ochronnego z wtyczki. Jeżeli konieczne jest użycie przedłużacza to należy użyć przedłużacza 3-żyłowego z prawidłowo połączonym przewodem ochronnym.
- System komputerowy należy zabezpieczyć przed nagłymi, chwilowymi wzrostami lub spadkami napięcia, używając eliminatora przepięć, urządzenia dopasowującego lub bezzakłócenowego źródła zasilania.
- Należy upewnić się, aby nic nie leżało na kablach systemu komputerowego, oraz aby kable nie były umieszczone w miejscu, gdzie można byłoby na nie nadeptywać lub potykać się o nie.
- Nie należy rozlewać napojów ani innych płynów na system komputerowy.
- Nie należy wpychać żadnych przedmiotów do otworów systemu komputerowego, gdyż może to spowodować pożar lub porażenie prądem, poprzez zwarcie elementów wewnętrznych.
- System Komputerowy powinien znajdować się z dala od grzejników i źródeł ciepła. Ponadto, nie należy blokować otworów wentylacyjnych. Należy unikać kładzenia luźnych papierów pod komputer oraz umieszczania komputera w ciasnym miejscu bez możliwości cyrkulacji powietrza wokół niego.

Index

A

Add Logical Drive (Add System Drive) option, 48
affinity. See LUN affinity and system drive.
airflow, 4, 10, 81
 filler panels in rack, 7
 See also temperature conditions, problems, solutions
alarm, 14, 16
 and troubleshooting, 80
 at power-on, 70
 mute pushbutton, 14, 16
 muting, 81
 cannot be muted, 83
 thermal, 82
 with all ESI/ops panel LEDs and all disk drive LEDs, 83
 with LED on power supply/cooling module, 76
 with power supply/cooling module LED and ESI/ops panel power supply/cooling module LED, 82
 with RAID module fault LED, 70
 with system status and ESI/ops panel power supply/cooling module LEDs, 75
 with system status and power supply/cooling module LEDs, 75
 with system status LED, 75
antitamper lock on disk drive module, 24
 disengaging, 87
 locking, 95
audible alarm. See alarm.

automatic rebuild, 66-67
availability, configuring for maximum, 57-59

B

bay numbering, 26-27
 enclosure, 26
 tower, 27
breaker in rack, 38

C

cable
 fibre channel, 104
 power, 31-32, 35-37, 40
 shielded, 109
 to host, 29
caching
 disabled, 50
 enabled, 47, 49, 58
 write-back, 47, 49, 58
 write-through, 50
CAP, 55-59
capacity, availability, and performance. See CAP.
capacity, configuring for maximum, 56
chassis grounding, 30
Class A, Class B, 105-112
COD, 44-45

configuration, 59-62, ??-63
RAID, 59
selecting RAID level, 55-59
configuration on disk, 44-45
controller parameters, take effect after reboot, 61
conventions, xv
current limit, 29
customer service, xv

D

data caching, 49-50
See also caching.
dead drive state, 65
depth, tower, 100
deskside. See tower.
device driver, 70
disk drive filler, 25
disk drive module, 23-25
adding, 86-95
antitamper lock, 24
disengaging, 87
locking, 95
capacity, on label, 23
carrier, 24
dual-ported, 24, 60
filler, 25
keying, none, 89, 90
label, 23
LEDs, 24
and troubleshooting, 79-80
at power-on, 71
loop, 24
maximum per loop
RAID and expansion enclosures, 59
tower, 59
order for installing, 86
replacing, 96-97

required, 43, 86
locations, 25
requirements for MORE options, 46
seating, 93
serial number, 23
specifications, 104
states, 65-66
status LED on ESI/ops panel, 74
total addressed by RAID controller, 2
disk. See disk drive module.
documentation, other, xiv
door of rack, opening and closing, 28
drive roaming, 45
and hot spare, 97
drive state reporting, 65-66
dummy module
and airflow, 4, 10
disk drive, 25

