CHALLENGE® RAID Installation and Maintenance Instructions

Document Number 108-0128-006

Contributors

Written by Carolyn Curtis
Illustrated by Cheri Brown, Dany Galgani, Dan Young, and Carolyn Curtis
Production by Heather Hermstad
Engineering contributions by Sammy Wilborn, Albert Lui, Henry Ortiz, Carl Strasen, and Paul Tsien

CHALLENGE® RAID Installation and Maintenance Instructions
Document Number 108-0128-006

Silicon Graphics, Inc.
Mountain View, California

Contents

About This Guide ..... **xiii**

1. Challenge RAID Kit Contents **1-1**

2. Challenge RAID Chassis Tour **2-1**

2.1 SCSI-2 Interface 2-6
2.2 Storage-Control Processor (SP) 2-6
2.3 Disk-Drive Modules 2-7
2.4 Power Supply Module (Voltage Semi-Regulated Converter) 2-8
2.5 Fan Module 2-9
2.6 Optional Battery Backup Unit (BBU) 2-10
2.7 Challenge RAID Configurations 2-11
2.7.1 Basic Configuration 2-11
2.7.2 Dual-Interface / Dual-Processor Configuration 2-12
2.7.3 Split-Bus Configuration 2-13
2.7.4 Dual-Bus / Dual-Initiator Configuration 2-13

3. Installing a Challenge RAID Storage System **3-1**

3.1 Setting Up a Challenge RAID Rack 3-2
3.2 Opening the Rear Door of a Challenge RAID Rack 3-5
3.3 Reseating Subassemblies 3-6
3.4 Labeling Disk Modules 3-7
3.5 Cabling Challenge RAID Configurations 3-9
3.5.1 Cabling the Basic Configuration 3-9
3.5.2 Cabling the Dual-Interface / Dual-Processor Configuration 3-13
3.5.3 Cabling the Split-Bus Configuration 3-14
3.5.4 Cabling the Dual-Bus / Dual-Initiator Configuration 3-15
3.6 Setting SCSI IDs for Storage-Control Processors 3-18
3.7 Cabling an Additional Challenge RAID Storage System (or Chassis Assembly) on a SCSI-2 Bus Channel 3-20

3.8 Setting SCSI IDs for SPs in a Second Challenge RAID on the Same SCSI-2 Bus Channel 3-25
3.9 Powering On Challenge RAID 3-25

4. Installing the Challenge RAID Software and Configuring Disks 4-1
4.1 Installing the Challenge RAID Software 4-1
4.2 Creating a Configuration File 4-2
4.3 Starting the Agent 4-2
4.4 Using *raidcli* Commands 4-3
4.5 Getting the Device Name With *raidcli getagent* 4-4
4.6 Unbinding Disks 4-4
4.7 Planning the Bind 4-5
4.7.1 Filesystem Capacity Considerations for RAID 4-5
4.7.2 Determining RAID Levels 4-5
4.7.3 Planning a RAID-3 Bind 4-7
4.8 Allocating Memory for RAID-3 LUNs 4-10
4.9 Binding Disks Into RAID Units 4-11
4.10 Enabling Command-Tagged Queuing 4-14
4.11 Understanding Dual Processors, Load Balancing, and Device Names 4-17
4.12 Setting Up Caching 4-17
4.12.1 Setting Cache Parameters 4-18
4.12.2 Upgrading Challenge RAID to Support Caching 4-20
4.12.3 Changing Caching Parameters 4-20

5. Maintaining and Upgrading Disk Modules 5-1
5.1 Handling FRUs 5-3
5.1.1 Avoiding Electrostatic Discharge (ESD) Damage 5-3
5.1.2 Precautions for Removing, Installing, or Storing FRUs 5-3
5.2 Identifying and Verifying a Failed Disk Module 5-4
5.3 Unbinding the Disk 5-5
5.4 Replacing a Disk Module 5-6
5.4.1 Removing a Failed Disk Module 5-6
5.4.2 Installing a Replacement Disk Module 5-9
5.4.3 Updating the Disk Module Firmware 5-11
5.5 Installing an Add-On Disk Module Array 5-12
5.5.1 Inserting the New Disk Module Array 5-12
5.5.2 Creating Device Nodes and Binding the Disks 5-15

6. Maintaining and Upgrading SPs 6-1
6.1 Opening and Closing the Fan Module 6-2

6.2 Locating the SPs ... 6-3
6.3 Replacing or Adding an SP ... 6-5
6.3.1 Using the Auto-Assign Capability in the Case of a Failed SP ... 6-5
6.3.2 Removing an SP ... 6-9
6.3.3 Installing an SP ... 6-11
6.4 Installing or Removing SP Memory Modules (SIMMs) ... 6-13
6.4.1 Installing SP Memory Modules ... 6-15
6.4.2 Removing SP Memory Modules ... 6-18
6.5 Upgrading FLARE 7.xx to FLARE 8.xx for Sauna ... 6-18
6.6 Upgrading Sauna to Phoenix ... 6-19
6.6.1 Upgrading the FLARE Code ... 6-19
6.6.2 Moving Sauna SIMMs to the Phoenix SP Board ... 6-20
6.6.3 Testing the Upgrade ... 6-22
6.7 Upgrading to FLARE 9.xx in a Storage System With RAID-3 LUNs 6-23

7. Maintaining and Upgrading Other Components ... 7-1
7.1 Opening and Closing the Fan Module ... 7-4
7.2 Replacing or Adding a Power Supply Module ... 7-5
7.2.1 Replacing a Power Supply Module ... 7-5
7.2.2 Adding a Power Supply Module ... 7-6
7.3 Replacing a Fan Module ... 7-6
7.4 Replacing or Adding a Battery Backup Unit ... 7-8
7.5 Replacing the Power Distribution Unit in a Challenge RAID Rack .. 7-10
7.6 Installing an Additional Chassis Assembly in a Challenge RAID Rack ... 7-13

A. Technical Specifications ... A-1

B. The raidcli Command-Line Interface ... B-1
B.1 bind ... B-2
B.2 chglun ... B-5
B.3 clearlog ... B-8
B.4 clearstats ... B-8
B.5 firmware ... B-9
B.6 getagent ... B-10
B.7 getcache ... B-11
B.8 getcontrol ... B-14
B.9 getcrus ... B-15
B.10 getdisk ... B-15
B.11 getlog ... B-18
B.12 getlun ... B-20

B.13 getsp B-24
B.14 setcache B-24
B.15 setstats B-26
B.16 trespass B-26
B.17 unbind B-27

C. Troubleshooting C-1
C.1 Determining Firmware Failure C-1
C.2 Restarting the Agent C-2
C.3 Reconfiguring the Challenge RAID Storage System as System Type 9 ..... C-2
C.4 Reenabling Command-Tagged Queuing C-5
C.5 Determining Caching Problems C-6
C.6 Reassigning LUN Ownership C-6

Index Index-1

Figures

Figure 2-1 Deskside Challenge RAID Storage System 2-1
Figure 2-2 Challenge RAID Rack 2-2
Figure 2-3 Disk Module Locations: Front View 2-3
Figure 2-4 Field-Replaceable Units: Deskside System 2-4
Figure 2-5 Field-Replaceable Units: Rack System 2-5
Figure 2-6 AMD-Based SP Removed From Challenge RAID Storage System 2-6
Figure 2-7 Typical Disk-Drive Module Removed from the Challenge RAID Storage System 2-7
Figure 2-8 Power Supply Module (VSC) Removed From Challenge RAID Storage System 2-8
Figure 2-9 Partially Opened Fan Module on Back of Challenge RAID Storage System 2-9
Figure 2-10 Battery Backup Unit (BBU) Removed From Challenge RAID Storage System 2-10
Figure 2-11 Basic Configuration 2-11
Figure 2-12 Dual-Interface / Dual-Processor Configuration 2-12
Figure 2-13 Split-Bus Configuration 2-13
Figure 2-14 Dual-Bus / Dual-Initiator Configuration Example 2-14
Figure 3-1 Ramp and Pallet (Shown With Bubblewrap Removed for Clarity) 3-2
Figure 3-2 Unlocking a Challenge RAID Rack Side Panel 3-3
Figure 3-3 Challenge RAID Rack Pallet Bracket Locations 3-4
Figure 3-4 Opening the Challenge RAID Rack Rear Door 3-5
Figure 3-5 Challenge RAID Storage System Power Switch 3-6
Figure 3-6 Unlocking the Fan Module 3-6
Figure 3-7 Disk Module Label 3-7
Figure 3-8 Marking the Label for Disk Module B0 3-7
Figure 3-9 Disk Drive Locations 3-8
Figure 3-10 Basic Configuration 3-9
Figure 3-11 SCSI-2 Bus Connectors on Back of Challenge RAID Chassis 3-10
Figure 3-12 Cabling the Basic Configuration 3-11
Figure 3-13 Connecting a SCSI Bus Cable 3-11
Figure 3-14 Connecting a SCSI Terminator Plug 3-11
Figure 3-15 Connecting SCSI Bus Cables to a Challenge RAID

Storage System and a Challenge S 3-12
Figure 3-16 Dual-Interface / Dual-Processor Configuration Example 3-13
Figure 3-17 Split-Bus Configuration Example 3-14
Figure 3-18 Connecting Challenge RAID: Split-Bus Configuration 3-15
Figure 3-19 Dual-Bus / Dual-Initiator Configuration Example 3-15
Figure 3-20 Connecting Challenge RAID: Dual-Bus / Dual-Initiator Configuration 3-16
Figure 3-21 Connecting SCSI Cables to a Challenge RAID and Two Challenge L Servers 3-17
Figure 3-22 Setting an SP's SCSI ID (Back of Challenge RAID) 3-18
Figure 3-23 Computing SCSI Cable Length Example: Single-Host Configuration Only 3-20
Figure 3-24 Two-Storage-System Cable-Planning Diagram 3-21
Figure 3-25 Example Two-Storage-System Cable-Planning Diagram: Second Challenge RAID 3-22
Figure 3-26 Connecting SCSI Buses A and B to Two RAID Chassis Assemblies in a Challenge RAID Rack Storage System 3-24
Figure 3-27 Connecting the Power Cord (Back of Challenge RAID Storage System) 3-25
Figure 3-28 Turning On the Storage System Power 3-26
Figure 3-29 Challenge RAID Indicator Lights 3-26
Figure 3-30 Unlocking the Fan Module 3-27
Figure 3-31 Enabling a Sauna SP's Power 3-27
Figure 4-1 Example RAID-3 Memory Allocation: One-SP System With RAID-3 LUNs Only 4-7
Figure 4-2 Example RAID-3 Memory Allocation for Failover System With RAID-3 LUNs Only 4-8
Figure 4-3 Example RAID-3 Memory Allocation for Failover System 4-9
Figure 5-1 Location of Disks (Front of Challenge RAID) 5-2
Figure 5-2 Disk Module Status Lights 5-4
Figure 5-3 Attaching the ESD Clip to the ESD Bracket on the Deskside Storage System 5-7
Figure 5-4 Attaching the ESD Clip to the ESD Bracket on a Rack Storage System 5-7
Figure 5-5 Pulling Out a Disk Module 5-8
Figure 5-6 Removing a Disk Module 5-8
Figure 5-7 Engaging the Disk Module Rail 5-9
Figure 5-8 Engaging the Disk Module Guide 5-10
Figure 5-9 Inserting the Replacement Disk Module 5-10
Figure 5-10 Marking the Label for Disk Module B0 5-13
Figure 5-11 Disk Drive Locations 5-14
Figure 5-12 Engaging the Disk Module Rail 5-14
Figure 5-13 Engaging the Disk Module Guide 5-15
Figure 6-1 Unlocking the Fan Module 6-2

Figure 6-2 Opening the Fan Module 6-2
Figure 6-3 Field-Replaceable Units (FRUs): Back of Challenge RAID Deskside System 6-3
Figure 6-4 Field-Replaceable Units: Rack System 6-4
Figure 6-5 Equipment View 6-7
Figure 6-6 Summary View 6-8
Figure 6-7 Disabling a Sauna SP 6-9
Figure 6-8 Ejectors for SP 6-10
Figure 6-9 Installing an SP 6-12
Figure 6-10 SIMM 1 and 3 Connectors on the Sauna SP Board 6-14
Figure 6-11 SIMM 1 and 2 Connectors on the Phoenix SP Board 6-14
Figure 6-12 Inserting an SP Memory Module: Sauna 6-16
Figure 6-13 Locking In an SP Memory Module: Sauna 6-16
Figure 6-14 Installing an SP Memory Module: Phoenix 6-17
Figure 6-15 Locking In an SP Memory Module: Phoenix 6-17
Figure 6-16 Removing a Memory Module From the SIMM Connector 6-18
Figure 6-17 Removing a Memory Module From the SIMM Connector 6-20
Figure 6-18 SIMM 1 and2 Connectors on the Phoenix SP Board 6-21
Figure 6-19 Inserting an SP Memory Module on the Phoenix SP Board 6-21
Figure 6-20 Locking the SIMM into the Phoenix SIMM Connector 6-22
Figure 7-1 Field-Replaceable Units (FRUs): Back of Challenge RAID 7-2
Figure 7-2 Field-Replaceable Units: Rack System 7-3
Figure 7-3 Unlocking the Fan Module 7-4
Figure 7-4 Opening the Fan Module 7-4
Figure 7-5 Removing a Power Supply Module (VSC) 7-5
Figure 7-6 Adding a Power Supply Module 7-6
Figure 7-7 Disconnecting the Fan Module Power 7-7
Figure 7-8 Unlatching the Fan Module 7-7
Figure 7-9 Lifting Out the Fan Module 7-7
Figure 7-10 Removing a Battery Backup Unit or Filler Panel 7-8
Figure 7-11 Removing or installing a Battery Backup Unit or Filler Panel 7-9
Figure 7-12 Clearing Access to the Power Distribution Unit in the Challenge RAID Rack 7-10
Figure 7-13 Removing the Power Distribution Unit Cover 7-11
Figure 7-14 Removing the Power Distribution Unit (Bird's-Eye View) 7-12
Figure 7-15 Installing an Additional Chassis Assembly in the Challenge RAID Rack 7-13
Figure 7-16 Securing the Chassis Assembly to the Rack 7-14
Figure 7-17 Attaching the Front Bezel 7-14
Figure B-1 Disk Module Locations B-16
Figure C-1 GridMgr Main Menu C-3
Figure C-2 Example Presentation Utility Screen C-5

Tables

Table 1-1 Challenge RAID Components 1-1
Table 1-2 Disk Module Part Numbers 1-1
Table 1-3 Challenge RAID Storage System Marketing Codes 1-2
Table 1-4 Deskside Power Cable Destination Kits 1-2
Table 1-5 Rack Power Cable Destination Kits 1-2
Table 1-6 Field-Replaceable Units 1-3
Table 3-1 Challenge RAID Configurations 3-9
Table 3-2 Setting SCSI IDs 3-19
Table 3-3 Recommended SCSI ID Settings for SPs in Second Challenge RAID on a SCSI-2 Bus Channel 3-25
Table 4-1 Binding LUNs 4-6
Table 4-2 Maximum CTQ Depths per LUN: Single-Hosted SPs 4-15
Table 4-3 Maximum CTQ Depths per LUN: Dual-Hosted SPs 4-15
Table 4-4 CTQ Performance Benefits for 2 K Random Read, 16 Threads 4-15
Table 4-5 Read and Write Partition Sizes (in MB): Firmware Revision 8.50 and Higher (No RAID-3 Memory Allocated) 4-19
Table 4-6 Read and Write Partition Sizes in MB: Firmware Revision 8.00-8.49 4-20
Table 5-1 Ordering Add-On Disk Module Sets 5-12
Table A-1 Deskside Challenge RAID Specifications A-1
Table A-2 Challenge RAID Rack Specifications A-2
Table B-1 *raidcli* Parameters B-1
Table B-2 Output of *raidcli getagent* B-10
Table B-3 Output of *raidcli getcache* B-12
Table B-4 Output of *raidcli getcrus* B-15
Table B-5 Output of *raidcli getdisk* B-17
Table B-6 *getlog* Error Codes B-18
Table B-7 Output of *raidcli getlun* B-22

About This Guide

The CHALLENGE® RAID storage system provides a compact, high-capacity, high-availability source of disk storage for the complete line of Silicon Graphics® CHALLENGE servers running IRIX™ 5.3 and 5.3 with XFS™: CHALLENGE S, CHALLENGE DM, CHALLENGE L, and CHALLENGE XL.

Note: For ease in reading, CHALLENGE is written as Challenge in the balance of this guide.

The Challenge RAID deskside storage system uses traditional or high-availability disk storage in as many as 20 disk modules that you or the customer can replace while the system continues to run. For even more storage, the Challenge RAID rack storage system offers up to four RAID chassis assemblies, each with as many as 20 disk modules. The chassis assemblies in a Challenge RAID rack can be connected to one or more SCSI buses on Challenge servers separately or in combination.

RAID levels 0, 1, 1_0 (0+1), 3, and 5 are supported, as well as disks configured as hot spares. In addition, the Challenge RAID storage system can provide storage-system caching.

Besides a command-line interface, a graphical user interface, RAIDGUI, is also available; it is documented extensively in the *CHALLENGE RAID Owner's Guide*, revision 007-2532-006 and later.

Note: This document is intended for use by Silicon Graphics System Support Engineers only.

Structure of This Document

This guide contains the following chapters:

Chapter 1 "Challenge RAID Kit Contents" lists the RAID cables and other components shipped with the unit.

Chapter 2 "Challenge RAID Chassis Tour" illustrates and describes all Challenge RAID components.

Chapter 3 "Installing a Challenge RAID Storage System" describes how to install, cable, and set up the storage system.

Chapter 4 “Installing the Challenge RAID Software and Configuring Disks” explains how to start the Challenge RAID software, use the command-line interface to group disks into RAID-5 units, and, if the customer desires, enable caching.

Chapter 5 “Maintaining and Upgrading Disk Modules” explains replacing and installing disk modules.

Chapter 6 “Maintaining and Upgrading SPs” explains replacement and upgrade procedures for storage-control processors, including FLARE code upgrades.

Chapter 7 “Maintaining and Upgrading Other Components” explains replacing and installing the remaining components, such as the power-supply module. It also explains how to install additional chassis assemblies in a Challenge RAID rack.

Appendix A “Technical Specifications” summarizes technical information for the Challenge RAID deskside storage-array system.

Appendix B “The raidcli Command-Line Interface” lists and explains all parameters of the *raidcli* command.

Appendix C “Troubleshooting” explains how to solve system problems.

An index completes this guide.

Note: It is advisable to have handy your copy of the *CHALLENGE/Onyx XL Rackmount Installation Instructions* (108-7042-020); Appendix D contains SCSI channel information. Depending on the Challenge server present at the site, you might also need the *CHALLENGE/Onyx Deskside Installation Instructions* (108-7039-020).

Conventions

These type conventions and symbols are used in this guide:

Helvetica Bold Hardware labels

Italics Executable names, filenames, IRIX commands, manual or book titles, new terms, program variables, tools, utilities, variable command-line arguments, variable coordinates, and variables to be supplied by the user in examples, code, and syntax statements

`Fixed-width type`
Error messages, prompts, and onscreen text

`Bold fixed-width type`
User input, including keyboard keys (printing and nonprinting); literals supplied by the user in examples, code, and syntax statements (*see also* <>)

“” (Double quotation marks) Onscreen menu items and references in text to document section titles

[] (Brackets) Surrounding optional syntax statement arguments

<> (Angle brackets) Surrounding nonprinting keyboard keys, for example, <**`Esc`**>, <**`Ctrl-D`**>

Chapter 1

Challenge RAID Kit Contents

Note: Before installing a Challenge RAID storage system, make sure that the installation site meets the AC power requirements and operating limits listed in Appendix A.

This chapter summarizes the contents of the Challenge RAID kit.

Note: If parts are missing, or if incorrect parts are included in the kit, please log a call with the Technical Assistance Center (TAC).

Table 1-1 lists components included with the unit.

Table 1-1 Challenge RAID Components

Component	SGI Part Number
5-foot SCSI external differential cable w / 2-56 jack screws	9290112
20-foot SCSI external differential cable w / 2-56 jack screws	9290111
Power cord for destination country	See Table 1-4
CD and document package:	013-1200-005
CD Challenge RAID software 2.1	813-0322-004
CD SCSI patches 8505XL	813-0412-001
CHALLENGE RAID Owner's Guide	007-2532-007

Table 1-2 lists disk module part numbers. Disk modules can be replaced by customers as well as by SSEs.

Table 1-2 Disk Module Part Numbers

Disk Module	Marketing Code	Part Number
Replacement 2 GB drive	P-S-RAID-1X2	9410117
Replacement 4 GB drive	P-S-RAID-1X4	9410118
Add-on five 2 GB drives	P-S-RAID-5X2	9410115
Add-on five 4 GB drives	P-S-RAID-5X4	9410116

Table 1-3 gives the marketing codes for the basic Challenge RAID storage system configurations.

Table 1-3 Challenge RAID Storage System Marketing Codes

Disk Array	Marketing Code	Part Number
Base array with five 2 GB drives	P-S-RAID-B5X2	9410100
Base array with five 4 GB drives	P-S-RAID-B5X4	9410101
Rack with no chassis assemblies	P-S-RAID-RACK	9470109
Export rack with no chassis assemblies	P-S-RAID-ERACK	9470110
Chassis assembly with five 2 GB drives	P-S-RAID-R5X2	9470107
Chassis assembly with five 4 GB drives	P-S-RAID-R5X4	9470108

Table 1-4 lists marketing codes for power cord destination kits for the deskside chassis.

Table 1-4 Deskside Power Cable Destination Kits

Marketing Code	Cable
DK-P4-001	115V (US and Japan)
DK-P4-002	220V (Europe: Austria, Belgium, Netherlands, Luxembourg, Finland, France, Germany, Norway, Sweden, Spain, Portugal, Greece, Iceland, Yugoslavia, Bolivia, Peru, Paraguay, Venezuela, Iran, Indonesia)
DK-P4-003	220V (UK, Singapore, Hong Kong, Malaysia)
DK-P4-004	220V (Denmark)
DK-P4-005	220V (Switzerland)
DK-P4-006	220V (Italy, Chile)
DK-P4-007	220V (India, South Africa)
DK-P4-008	220V (Israel)
DK-P4-009	220V (Australia)
DK-P4-010	115V (Japan, Korea, Taiwan)

Note: All destination kits include the SCSI Differential terminator plug with 2-56 jack screws (p / n 9660006).

Table 1-5 lists marketing codes for power cord destination kits for the rack.

Table 1-5 Rack Power Cable Destination Kits

Marketing Code	Cable
DK-T4-001	115V (US)
DK-T4-004	220V (international)

Table 1-6 lists field-replaceable units (FRUs), that is, units only a Silicon Graphics SSE is qualified to replace.

Table 1-6 Field-Replaceable Units

Component	Marketing Code	Part Number
Storage-control processor (SP), no cache	P-S-RAID-SP0	9470119
Storage-control processor (SP)	P-S-RAID-SP	9470111
Storage-control processor (SP) with additional 8 MB mirroring	P-S-RAID-SP8	9470112
Storage-control processor (SP) with 32 MB mirroring	P-S-RAID-SP32	9470120
Storage-control processor (SP) with 64 MB mirroring	P-S-RAID-SP64	9470113
Power supply module (VSC)	P-S-RAID-PWR	9470114
Power receptacle (rack only)	N / A	005-38429
Power distribution unit for rack (domestic only)	N / A	005-042667
Rip controller	N / A	005-041325
Midplane	N / A	005-041324
SCSI internal I / O cable	N / A	005-040575
Storage-control processor (SP)	P-S-RAID-SP	9400103
Battery backup unit (BBU) (optional)	P-S-RAID-BBU	9341100
Fan module with 6 fans	P-S-RAID-FAN	9310103
Add-on 8 MB mirrored cache	P-S-RAID-C8	9390100
4 MB SIMM (included in 8 MB mirrored cache)	N / A	variable
Add-on 64 MB mirrored cache	P-S-RAID-C64	9390101
16 MB SIMM (included in 64 MB mirrored cache)	N / A	variable

Chapter 2

Challenge RAID Chassis Tour

This chapter describes the hardware components that make up the Challenge RAID storage system and the different ways to connect the storage system to a Challenge server. Figure 2-1 is an external view of the deskside version.

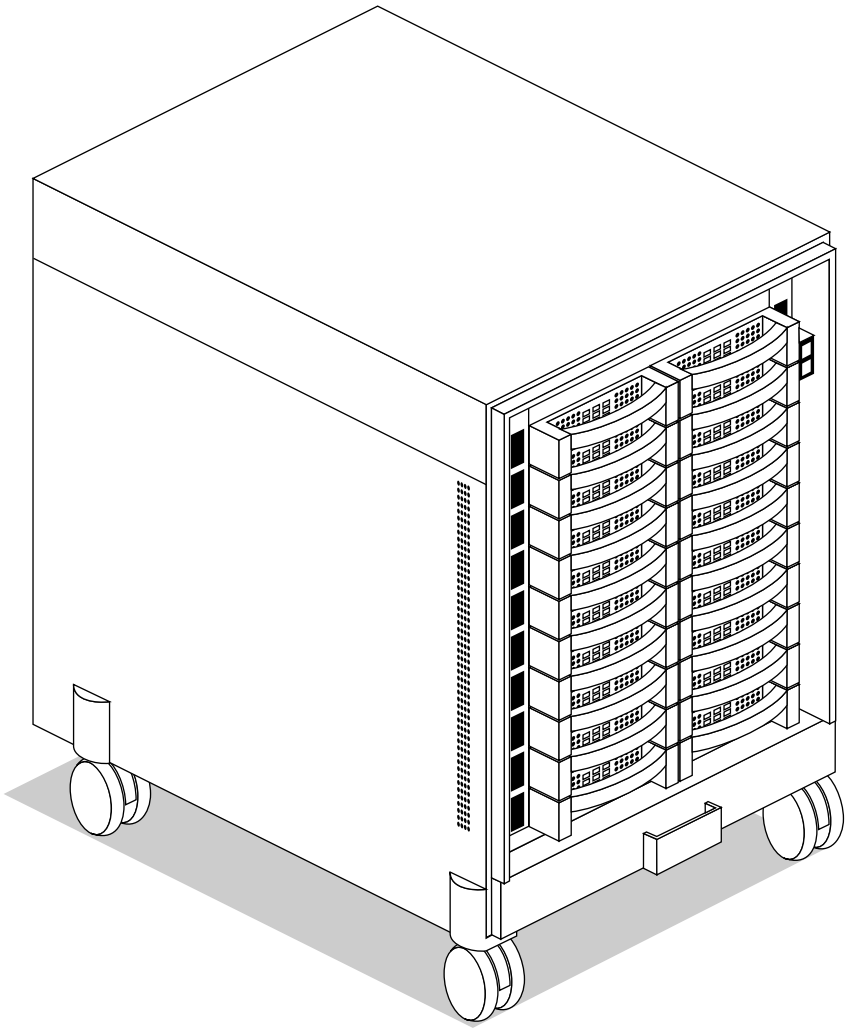

Figure 2-1 Deskside Challenge RAID Storage System

Note: This guide is written for Silicon Graphics qualified System Support Engineers. Only SSEs should install or replace any components other than disk modules.

Figure 2-2 is an external view of the Challenge RAID rack, with the maximum of four chassis assemblies installed. Each chassis assembly in a Challenge RAID rack corresponds to one deskside Challenge RAID chassis.

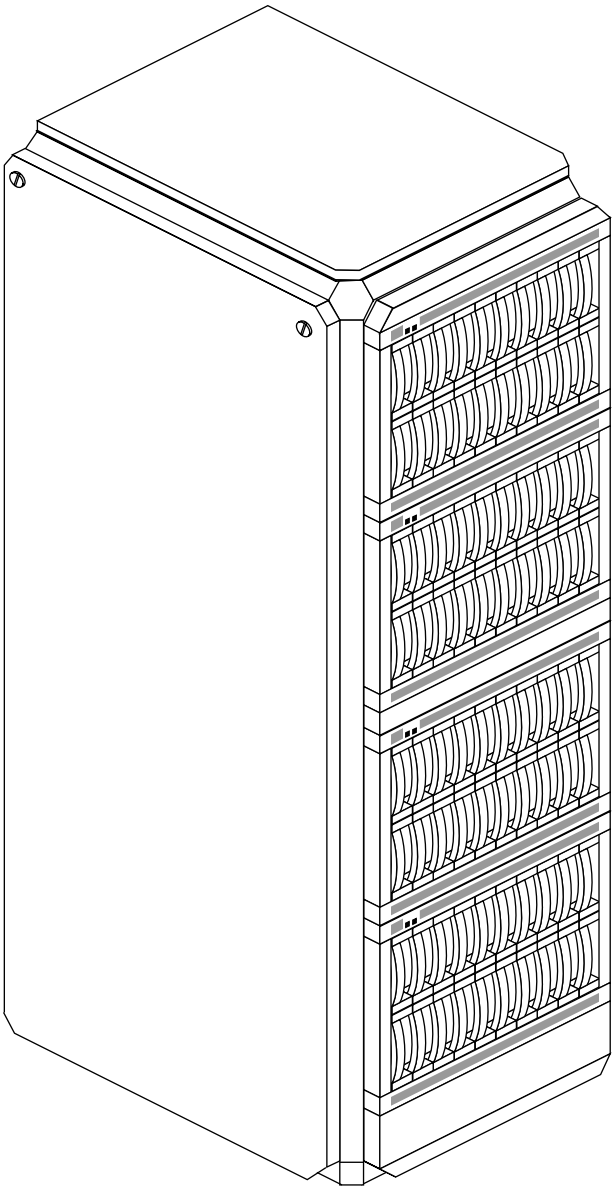

Figure 2-2 Challenge RAID Rack

The Challenge RAID storage system consists of a SCSI-2 (small computer system interface 2) interface in the host and a storage-system chassis assembly that contains compartments for these modules:

- one or two disk-array storage-control processors (SPs)
- five to twenty disk-drive modules in groups of five
- one fan module
- two or three power supplies (voltage semi-regulated converters)
- battery backup unit (optional)

For a Challenge RAID storage system to support storage-system caching, it must have these components:

- two SPs, each with 8 or 64 MB of memory
- a battery backup unit
- disk modules in compartments A0, B0, C0, D0, and E0

Figure 2-3 diagrams the location of the Challenge RAID disks, which can be replaced by the customer or by an SSE.

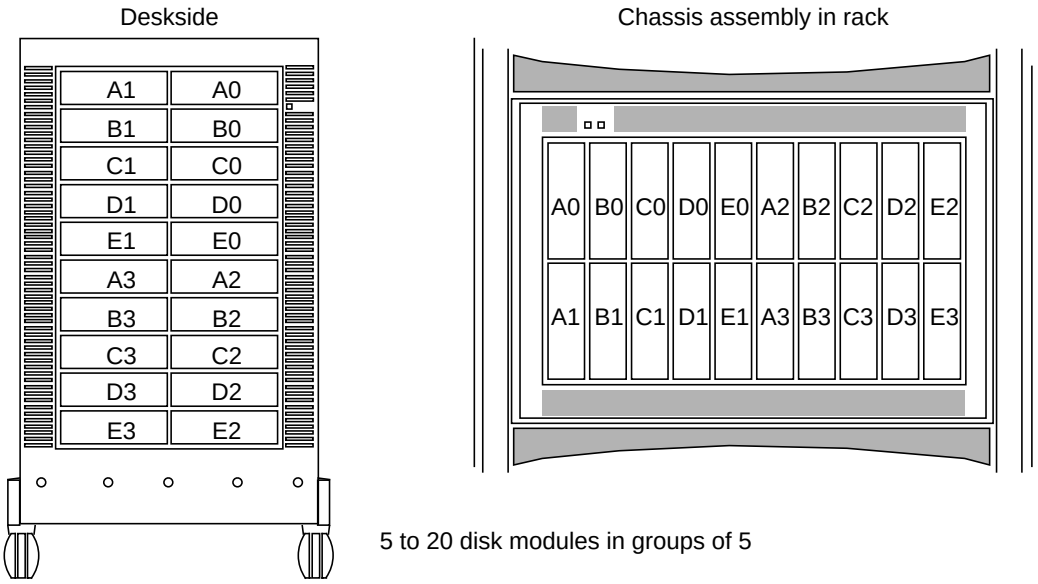

Figure 2-3 Disk Module Locations: Front View

Figure 2-4 shows the Challenge RAID field-replaceable units (FRUs) in the deskside Challenge RAID storage system. Only SSEs replace these.