E

electromagnetic emissions, 109
electrostatic discharge, 109
enclosure, 2-8
bay numbering, 26
expansion, 2
cabling, 8
in rack, 5-8
system ID, 21, 70
height, 4, 99
housing, 4
RAID, 2
cabling, 8
in rack, 5-8
system ID, 21, 70
total in rack, 7
weight, 99

environmental
 device, 66
 drive state, 66
 requirements, 101
ESI/ops panel module, 13-16
 alarm, 14, 16
 and SES, 69
 key switch, 14, 15
 LEDs, 14-15
 and SES, 44
 and troubleshooting, 72-75
Expand Logical Drive (Expand System Drive)
 option, 48-49
expansion enclosure. See enclosure, expansion.

F

fan
 failure, 75
 and thermal shutdown, 82
 fault, 74
 increased noise level, 81
 slowing, 70, 80
 speeding up, 81
FC-AL loop resiliency circuit (LRC) I/O module. See
 I/O module.
filler panels in rack, 7
front dummy fascia, 25

G

GAM, 52, 52-53
Global Array Manager, 52, ??-53
grounding, 30
grounding, checking
 rack, 37
 tower, 32

H

HBA, cabling to, 29
height
 enclosure, 99
 rack, 99
 tower, 100
host
 cabling to, 29
 does not recognize storage system, 70
hot spare, 45, 96
 and automatic rebuild, 66
 and availability, 57
 and drive roaming, 97
 creating, 67
hot swap disk drive replacement, 67, 85
hub
 cabling to, 29
 in rack, 7
humidity requirements, 101

I

ID
 selector switch, 21
 and troubleshooting, 70
 system
 expansion enclosure, 21, 70
 RAID module, 21, 70
 tower, 21
I/O module, 19-22
 and loops, 19, 60
 LEDs, 22
 and troubleshooting, 77
 ports, 22

J

- JBOD, 54
 - and availability, 59
 - and capacity, 56
 - and performance, 59

K

- key switch, 14, 15
 - starting enclosure, 37
 - starting tower, 33
 - stopping enclosure, 40
 - stopping tower, 40
 - when power is restored, 41

L

- label on disk drive module, 23
- LED
 - and troubleshooting, 71-80
 - checking at power-on, 70
 - disk drive module, 24
 - and troubleshooting, 79-80
 - ESI/ops panel module, 14-15
 - and SES, 44
 - and troubleshooting, 72-75
 - I/O module, 22
 - and troubleshooting, 77
 - power supply/cooling module, 12
 - and troubleshooting, 76
 - RAID controller module, 18, 19
 - and troubleshooting, 78-79
- logical drive. See system drive.
- loop, 19, 24, 62, 63
 - configuration, 59-62, ??-63
- LUN
 - affinity and system drive, 64

- integrity and disk drive module failure, 96
- mapping, 64

M

- manuals, other, xiv
- model number, 108
- module, dummy, 4, 10
- MORE, 46-49
 - caveats, 47-48
 - disk drive requirements and restrictions, 46
- Mylex Online RAID Expansion. See MORE.

O

- online drive state, 65
- overheating, overtemperature
 - See temperature conditions, problems, solutions.

P

- performance, configuring for maximum, 59
- port
 - I/O module, 22
 - RAID controller module
 - to host, 29
- power
 - checking, 34, 39
 - cord, 31-32, 35-37
 - requirements, 102
 - PDU, 102
 - rack, 102
 - restored after failure, and key switch positions, 41
 - voltage requirement
 - rack, 35
 - tower, 31

power supply/cooling module, 12
 LED, 12
 and troubleshooting, 76
 on ESI/ops panel, 74, 75

powering
 off

 enclosure in rack, 40
 rack, 40
 tower, 40

 on

 alarm sounds, 70
 checking system status, 34, 39
 problems, 70
 rack, 35-39
 tower, 31-35

programmable LUN mapping, 64

R

rack, 5-8

 breaker, 38
 cabling, 35-37
 height, 99
 in U, 7
 power requirements, 102
 powering on, 35-39
 rear door, opening and closing, 28
 weight, 99