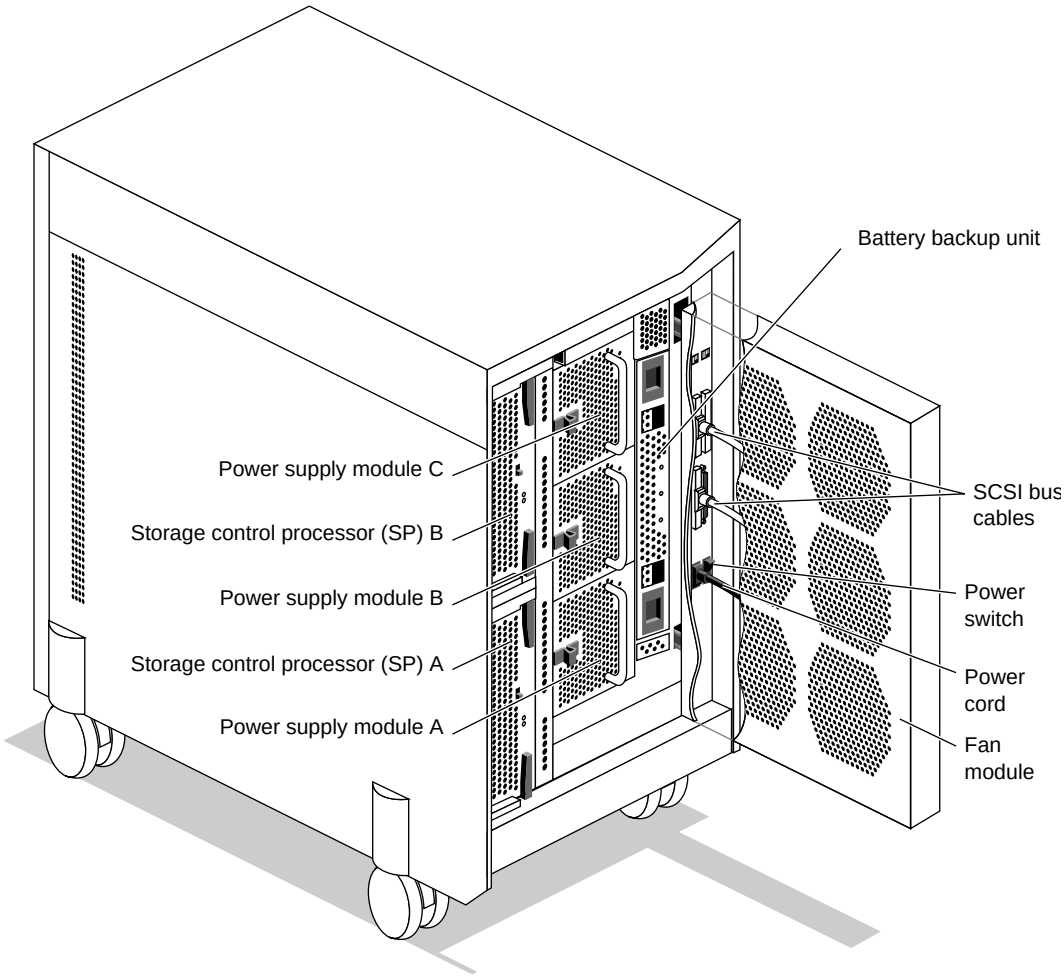

Figure 2-4 Field-Replaceable Units: Deskside System

See Chapter 3 for instructions on opening the storage system.

Figure 2-5 shows the Challenge RAID field-replaceable units (FRUs) in one RAID chassis assembly of a Challenge RAID rack storage system. Only SSEs replace these.

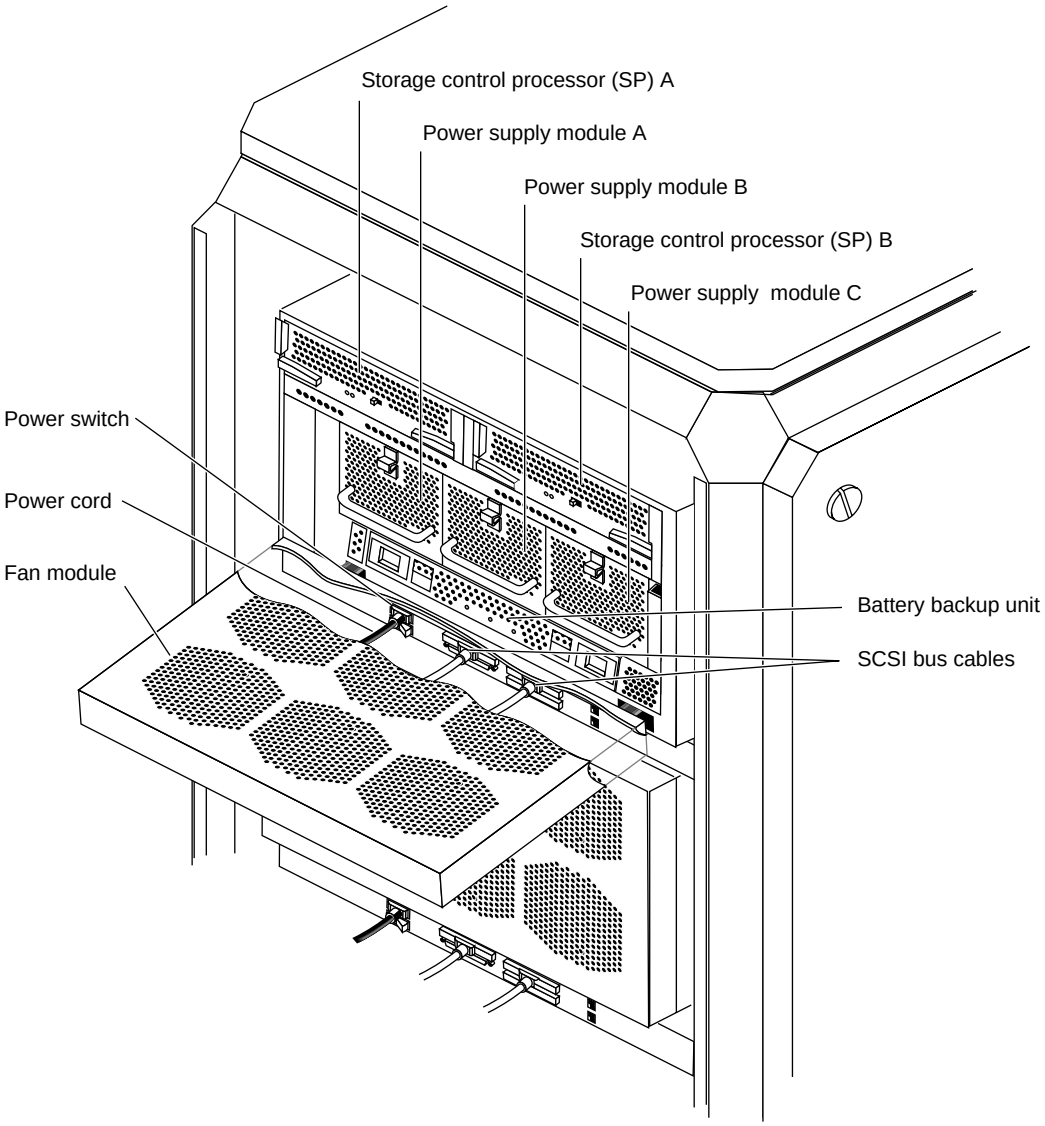

Figure 2-5 Field-Replaceable Units: Rack System

See Chapter 3 for instructions on opening the storage system.

2.1 SCSI-2 Interface

The SCSI-2 interface (or adapter, or controller) is a printed-circuit board that resides in an I/O slot in the Challenge cabinet. It connects to an SP in the storage-system cabinet by a SCSI-2 differential bus, and transfers data between the Challenge memory and the SCSI-2 bus. The interface is included in the Challenge system or can be ordered as a kit.

2.2 Storage-Control Processor (SP)

The storage-control processor consists of a printed-circuit board with two or four memory modules (SIMMs) and a bezel with status lights, and latches to secure the SP in place. The memory modules provide either 8 or 64 MB of SP memory. The SPs are visible when you swing away the fan module from the back of the storage system.

Figure 2-6 shows the AMD®-based SP (Sauna). The PowerPC®-based SP (Phoenix) lacks a power switch, and the SIMM connectors are aligned differently. SIMM connectors on the two types of SP are explained in detail in "Installing or Removing SP Memory Modules (SIMMs)" in Chapter 6.

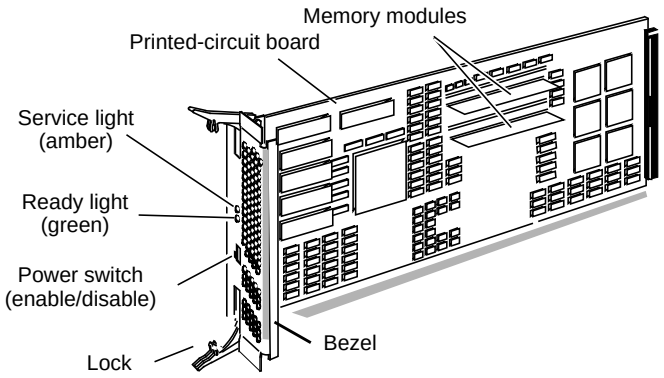

Figure 2-6 AMD-Based SP Removed From Challenge RAID Storage System

Two status lights indicate the following:

- ready light (green): lights while the SP is powered on and operating normally
- service light (amber): lights when the SP is not working properly

Challenge RAID can support two SPs. Caching requires two SPs, each with the same amount of memory, either 8 or 64 MB. Minimum system memory allowed is 2 MB for Sauna, and 4 MB for Phoenix.

Note: In storage systems with two SPs, the SPs must be of the same type (Sauna or Phoenix). For caching and failover, the SPs must have the same amount of memory.

The SP processes data written to or read from the disk modules, and controls the disk modules in the storage system through a synchronous SCSI-2 bus. It has five internal SCSI-2 buses, each supporting as many as four disk modules for a total of 20 disk e modules.

If the Challenge RAID has one SP, you can install a second one while the storage system is running and configure it into the system. When both SPs are installed and configured, you can replace either SP while the storage system is running.

Note: Never attempt to replace any SP components except memory modules (SIMMs).

2.3 Disk-Drive Modules

A disk-drive module, also called a disk module, consists of a disk drive, a power regulator board, internal cabling, and a plastic carrier. The carrier has a handle and guides for inserting and holding the module in the storage system's chassis. A label attached to the carrier's top shows the drive module's model number and capacity. Figure 2-7 shows a disk-drive module.

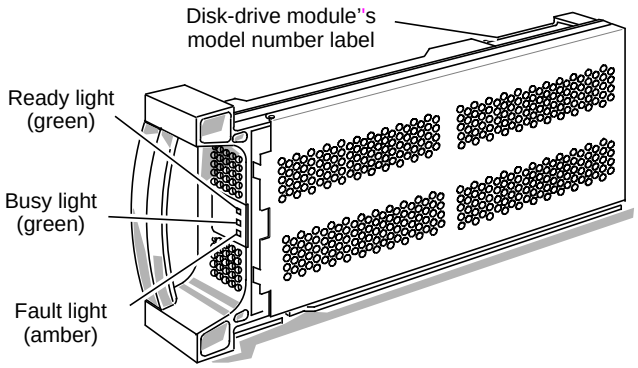

Figure 2-7 Typical Disk-Drive Module Removed from the Challenge RAID Storage System

Three status lights on the module indicate the following:

- ready light (green): lights while the disk-drive module is powered up and ready for use
- busy light (green): lights while the drive is in use; for example, during formatting or user I/O operations
- fault light (amber): lights when the module is shut down by the SP because the module failed; also lights after you replace the drive, while the replacement drive spins up to speed

Caution: Use only Challenge RAID disk modules to replace failed disk modules. Challenge RAID disk modules contain proprietary firmware that the storage system requires for correct functioning. Using any other disks, including those from other Silicon Graphics systems, can cause failure of the storage system.

You can remove or install any one module within an array group while the storage system is running.

Note: Never open a disk-drive module or attempt to replace any of its internal components. Never remove more than one disk module or filler module at a time. Always use ESD precautions when handling disk modules.

2.4 Power Supply Module (Voltage Semi-Regulated Converter)

The power supply modules, or voltage semi-regulated converters (VSCs) convert the installation site's AC line voltage to the 48, 24, and 18 VDC required to power the modules in the storage--system chassis. A storage system normally contains two power supply modules. Three power supply modules provide high-availability operation and might be required for certain Challenge RAID configurations.

The power supplies are visible when you swing open the fan module on the back of the storage system. Figure 2-8 shows a power supply module.

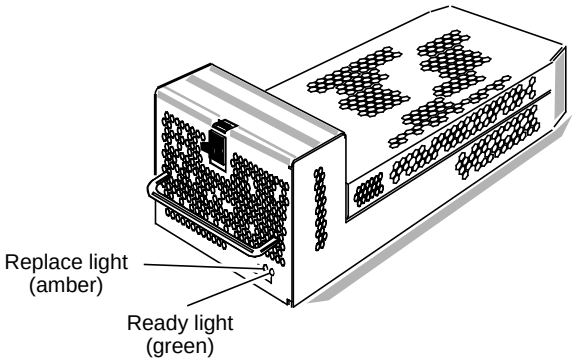

Figure 2-8 Power Supply Module (VSC) Removed From Challenge RAID Storage System

Two status lights on the power supply module indicate the following:

- ready light (green): lights while power supply is operating normally
- replace light (amber): lights when the SP determines that the power supply has failed

You can remove or install a power supply while the storage system is running; however, a minimum of two or three power supply modules is required, depending on the configuration.

Note: Never open a power supply module or attempt to replace any of its internal components.

If a power supply module fails, replace it as quickly as possible. If the customer has only two power supply modules and one of them fails, the storage system shuts down. If the customer has three power supplies and one of them fails, the storage system continues to operate, but the failed power supply module should be replaced as soon as possible to restore high-availability operation. Failure to replace the power supply module could cause the storage system to shut down if either the fan module or another power supply module were to fail.

2.5 Fan Module

Challenge RAID uses a single fan module, also called a fan pack, that attaches to the back of the storage system's chassis. The fan module contains six high-capacity, multiple-speed fans and a control / monitor board. Figure 2-9 shows the fan module.

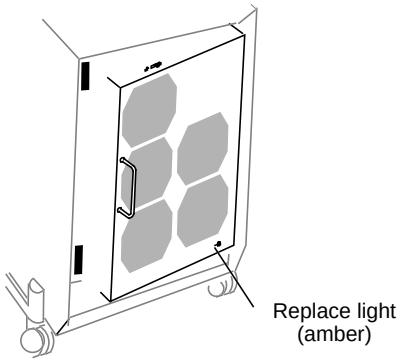

Figure 2-9 Partially Opened Fan Module on Back of Challenge RAID Storage System

The amber status light (replace light) comes on when a fan in the module is not working.

For high-performance operation, if a fan stops working, the remaining five fans speed up to maintain air flow, and the replace light turns on to indicate that you must replace the fan module as soon as possible. If one of the fans in the module fails, replace the fan module as quickly as possible to maintain high-performance operation. Failure to replace the module could cause the Challenge RAID storage system to shut down if either a power supply module or another fan fails.

You can remove a fan module while the storage system is running. The fan module is mounted on hinges so you can swing it away from the storage system. A mechanical lock and magnetic catches hold the fan module closed. If the fan module remains open for more than about two minutes, the storage system automatically shuts down to prevent overheating.

Note: Never open the module or attempt to replace any of its internal components unless you are specifically trained on this unit.

2.6 Optional Battery Backup Unit (BBU)

The battery backup unit is required for the Challenge RAID cache. In the event of a power outage, the battery backup unit allows the SP to shut down the storage system in an orderly way. The battery backup unit can provide up to 60 seconds of power, which gives the SP enough time to empty its cache by writing data from the cache to the physical disk units.

To conserve power, the SP shuts down all unnecessary disk modules while emptying the cache. Once cache data is safely stored on the physical disk units, the SP disables the battery backup unit to preserve the life of its internal power cells. When power returns, the battery backup unit recharges its cells automatically.

The battery backup unit is visible when you swing open the fan module on the back of the storage system. Figure 2-10 shows the battery backup unit.

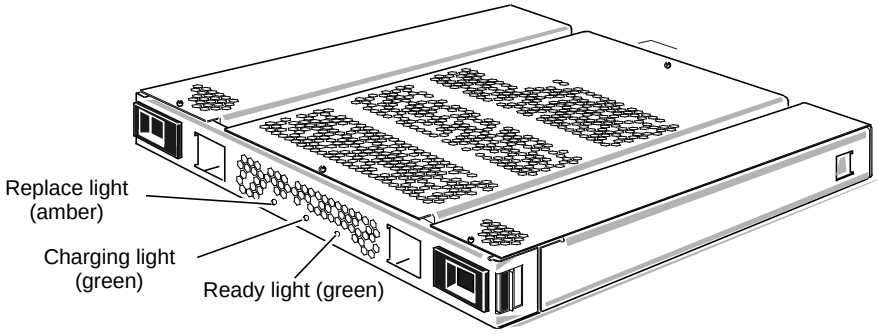

Figure 2-10 Battery Backup Unit (BBU) Removed From Challenge RAID Storage System

Three status lights indicate the following:

- ready light (green): lights when the battery backup unit is fully charged and ready to handle a power outage, and has been enabled by a SP; blinks when the battery backup unit is supplying power during an outage
- charging light (green): lights while the battery backup unit is recharging its internal power cells or when the battery backup unit has been disabled by a SP
- replace light (amber): lights when the battery backup unit's self diagnostics determines that the battery backup unit has failed, and momentarily when the battery backup unit powers on

You can install or remove a battery backup unit while the disk-array storage system is running.

Note: Never open a battery backup unit or attempt to replace any of its internal components.

2.7 Challenge RAID Configurations

A Challenge RAID storage system connects to one or two Challenge servers by a wide SCSI-2 differential bus to a SCSI-2 interface. It can use one or two SCSI buses. Four configurations are available:

- basic
- dual-interface / dual-processor
- split-bus
- dual-bus / dual-initiator

This section lists the basic components of each configuration and shows their interconnections.

For a storage system to support storage-system caching, it must have

- two SPs, each with 8 or 64 MB of memory
- a battery backup unit
- disk modules in compartments A0, B0, C0, D0, and E0

2.7.1 Basic Configuration

Components of the basic configuration are

- Challenge server with one SCSI-2 interface
- one SCSI-2 bus
- Challenge RAID storage system with one storage-control processor (SP A)

Figure 2-11 diagrams the basic configuration.

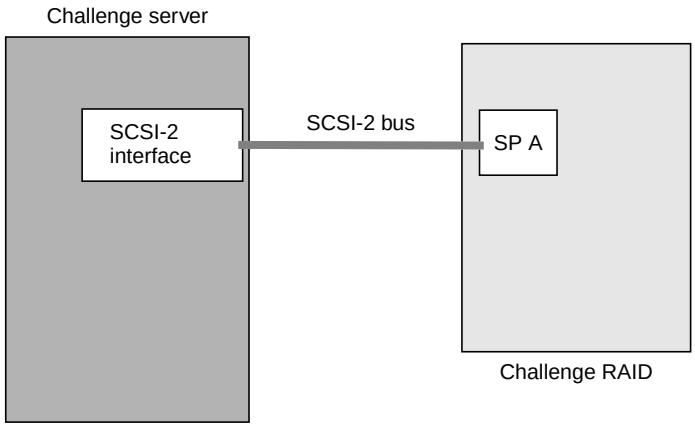

Figure 2-11 Basic Configuration

2.7.2 Dual-Interface/Dual-Processor Configuration

Components of the dual-interface / dual-processor configuration are

- Challenge server with two SCSI-2 interfaces
- two SCSI-2 buses
- Challenge RAID with two storage-control processors (SP A and SP B)

Figure 2-12 diagrams this configuration.

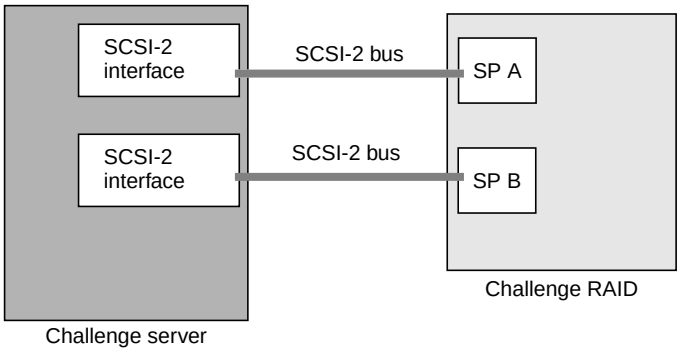

Figure 2-12 Dual-Interface / Dual-Processor Configuration

In this example, if one SP or SCSI-2 interface fails, stored data can be accessed through the alternate path. If the customer uses XLV volumes and applicable patches, this configuration provides simple alternate path switching (failover).

Note: After you replace the failed SP or SCSI-2 interface, you must reassign LUN ownership, as explained in "Reassigning LUN Ownership" in Appendix C.

2.7.3 Split-Bus Configuration

Components of the split-bus configuration are

- two Challenge servers, each with one SCSI-2 interface
- two SCSI-2 buses (one per Challenge server)
- Challenge RAID with two SPs

Figure 2-13 diagrams this configuration.

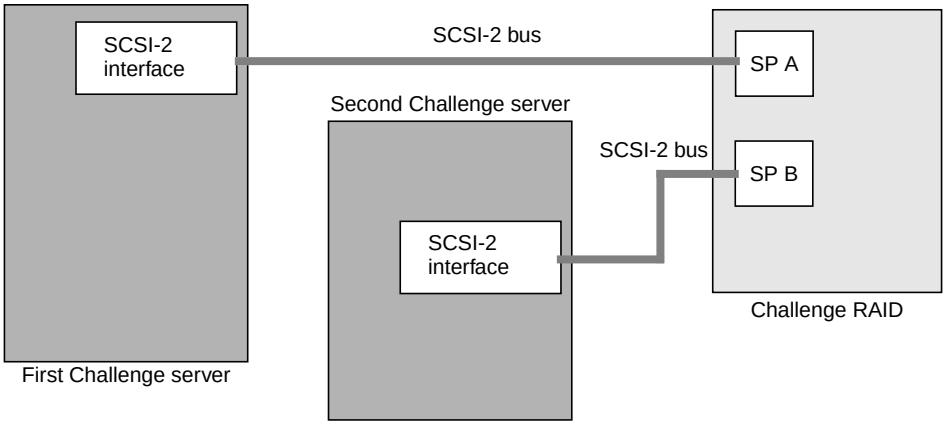

Figure 2-13 Split-Bus Configuration

In this configuration, if SP A fails, SP B takes over ownership of SP A's LUNs if auto-assign is enabled (which it is by default). After you replace SP A, you must use the *trespass* command to transfer ownership of SP A's LUNs back to the new SP A. "Reassigning LUN Ownership" in Appendix C in this guide gives instructions.

2.7.4 Dual-Bus/Dual-Initiator Configuration

Components of the dual-bus / dual-initiator configuration are

- two Challenge servers, each with two SCSI-2 interfaces
- four SCSI-2 buses (two per Challenge server)
- Challenge RAID with two SPs

The dual-bus / dual-initiator configuration provides the highest availability. Each host has two SCSI-2 adapters, each of which connects by a separate SCSI-2 bus to a separate SP in the storage system. Since this configuration protects against a SCSI-bus cable failure, it provides higher availability than the dual-initiator configuration. It is for enterprises requiring the highest level of availability, such as the Oracle Parallel Server™ and IRIS FailSafe™ products. Note that specific software is included in these products that protects the integrity of the data.

For better performance with this configuration, some physical disk units are bound on one SP and the other physical disk units on the other SP. The SP that binds a physical disk unit is the default owner of that physical disk unit. The route through the SP that owns a physical disk unit is the primary route to the physical disk unit. The route through the other SP is the secondary route to the physical disk unit, and is available if a component in the primary route fails. Figure 2-14 diagrams this configuration.

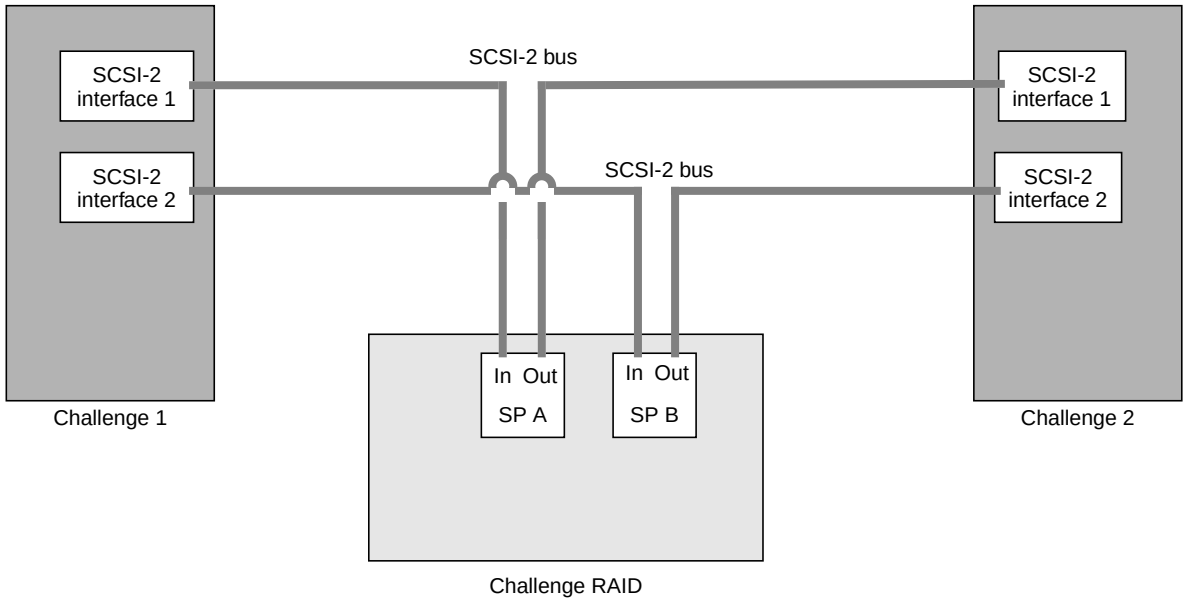

Figure 2-14 Dual-Bus / Dual-Initiator Configuration Example

Caution: Because both hosts can access all disks in this configuration, specific software is required to protect the integrity of the data, which is not included in the Challenge RAID storage-system software. The OPS and IRIS FailSafe products include such software.

Note: To work reliably, this configuration requires SCIP mezzanine boards.

Chapter 3

Installing a Challenge RAID Storage System

This chapter explains:

- setting up a Challenge RAID rack
- opening the rear door of a Challenge RAID rack
- reseating subassemblies
- labeling disk modules
- cabling RAID configurations
- setting SCSI IDs for SPs
- cabling an additional Challenge RAID storage system on a SCSI-2 bus channel
- setting SCSI IDs for SPs in a second Challenge RAID on the same SCSI-2 bus channel
- powering on the Challenge RAID storage system

A Phillips-head and a medium or thin flat-blade screwdriver are required. For maintenance and troubleshooting, a laptop (or ASCII terminal) is required. For unbolting the Challenge RAID rack, you need a 9/16-inch nut driver; the ratcheting type is recommended.

Note: For instructions on installing an additional chassis assembly in a Challenge RAID rack, see Chapter 5, “Maintaining and Upgrading Disk Modules.”

3.1 Setting Up a Challenge RAID Rack

Warning: **The fully loaded Challenge RAID rack weighs about 900 lb (407.96 kg); the weight is distributed within a relatively small footprint. Some sites have floor loadings of 50 lbs/ft^2 (244 kg/m^2) or less. Check that the flooring, especially raised flooring, can properly support the distribution of this weight. Particular care must be given to installations with more than one Challenge RAID rack placed close together.**

At least two persons are required to move a Challenge RAID rack. Do not undertake this procedure unless another qualified Silicon Graphics SSE is available to help.

To set up a Challenge RAID rack, follow these steps:

1. Remove the cardboard cover and verify that all parts are present and in good condition. If they are not, notify the shipper of the missing parts.

   Notice the ramp attached to the pallet by a piece of flexible plastic. You'll use this ramp later to move the rack off the pallet.

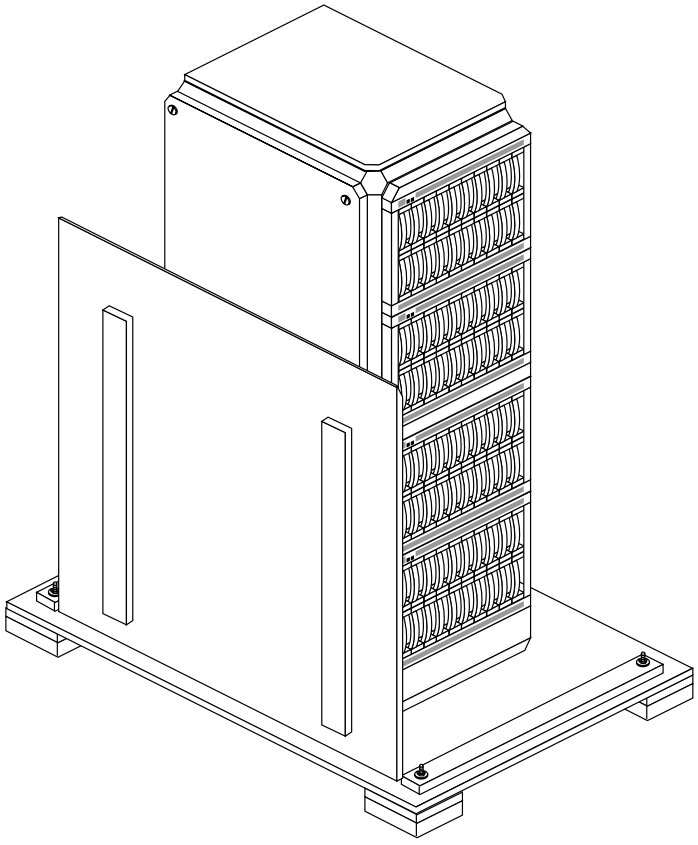

Figure 3-1 Ramp and Pallet (Shown With Bubblewrap Removed for Clarity)

2. Leave the rack still bolted to the pallet and move it to or near the position where the customer wants it.

   Note: The rack can be very difficult to move on rugs; it is much easier to move it while it is still bolted to the pallet. A fully loaded rack system weighs just under 900 pounds.

3. Remove the bubblewrap. Remove the flexible plastic that attaches the ramp to the pallet; move the ramp out of the way if necessary.

4. To remove the Challenge RAID rack side panels, which is necessary for unbolting the rack from the pallet, use a coin to open the latches that hold the side panels in place, as shown in Figure 3-2. Each panel has two latches, one in each upper corner.

   Remove both side panels. It is not necessary to remove the rear door for this procedure.

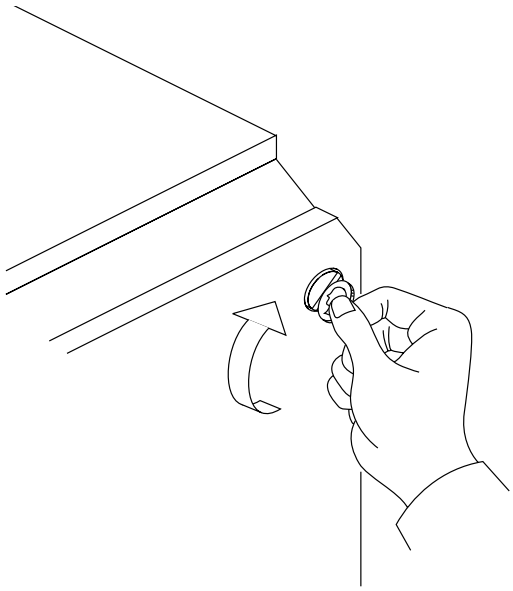

Figure 3-2 Unlocking a Challenge RAID Rack Side Panel

The rack is attached to the pallet with four brackets, two on each side; Figure 3-3 shows brackets on one side.

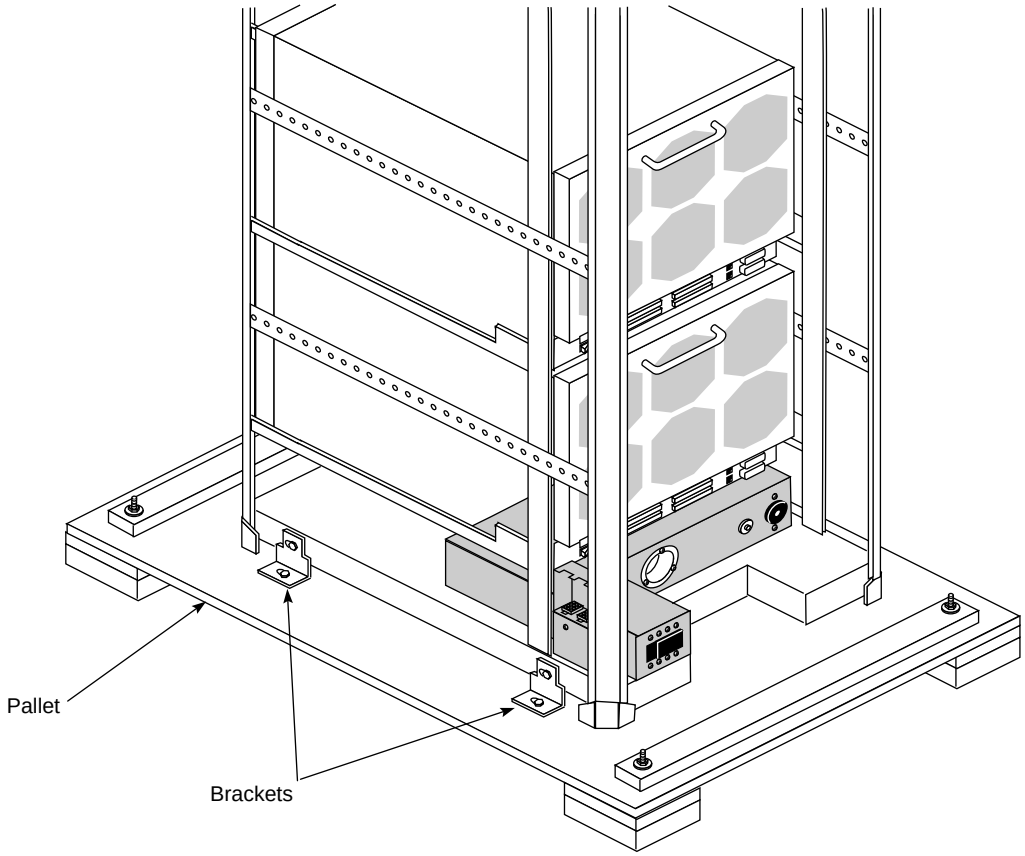

Figure 3-3 Challenge RAID Rack Pallet Bracket Locations

Note: Figure 3-3 shows the rack's rear door removed, but it is not necessary to remove it for this procedure.