RAID Assist, 52

 and hot spare, 67

RAID controller module, 17-19

 and loops, 19, 60
 drives addressed, 2
 LEDs, 18, 19
 Activity, 18, 79
 and troubleshooting, 78-79
 specifications, 103
 write cache size, 58

RAID enclosure, 2

 cabling, 8
 housing, 4
 in rack, 5-8

RAID level

 RAID 0, 53

 and availability, 58
 and capacity, 56
 and disk failure, 96
 and performance, 59

 RAID 0+1, 54

 and automatic rebuild, 66
 and availability, 59
 and capacity, 56
 and disk failure, 96
 and performance, 59
 and rebuild, 65

 RAID 1, 53

 and automatic rebuild, 66
 and availability, 58
 and capacity, 56
 and disk failure, 96
 and performance, 59
 and rebuild, 65

 RAID 3, 54

 and automatic rebuild, 66
 and availability, 58
 and capacity, 56
 and disk failure, 96
 and performance, 59
 and rebuild, 65
 maximum drives, 59

 RAID 30, 54

 RAID 5, 54

 and automatic rebuild, 66
 and availability, 58
 and capacity, 56
 and disk failure, 96
 and performance, 59
 and rebuild, 65
 maximum drives, 59

- RAID 50, 54
- RAID 6. See RAID 0+1.
- RAID 7, 54
 - See also JBOD.
- strategy for selecting, 55-59
- supported, 53-55
- rebuild
 - automatic, 66-67
 - drive state, 65
- regulatory information, 105-112
- RUN position of key switch
 - at power-on, 33, 38
 - when power restored, 41

- S**
- SCSI Enclosure Services. See SES.
- serial number, 23
- server does not recognize storage system, 70
- service, xv
- SERVICE position of key switch
 - at power-on, 34, 38
 - service mode, 34, 38
 - standby mode, 34, 38
 - when power restored, 41
- SES, 15, 43-44, 50
 - and ESI/ops panel LEDs, 44
 - and ESI/ops panel module, 69
- SGL Fibre Channel Hub
 - cabling to, 29
 - in rack, 7
- SGL switch
 - cabling to, 29
 - in rack, 7
- SGL, contacting, xv
- shielded cables, 109
- shipping dimensions, 99
- shutdown, thermal, 82

- slot numbering, 26-27
 - enclosure, 26
 - tower, 27
- specifications
 - disk drive module, 104
 - RAID controller module, 103
 - storage system physical, 99-100
- standby
 - drive state, 65
 - LED on ESI/ops panel, 74
 - mode, 34, 38, 40
 - pushbutton
 - and powering off, 40
 - and powering on in service mode, 34, 38
 - disabled in run mode, 34, 38
- starting storage system. See powering on.
- stopping storage system. See powering off.
- support, xv
- switch
 - cabling to, 29
 - in rack, 7
- system drive, 46-49, 64-65
 - and data caching, 64
 - and LUN affinity, 64
 - maximum, 60, 64
 - size, 64
- system status LED, 74, 75

- T**
- temperature
 - conditions, problems, solutions, 81-83
 - requirements, 101
- test mode, 83
- thermal
 - alarm, 82
 - control, 81
 - shutdown, 82

- tower, 9-10
 - adapting for rackmounting, 10
 - bay numbering, 27
 - cabling, 31-32
 - depth, 100
 - height, 100
 - powering on, 31-35
 - system ID, 21
 - weight, 100
 - width, 100

U

- U, height, 100
- unconfigured drive state (unconfigured location), 66
- uninterruptible power supply, 74
- UPS, 74

V

- voltage, 102
 - out of range, 70
 - requirement
 - rack, 35
 - tower, 31

W

- weight
 - components, 100
 - enclosure, 99
 - rack, 99
 - shipping, 99
 - tower, 100
- width, tower, 100
- worldwide name, 23

- write cache
 - disable, 50
 - enable, 47, 49, 58
 - size, RAID controller, 58
- write-back caching, 47, 49, 58
 - and system drive, 64
- write-through caching, 50
 - and system drive, 64