5. Using a 9/16-inch nut driver, remove the nuts on all brackets; remove the brackets.
6. Position the ramp at the side of the pallet with one edge resting securely on the two feet of the pallet. Carefully move the rack off the pallet using the ramp.

Warning: **At least two persons are required to move a Challenge RAID rack. Do not undertake this procedure unless another qualified Silicon Graphics SSE is available to help.**

3.2 Opening the Rear Door of a Challenge RAID Rack

To check chassis assemblies in the rack, you must open the rear door.

To open the rear door of the rack, use a thin or medium flat-blade screwdriver, as shown in Figure 3-4.

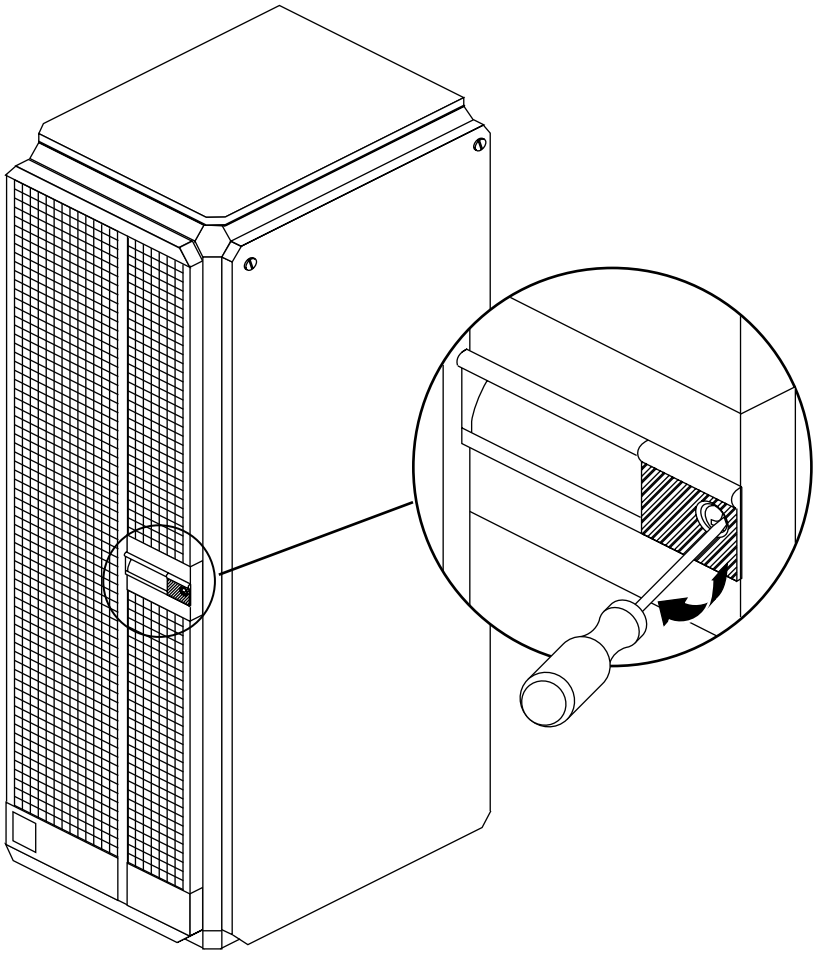

Figure 3-4 Opening the Challenge RAID Rack Rear Door

Move the screwdriver to the left to open the door.

Note: If you must remove the rear door entirely for any procedure, rehanging the rear door might require a second person because of the difficulty of lining up the door's three pinlocks.

3.3 Reseating Subassemblies

Make sure all components of the Challenge RAID storage system are firmly seated in their sockets. Follow these steps:

1. If the Challenge RAID storage system is turned on, power it off. See Figure 3-5.

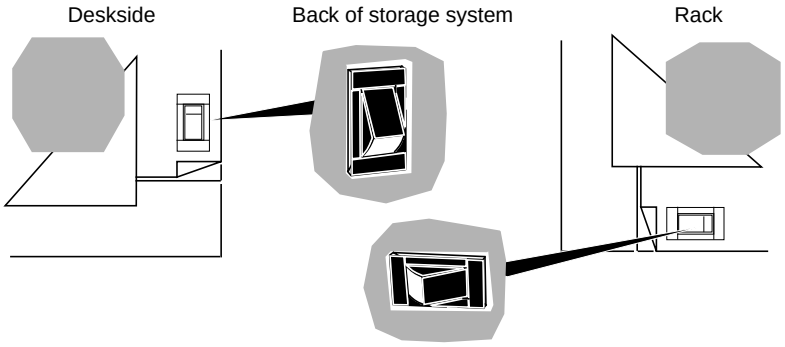

Figure 3-5 Challenge RAID Storage System Power Switch

2. On the back of the Challenge RAID chassis, move the fan module's latch to the **UNLOCK** position, as indicated in Figure 3-6.

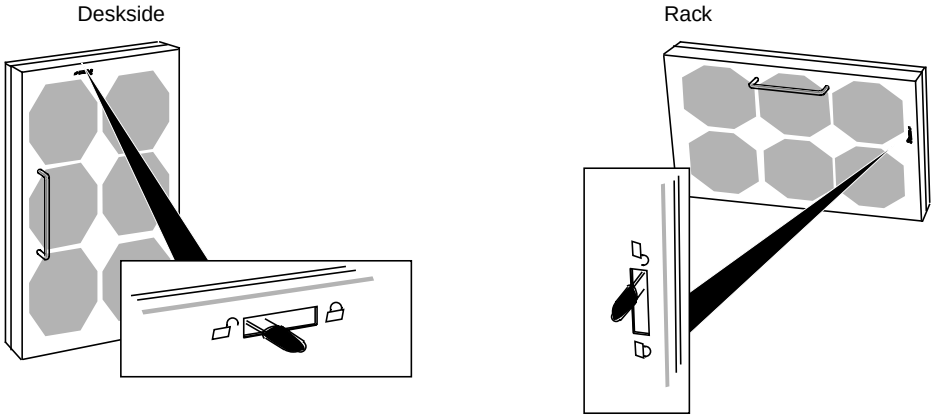

Figure 3-6 Unlocking the Fan Module

3. Swing open the fan module.
4. Check that these units are firmly seated:
   - power supply modules A, B, and (optionally) C
   - SP A; optional SP B
   - optional battery backup unit

   Figure 2-4 and Figure 2-5 in Chapter 2 diagram the location of these components for a deskside and rack system, respectively.

   Disk modules are discussed in the next section in this chapter.
5. Close the fan module by moving its latch to the **LOCK** position.

3.4 Labeling Disk Modules

Each disk module has a sticky label so that the installer or owner can label the disk module according to its the slot it occupies in the chassis. On the disk module label is a matrix that corresponds to the possible disk positions in the chassis. Figure 3-7 shows the disk module label.

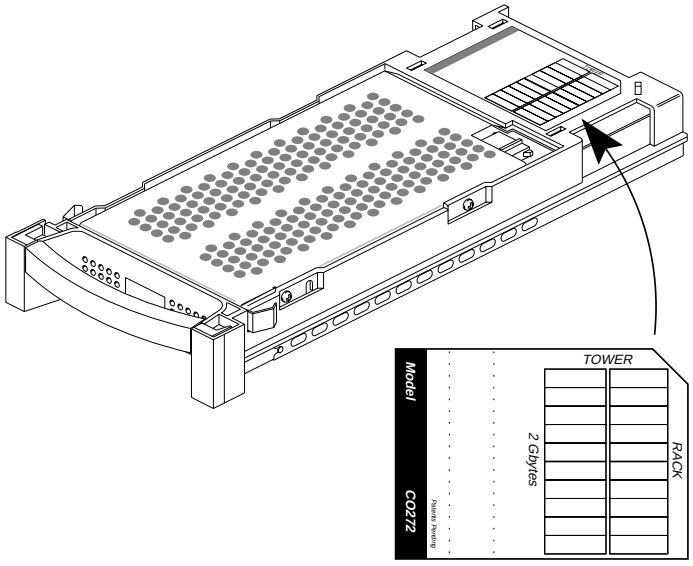

Figure 3-7 Disk Module Label

Labeling disk modules consists of marking this matrix. You can either write the slot position on the label in the corresponding place on the matrix or make a check mark in the position to indicate the slot that the disk module occupies. Figure 3-8 shows these two ways of labeling disk module B0.

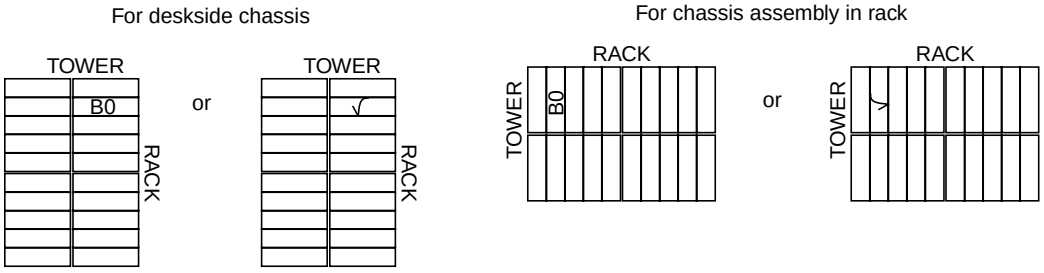

Figure 3-8 Marking the Label for Disk Module B0

For reference, Figure 3-9 diagrams all disk module locations.

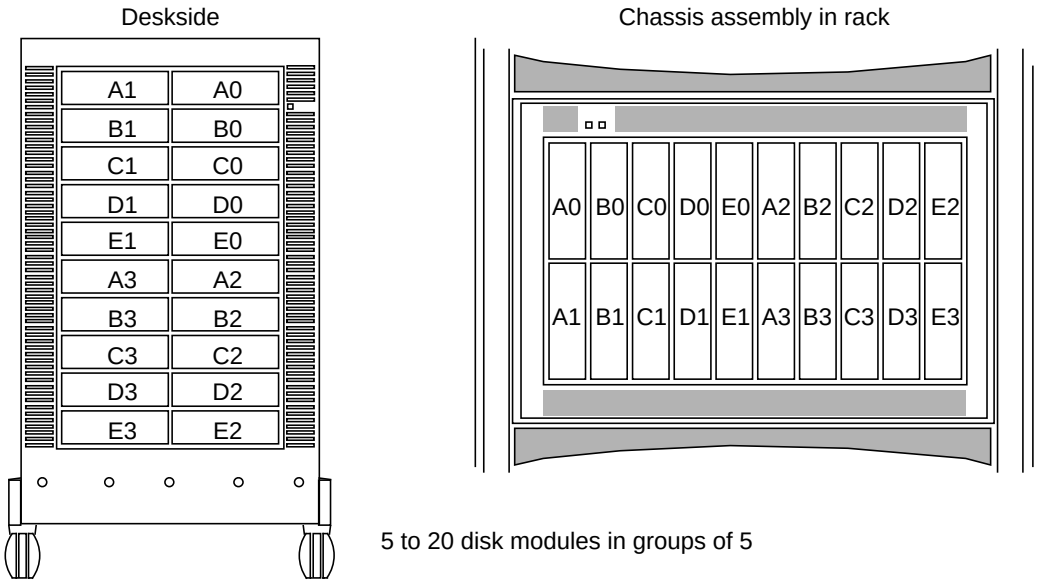

Figure 3-9 Disk Drive Locations

To label disk modules, follow these steps:

1. Pull out the disk drive module that is in slot A0. Mark the label **A0** in the space indicated in Figure 3-7. Reinsert and reseat disk drive module A0.
2. Pull out the disk drive module directly beneath it, that is, the drive module on the right side that is second from the top. In the corresponding space on the matrix, mark its label **B0**, as shown in Figure 3-8.
3. Reinsert and reseat disk drive module B0. Pull out the disk drive module directly beneath it, that is, the drive module on the right side that is third from the top. In the corresponding space on the matrix, mark its label **C0**.
4. Pull out, label, and reseat disk drive module D0; pull out, label, and reseat disk drive module E0.

   Caution: Disk modules A0, B0, C0, and, if present, A3 contain special system firmware for Challenge RAID storage systems. Do not remove them or disk modules from bus 0 for use in other disk module positions.

5. Pull out the disk drive module that is in the upper left slot. In the upper left space on the matrix, mark its label **A1**; reinsert and reseat it. Repeat this procedure for disk modules in slots B1, C1, D1, and E1.
6. If there are disk modules in the lower half of the chassis, pull out the disk drive module that is in the sixth slot from the top, on the right. In the corresponding space on the matrix, mark its label **A2**; reinsert and reseat it. Repeat this procedure for disk modules in slots B2, C2, D2, and E2.
7. If there are disk modules in the remaining bus in the chassis, pull out the disk drive module that is in the sixth slot from the top, on the left. In the corresponding space on the matrix, mark its label **A3**; reinsert and reseat it. Repeat this procedure for disk modules in slots B3, C3, D3, and E3.

3.5 Cabling Challenge RAID Configurations

Table 3-1 summarizes Challenge RAID configurations.

Table 3-1 Challenge RAID Configurations

Configuration	Host	SCSI-2 Interface	SCSI-2 Bus	SP
Basic	1	1	1	1
Dual-interface / dual-processor	1	2	2 (1 per server)	2
Split-bus	2	2 (1 per server)	2 (1 per server)	2
Dual-bus / dual-initiator	2	4 (2 per server)	2 (1 per server)	2

This section explains cabling each of these configurations.

3.5.1 Cabling the Basic Configuration

The Challenge RAID basic configuration consists of one Challenge server with one SCSI-2 interface, one SCSI-2 bus, and one Challenge RAID storage system or chassis assembly with one storage-control processor (SP A). Figure 3-10 diagrams the basic configuration.

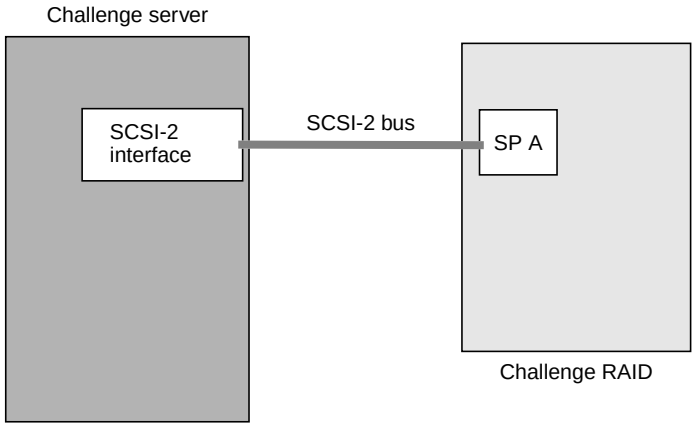

Figure 3-10 Basic Configuration

Figure 3-11 shows the location of SCSI-2 bus connectors on the Challenge RAID chassis.

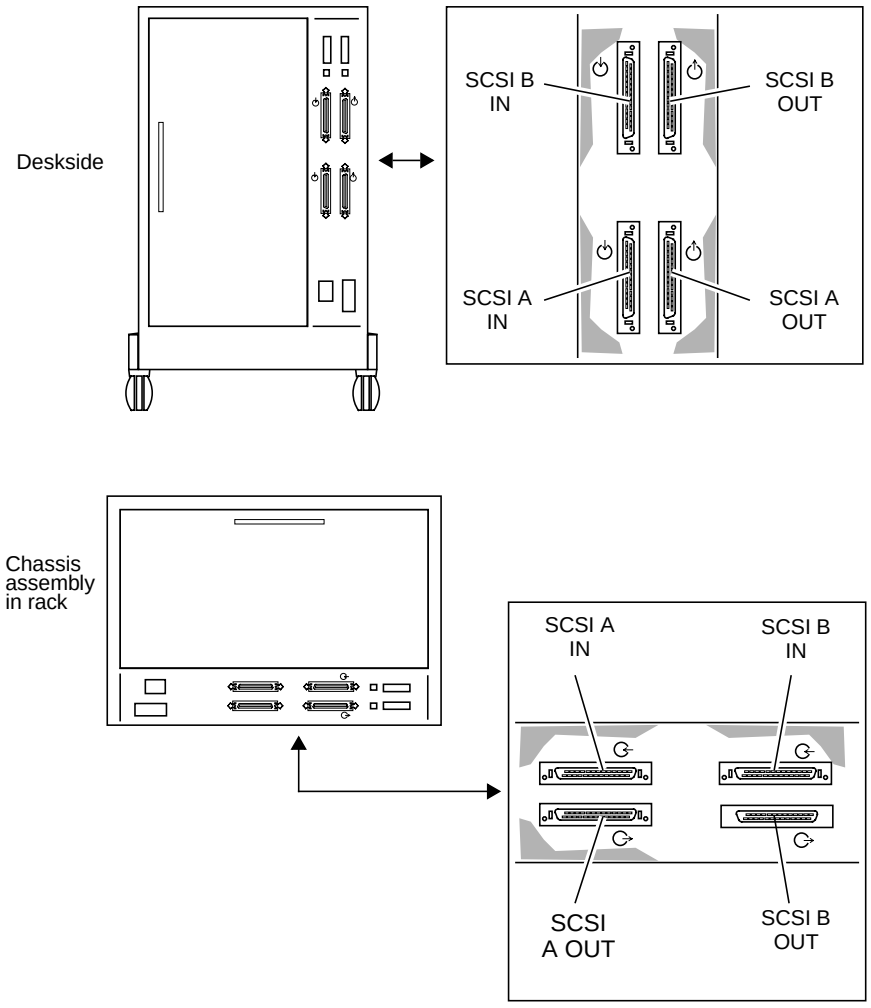

Figure 3-11 SCSI-2 Bus Connectors on Back of Challenge RAID Chassis

To cable the basic configuration, follow these steps:

1. Power off the Challenge server, if it is not already powered off.
2. Connect the SCSI cable to the **SCSI A** in connector as shown in Figure 3-12 and Figure 3-13, making sure the cable is fastened securely.

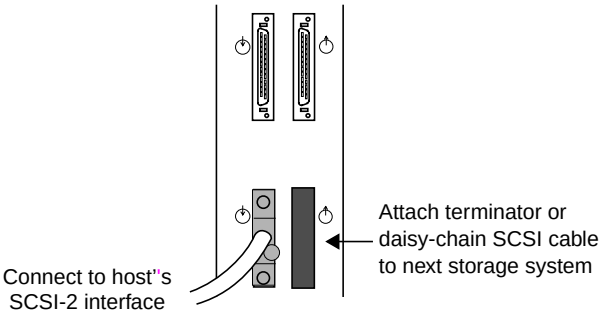

Figure 3-12 Cabling the Basic Configuration

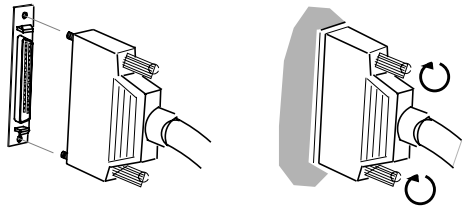

Figure 3-13 Connecting a SCSI Bus Cable

3. Connect a terminator plug to the **SCSI B** in connector as shown in Figure 3-14, making sure the terminator plug is fastened securely.

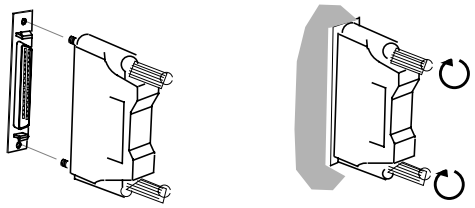

Figure 3-14 Connecting a SCSI Terminator Plug

Caution: SCSI terminators must be in place at all times during operation. Removing them causes the storage system to crash.

Figure 3-15 shows a typical hookup for a Challenge S.

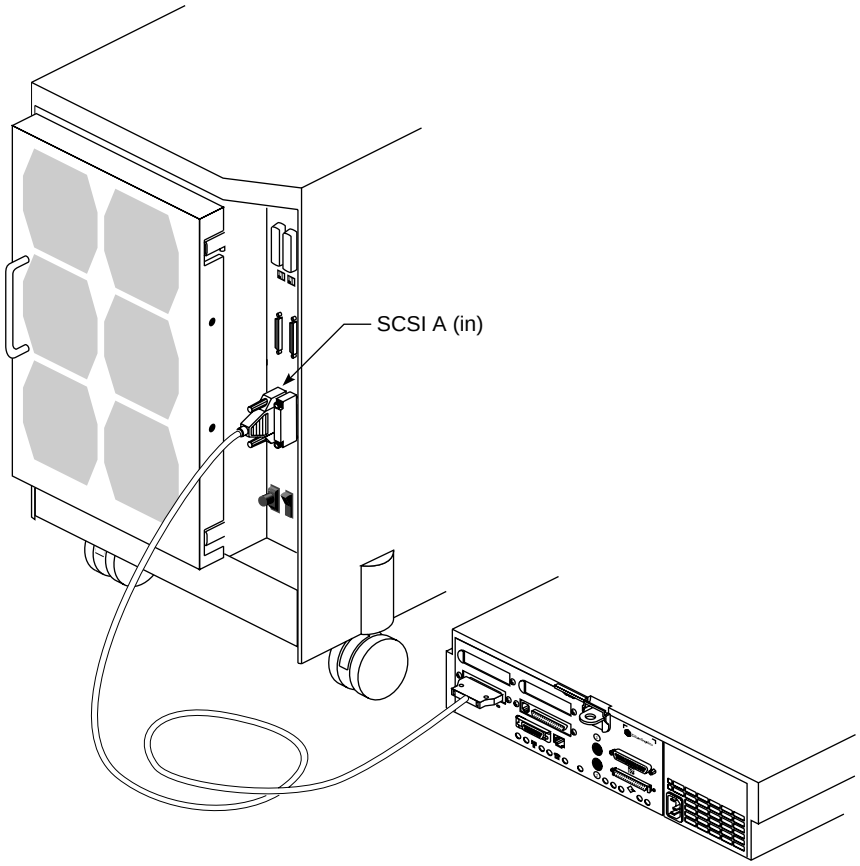

Figure 3-15 Connecting SCSI Bus Cables to a Challenge RAID Storage System and a Challenge S

3.5.2 Cabling the Dual-Interface/Dual-Processor Configuration

The Challenge RAID dual-interface / dual-processor configuration consists of a Challenge server with two SCSI-2 interfaces, two SCSI-2 buses, and a Challenge RAID storage system or chassis assembly with two storage-control processors (SP A and SP B). Figure 3-16 diagrams an example of this configuration.

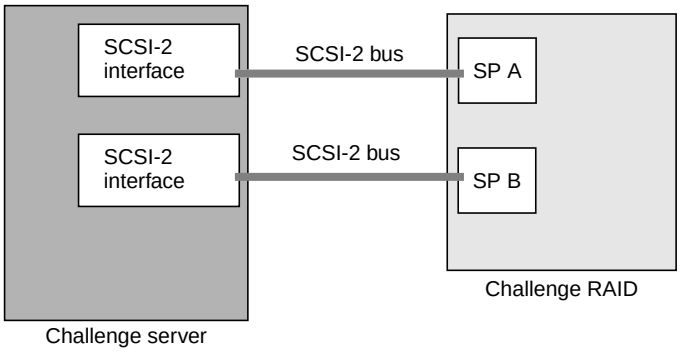

Figure 3-16 Dual-Interface / Dual-Processor Configuration Example

To cable the dual-interface / dual-processor configuration, follow these steps:

1. Power off the Challenge server (if necessary).
2. Arbitrarily name one interface in the Challenge server *first interface* and the other *second interface*.
3. Cable the first interface to the SCSI A in connector (the lower SP) in the Challenge RAID storage system. See Figure 3-11 for the location of the SP sockets.
4. Cable the second interface to the SCSI B in connector (the upper SP) in the Challenge RAID storage system. Connect the SCSI cable as shown in Figure 3-13, making sure all cables are fastened securely.

3.5.3 Cabling the Split-Bus Configuration

The Challenge RAID split-bus configuration consists of two Challenge servers, each with one SCSI-2 interface, two SCSI-2 buses (one per Challenge server), and a Challenge RAID storage system or chassis assembly with two SPs. Figure 3-17 diagrams an example of this configuration.

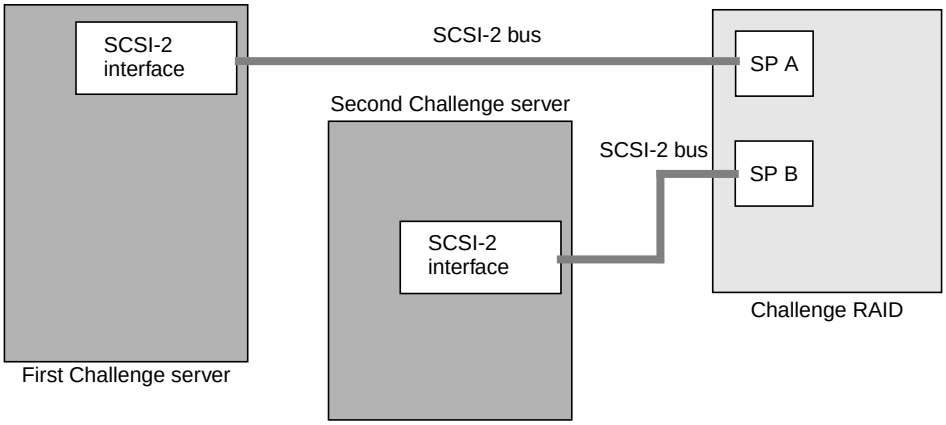

Figure 3-17 Split-Bus Configuration Example

To cable the split-bus configuration, follow these steps:

1. Power off the Challenge servers (if necessary).
2. Consulting the customer, designate one Challenge server *first server* and the other *second server.*
3. Cable the first server to the SCSI A in connector (the lower SP) in the Challenge RAID storage system. See Figure 3-11 for the location of the SP sockets.
4. Cable the second server to the SCSI B in connector (the upper SP) in the Challenge RAID storage system. Connect the SCSI cable as shown in Figure 3-13, making sure all cables are fastened securely.

Figure 3-18 diagrams connecting Challenge RAID to the SCSI bus in a split-bus configuration.

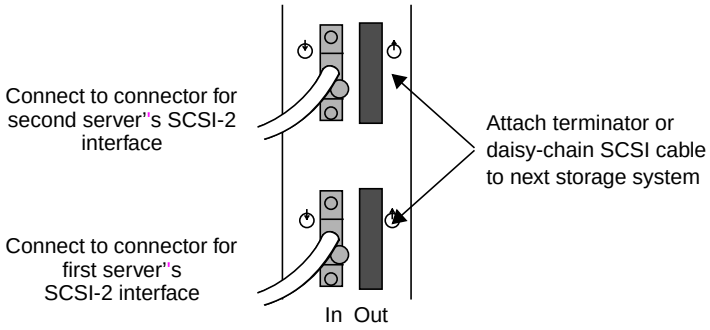

Figure 3-18 Connecting Challenge RAID: Split-Bus Configuration

3.5.4 Cabling the Dual-Bus/Dual-Initiator Configuration

The Challenge RAID dual-bus / dual-initiator configuration consists of

- two Challenge servers, each with two SCSI-2 interfaces
- four SCSI-2 buses (two per Challenge server)
- Challenge RAID storage system or chassis assembly with two SPs

Figure 3-19 diagrams an example of this configuration, which is required for the Oracle Parallel Server and IRIS FailSafe products.

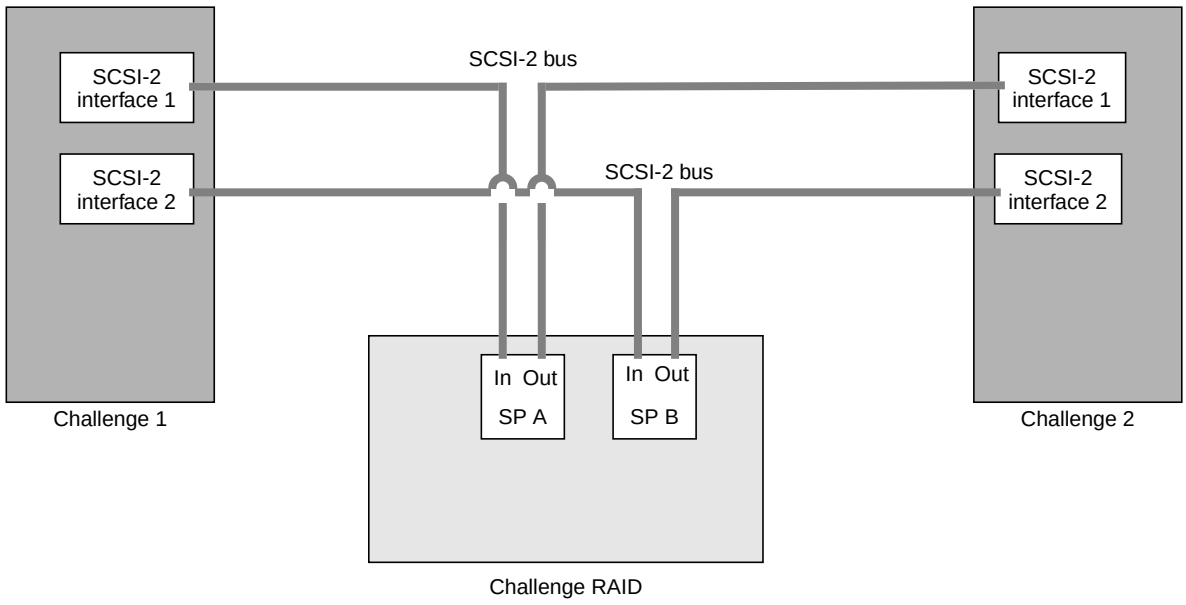

Figure 3-19 Dual-Bus / Dual-Initiator Configuration Example

To cable the dual-bus / dual-initiator configuration, follow these steps:

1. Power off the Challenge servers (if necessary).
2. Consulting the customer, designate one Challenge server as the *first server* and the other as the *second server.*
3. Cable the first SCSI interface of first server to the SCSI A in connector (the lower SP) in the Challenge RAID storage system. See Figure 3-11 for the location of the SP sockets.
4. Cable the first SCSI interface of the second server to the SCSI A out connector.
5. Cable the second server to the SCSI B in connector (the upper SP) in the Challenge RAID storage system.
6. Cable the second SCSI interface of the second server to the SCSI B out connector.

   Connect the SCSI cable as shown in Figure 3-13, making sure all cables are fastened securely. Figure 3-20 diagrams connecting Challenge RAID to the SCSI bus in a dual-bus / dual-initiator configuration.

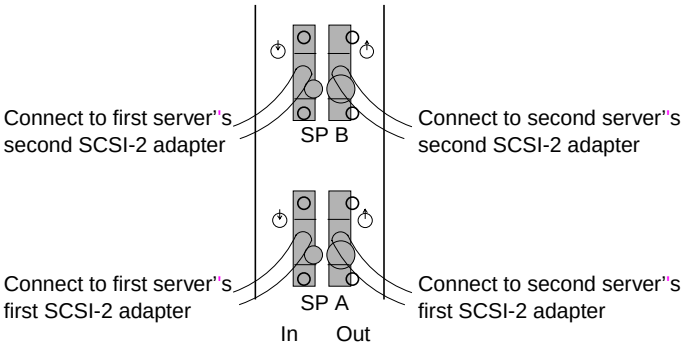

Figure 3-20 Connecting Challenge RAID: Dual-Bus / Dual-Initiator Configuration

Figure 3-21 shows a typical hookup for an OPS configuration with one Challenge RAID storage system.

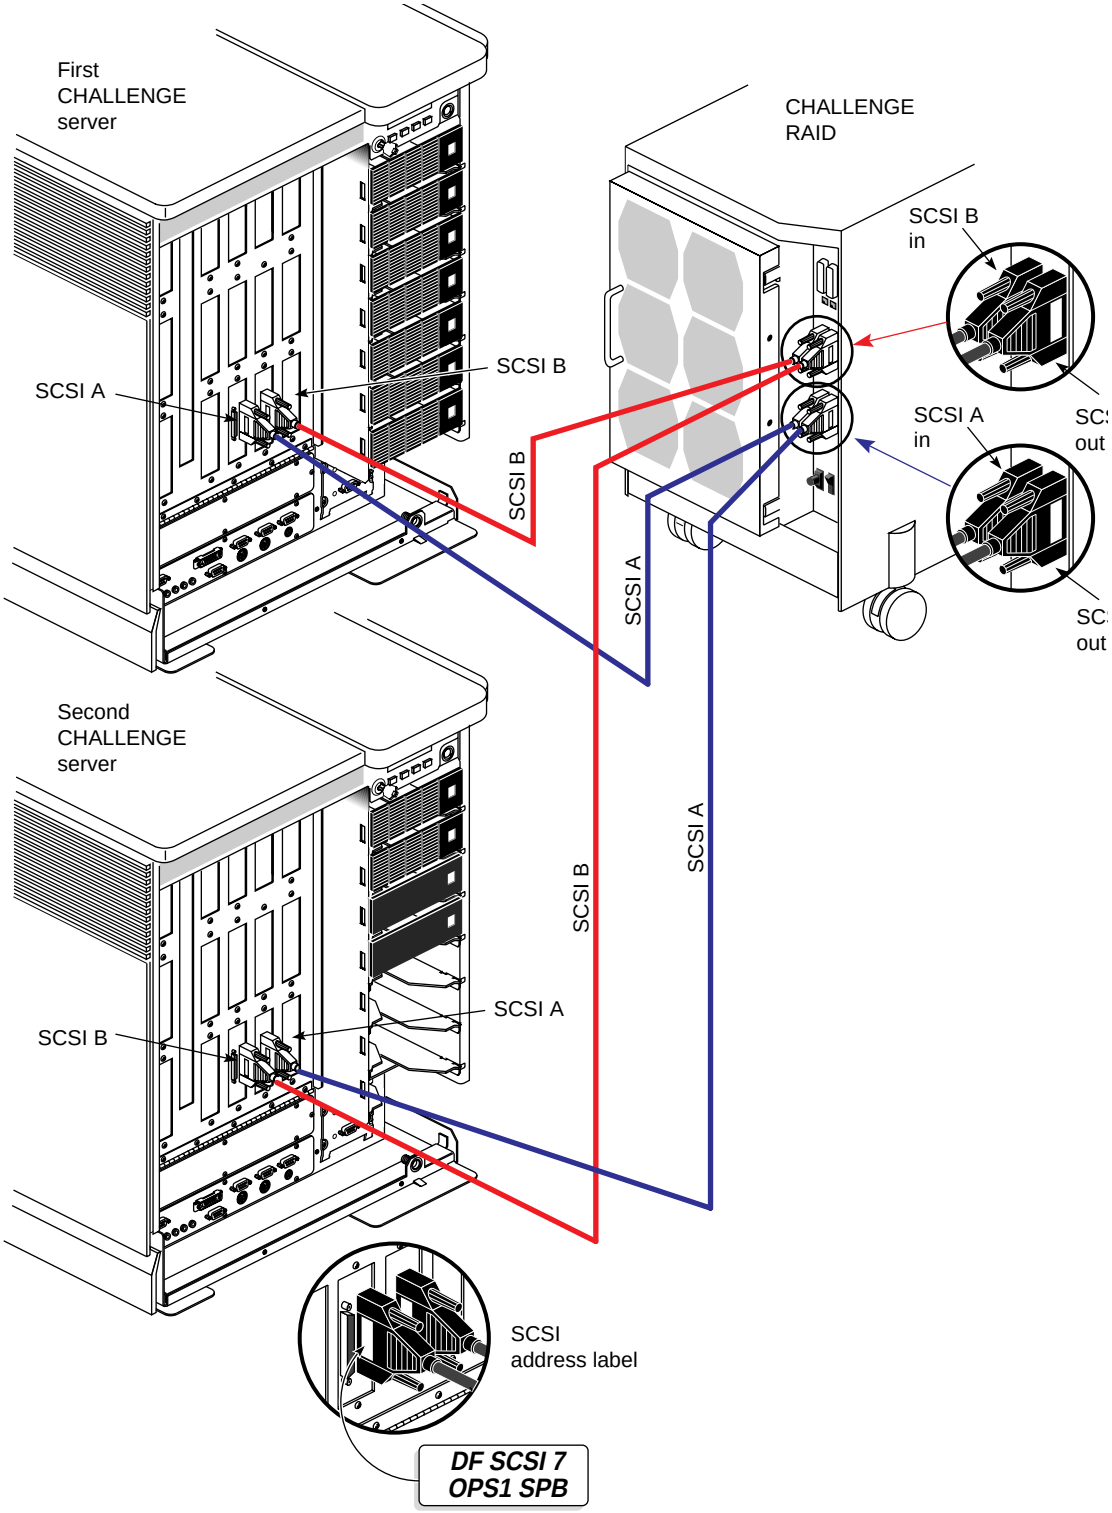

Figure 3-21 Connecting SCSI Cables to a Challenge RAID and Two Challenge L Servers

3.6 Setting SCSI IDs for Storage-Control Processors

Each SP connected to the same SCSI bus channel requires a unique SCSI ID number. The back of the Challenge RAID storage system has two SCSI ID switch packs, as shown in Figure 3-22. One switch pack sets the SCSI ID for SP A, and the other sets the SCSI ID for SP B. Each switch pack has four switches (ID0-ID3).

Note: To set SCSI IDs for a second Challenge RAID storage system on the same bus channel, see "Setting SCSI IDs for SPs in a Second Challenge RAID on the Same SCSI-2 Bus Channel," later in this chapter.

Set the SCSI ID switches for SP A and SP B to the lowest unused ID number. Figure 3-22 indicates the ID switches. In this example, SP A is set to 0 and SP B is set to 1.

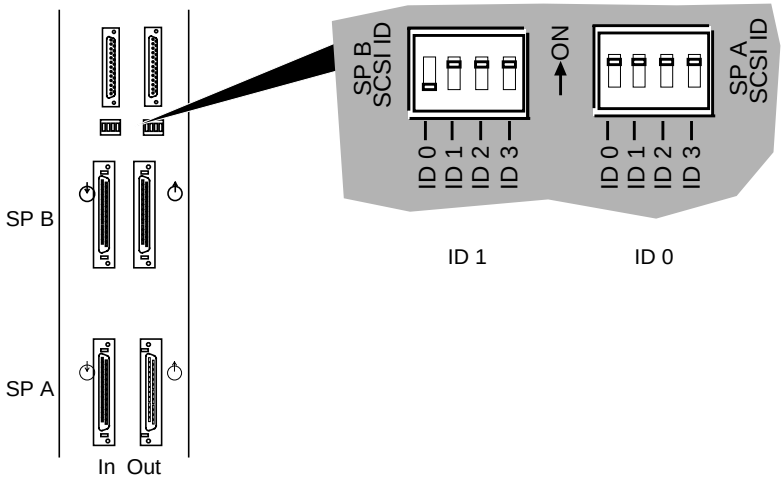

Figure 3-22 Setting an SP's SCSI ID (Back of Challenge RAID)

Table 3-2 charts ID switch settings.

Table 3-2 Setting SCSI IDs

SCSI ID Number	Switch Number			
	ID 0	ID 1	ID 2	ID 3
1	Off	On	On	On
2	On	Off	On	On
3	Off	Off	On	On
4	On	On	Off	On
5	Off	On	Off	On
6	On	Off	Off	On
7	Off	Off	Off	On
8	On	On	On	Off
9	Off	On	On	Off
10	On	Off	On	Off
11	Off	Off	On	Off
12	On	On	Off	Off
13	Off	On	Off	Off
14	On	Off	Off	Off
15	Off	Off	Off	Off

Note: Make sure that you set each SP to its own ID number—an ID number that is not used by any other device on the same SCSI bus. You can use any unused SCSI device number except 0.

3.7 Cabling an Additional Challenge RAID Storage System (or Chassis Assembly) on a SCSI-2 Bus Channel

For Challenge RAID rack storage systems, or for customer sites with more than one deskside Challenge RAID storage system, follow instructions in this section to cable more than one Challenge RAID storage system on the same SCSI-2 bus channel.

Once you have connected the first Challenge RAID to the host, connect a second one by following these steps:

1. Add the length of all SCSI cables to be used in the final configuration to make sure that it does not exceed the recommended maximum of 60 feet (18.3 meters). Include 3 feet for the Challenge chassis and 8 feet for each RAID chassis assembly. Use Figure 3-23 as a guide.

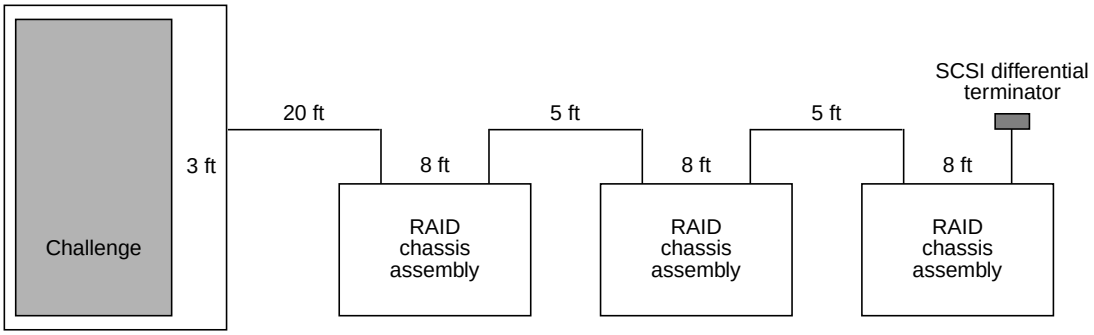

Total length in this example: 57 feet
Recommended maximum length: 60 feet

Figure 3-23 Computing SCSI Cable Length Example: Single-Host Configuration Only

Caution: Although the SCSI bus absolute length limit is 80 feet, exceeding 60 feet on the SCSI bus is not recommended. When cable lengths exceed 60 feet, problems can occur on the SCSI mezzanine card, the IO4B board, or both.

2. In Figure 3-24, locate the column of diagrams labeled “First Challenge RAID System.”
3. Circle the configuration for the Challenge RAID system that is already connected to the SCSI-2 bus channel.

   For example, if the first storage system configuration is the basic configuration, circle this configuration under “First Challenge RAID System,” as shown in Figure 3-25.

4. In the column of diagrams labeled “Second Challenge RAID System” in Figure 3-24, circle the configuration for the second Challenge RAID to which you want to connect the first storage system.

   For example, circle the dual-interface / dual-processor configuration under “Second Challenge RAID System,” as shown in Figure 3-25.

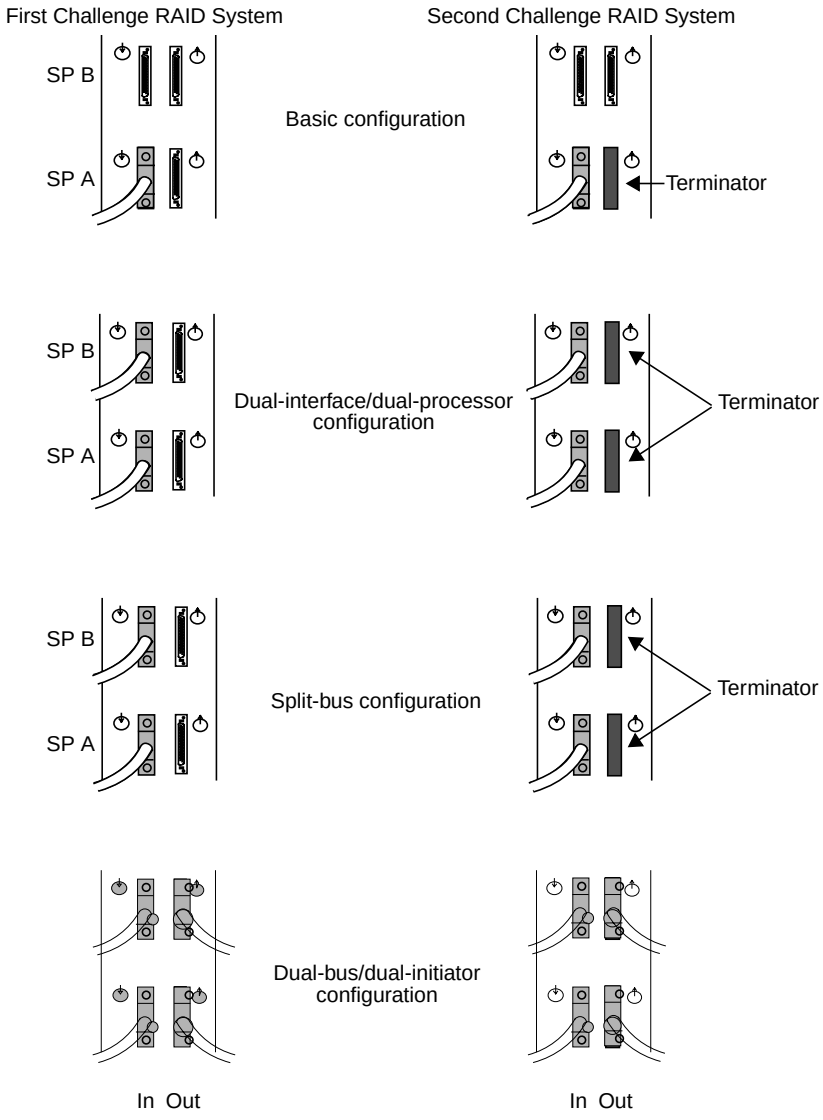

Figure 3-24 Two-Storage-System Cable-Planning Diagram

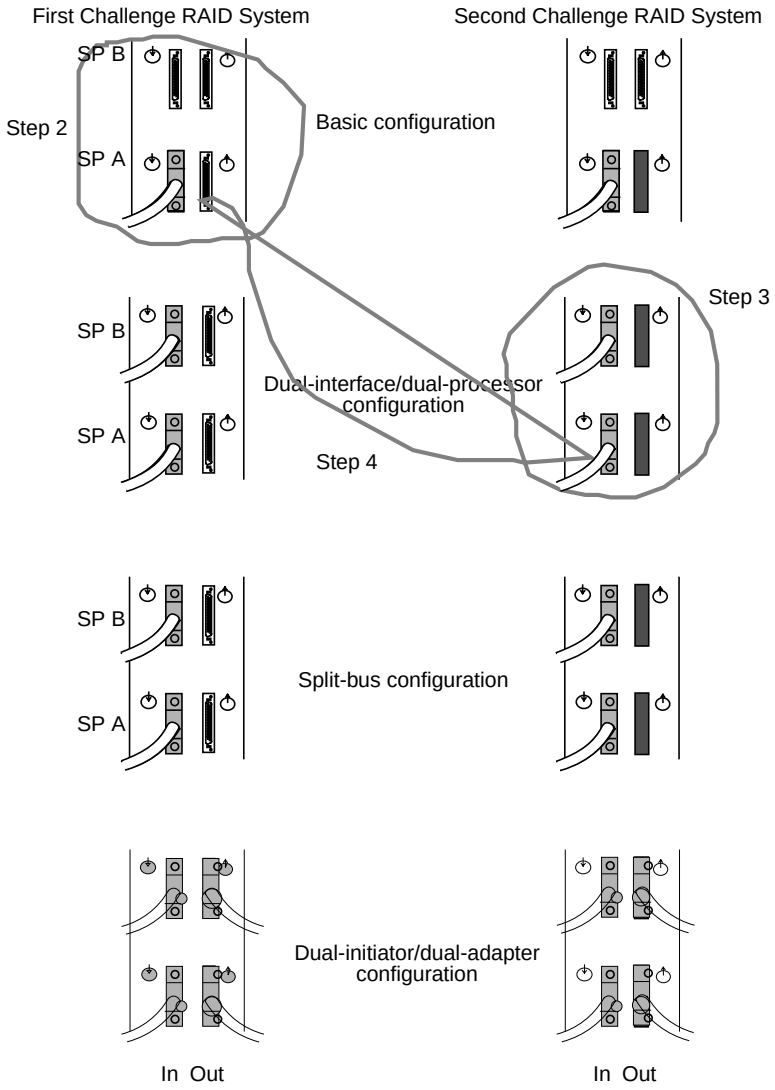

Figure 3-25 Example Two-Storage-System Cable-Planning Diagram: Second Challenge RAID

In your configuration, if the SCSI A channel of the first storage system is used (a SCSI cable is plugged into its in connector), but its out connector is not used (no SCSI cable or terminator plug is plugged into it), draw a line from the unused out connector of SCSI A to the end of an input cable of the second storage system. If the second storage system configuration has two available input cables, draw the line to either SCSI A or SCSI B. Which connector you use is entirely up to you.

For example, connect the SCSI A output of the first storage system to the SCSI A input of the second storage system, as shown in Figure 3-25.

In your configuration, if the SCSI B channel of the first storage system is used (a SCSI cable is plugged into its in connector), but its out connector is not used (no SCSI cable or terminator plug is on it), draw a line from the unused out connector of SCSI B to the end of an input cable of the second storage system. If no input cable is available on the second storage system, circle the SCSI B out connector of the first storage system to indicate that you must install a SCSI terminator plug in the connector.

For example, the SCSI B in connector of the first storage system has no input bus connected to it; therefore you cannot use it, as shown in Figure 3-25.

If the second storage system has an unconnected input cable, connect it directly to a SCSI bus connector on the server. Write “To server” at the end of this cable.

The basic configuration requires one SP installed in slot SP A. The others require two SPs. If you need to install or move an SP, follow the steps in Section 6.3, “Replacing or Adding an SP,” in Chapter 6.

Figure 3-26 shows SCSI buses A and B hooked up to two RAID chassis assemblies in a Challenge RAID rack storage system.

Caution: When planning to cable more than one Challenge RAID storage system or chassis assembly, take the differential SCSI cabling usable length limit of 60 feet (18.3 meters) into account. In the total you must include 8 feet for each Challenge RAID deskside unit or chassis assembly plus 3 feet for the server.

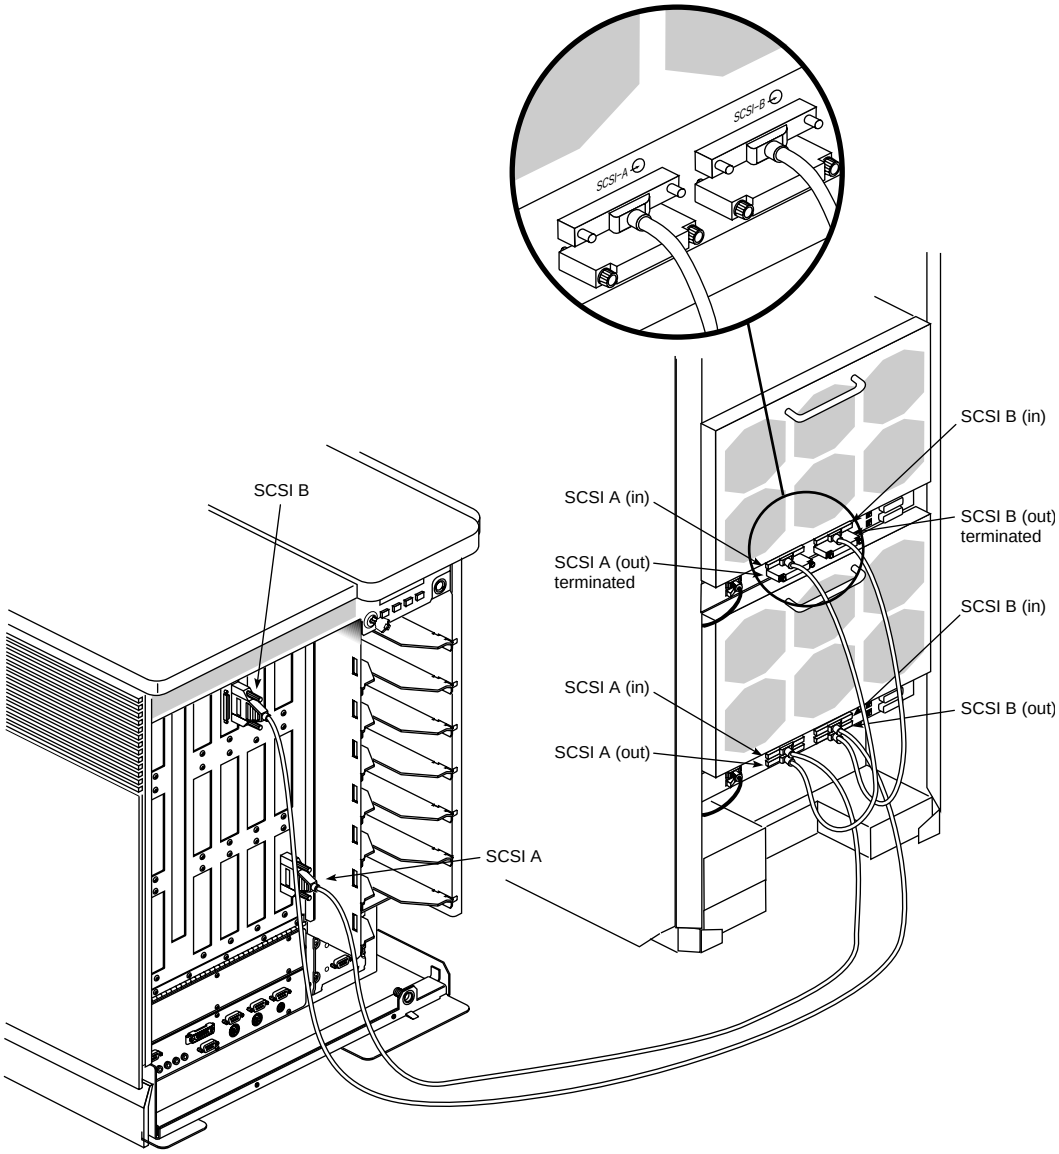

Figure 3-26 Connecting SCSI Buses A and B to Two RAID Chassis Assemblies in a Challenge RAID Rack Storage System

3.8 Setting SCSI IDs for SPs in a Second Challenge RAID on the Same SCSI-2 Bus Channel

Set the SCSI ID switches for the configuration you are installing; Table 3-3 presents suggested settings.

Table 3-3 Recommended SCSI ID Settings for SPs in Second Challenge RAID on a SCSI-2 Bus Channel

Dual-SP and Dual-Initiator Configuration for First Storage System: SPs on Same SCSI-2 Channel	Any Other Configuration for the First Storage System
SCSI ID 1 for the first SP A	SCSI ID 1 for the first SP A and SP B
SCSI ID 2 for the second SP B	SCSI ID 2 for the second SP
SCSI ID 3 for the third SP	SCSI ID 3 for the third SP
SCSI ID 4 for the fourth SP	SCSI ID 4 for the fourth SP

The recommended SCSI ID settings are just suggestions. You can set the SCSI ID for an SP to any number you want except for the following:

- Never set two SPs connected on the same bus to the same ID number. Each SP or device on a SCSI-2 bus channel must have a unique SCSI ID number.
- Use any unused SCSI device number except 0.

3.9 Powering On Challenge RAID

To power on the system, follow these steps:

1. Make sure that the Challenge RAID power switch is off.
2. Connect the power cord to the storage system. Plug the other end of the power cord into an AC power outlet. See Figure 3-27.

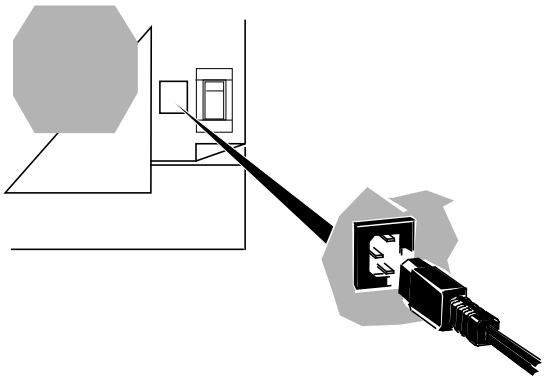

Figure 3-27 Connecting the Power Cord (Back of Challenge RAID Storage System)

3. Turn on the storage system's power.

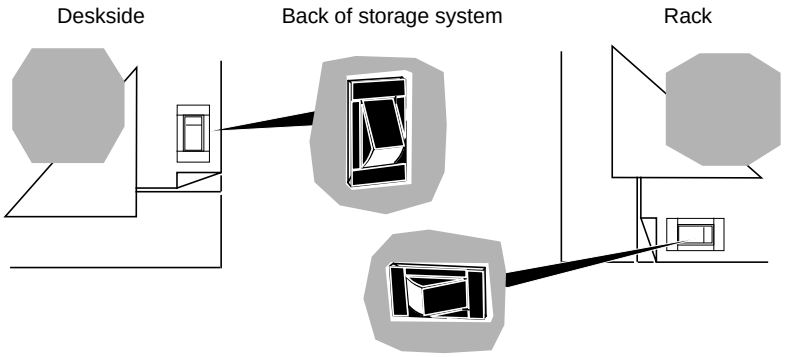

Figure 3-28 Turning On the Storage System Power

The green power light on the front of the Challenge RAID storage system turns on and the fans rotate. See Figure 3-29.

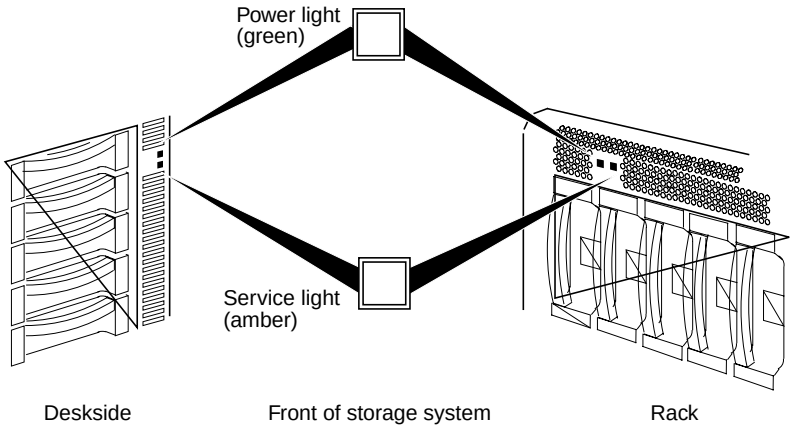

Figure 3-29 Challenge RAID Indicator Lights

4. If the busy light on *none* of the drive modules lights up, make sure that the power for each SP is enabled. Move the fan module's latch to the **UNLOCK** position, as indicated in Figure 3-30.

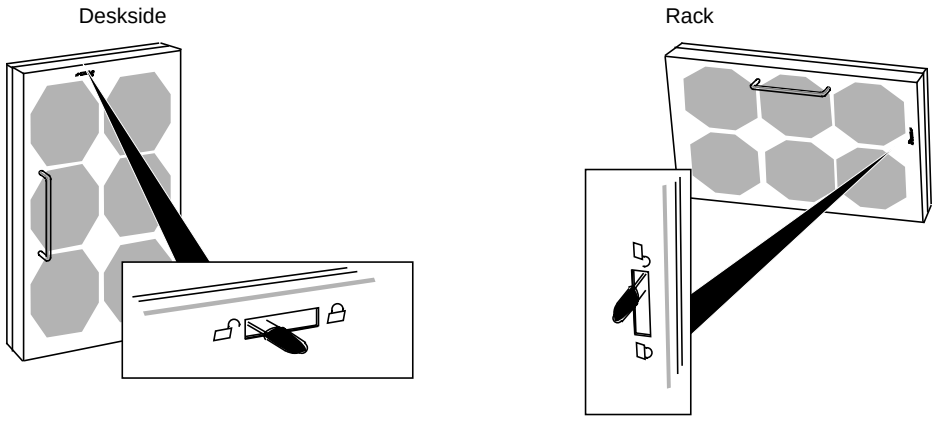

Figure 3-30 Unlocking the Fan Module

5. Swing open the fan module. On a Sauna SP, move the power switch to the enable position, as shown in Figure 3-31.

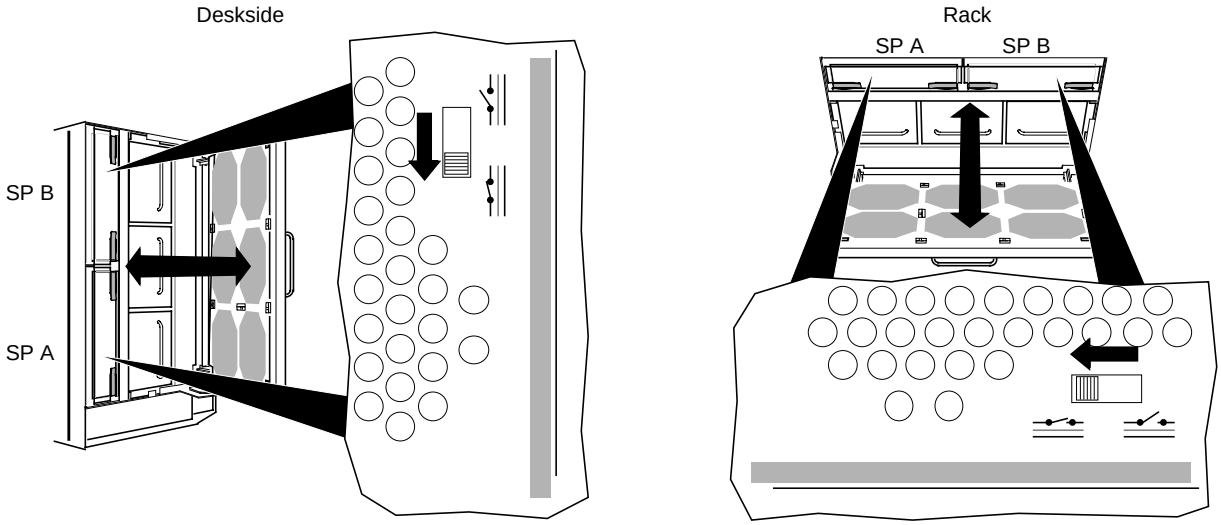

Figure 3-31 Enabling a Sauna SP's Power

Note: The Phoenix SP does not have a power switch.

6. Close the fan module by swinging it closed, and then moving the latch to the **LOCK** position.
7. Power on the Challenge server(s).

Chapter 4

Installing the Challenge RAID Software and Configuring Disks

This chapter explains how to label and partition physical disks, and create and mount disk filesystems. It explains how to use *raidcli*, the command-line interface (formerly *raid5*), in an IRIX shell to accomplish disk management tasks. This chapter explains

- installing the Challenge RAID software
- creating a configuration file
- starting the agent
- using *raidcli* commands
- getting the device name with *raidcli getagent*
- unbinding disks
- planning the bind
- binding disks into RAID units
- enabling command-tagged queuing
- understanding dual processors, load balancing, and device names
- setting up caching, if the customer desires (system with two SPs only)

A laptop or ASCII terminal is required for maintenance or troubleshooting.

4.1 Installing the Challenge RAID Software

Note the following:

- If the Challenge RAID storage system is not connected to Challenge server during the software install process, the device nodes will not be created.
- If a second (add-on) Challenge RAID is connected later, the device nodes must be created by changing to the */dev* directory and entering *./MAKEDEV dks scsi*.

Follow these steps:

1. Make sure the Challenge server is running IRIX version 5.3 or IRIX 5.3 with XFS and applicable patches. If it is not, upgrade the software using patches included in the current Required Patches List, including those on the patches CD-ROM (813-0412-001) included with the Challenge RAID product.
2. Load the Challenge RAID software from the RAID CD-ROM (813-0322-004). The Challenge RAID software images are created in *inst* format; thus, they require that the *inst* program be used. Installation software and instructions are included with the standard Silicon Graphics software products.
3. After quitting *inst*, reboot the system to start the software.

4.2 Creating a Configuration File

The Challenge RAID agent, which interprets commands from the command line interface for the Challenge RAID storage system, uses the configuration file */usr/raid5/raid5-agent.config*. This file contains

- description specification: list of objects that can be specified in the file
- RAID specification: format for RAID devices
- user specification: list of users allowed to perform write operations; root is the default
- poll specification: a value that determines whether polling is performed and whether the agent stores a copy of the log in its memory

This file is self-documenting and contains samples. Edit it to suit the customer's needs. */usr/raid5/raid5-agent.config.proto* is a example configuration.

4.3 Starting the Agent

RAID administration is through a daemon called the RAID agent. This agent communicates to the Challenge RAID storage system through */dev/scsi* device nodes and interprets commands from the command line interface for the SCSI device, namely, the storage system. When the storage system is installed, the startup script starts the agent.

Note: Because the agent opens the */dev/scsi* device nodes on an exclusive basis, no other process can open them.

The agent interprets commands from the command line interface for the SCSI device, namely, the Challenge RAID storage system. The agent can determine the existence and ownership of all LUNs in the RAID, regardless of ownership.

To start the agent manually, turn on the Challenge RAID storage system as explained in Chapter 3. Start the Challenge RAID software:

```
/etc/chkconfig -f raid5 on
/etc/init.d/raid5 start
```

If the agent is not active, the Challenge RAID storage system is not available to the server and an error message appears. For error conditions and solutions, see Appendix C, "Troubleshooting." Error messages from the agent are not reported to the console, but are stored in the *syslog* file.

4.4 Using *raidcli* Commands

The *raidcli* command sends Challenge RAID storage management and configuration requests to an application programming interface (API) on the Challenge server. For the *raidcli* command to function, the agent—an interpreter between the command line interface and the Challenge RAID storage system—must be running.

Note: A graphical user interface, RAIDGUI, is also available; it is documented extensively in the *CHALLENGE RAID Owner's Guide*, revision 007-2532-006 on.

The synopsis of the *raidcli* command is

`raidcli [-vp] [-d` *device*`]` *parameter* `[`*optional_arguments*`]`

In this syntax, variables and options mean:

-v — Enables verbose return.

-p — Parses the *raidcli* command without calling the API. If the string does not parse correctly, an error message is printed to stderr; otherwise there is no output.

-d *device* — Target RAID device. Use *raidcli getagent* for a list of RAID devices. This switch must be present for all *raidcli* management and configuration commands unless the environment variable indicates otherwise. This switch overrides an environment variable.

The environment variable RaidAgentDevice is the default value for the device when none is specified with the *-d* flag. If RaidAgentDevice is not set and no *-d* switch is present, an error is generated on all commands that need device information. For example:

```
setenv RaidAgentDevice sc4d2l0
setenv RaidAgentDevice /dev/scsi/sc4d2l0
```

The */dev/scsi* prefix is optional.

4.5 Getting the Device Name With *raidcli getagent*

Use the *getagent* parameter with the command line interface *raidcli* to display information on devices controlled by the API:

```
raidcli getagent
```

Following is sample output for one device; normally, the output would give information on all devices.

```
Agent Rev: 1.55
Name: phx-SPB-sc6d2
Desc: TOWER 1
Node: sc6d2l0
Signature: 0xc4aa1c00
Peer Signature: 0x0
Revision 9.2.91, Model: 7305
SCSI Id: 1
Prom Rev: 1.55.0
SP Memory: 32
Serial No: 96-0525-6
```

All *raidcli* commands except *raidcli getagent* require specifying the device; use the string at Node in the *getagent* output to specify the device. For a complete explanation of this output, see Appendix B, “The raidcli Command-Line Interface.”

4.6 Unbinding Disks

The disk modules in the Challenge RAID are configured as RAID-5 units at the factory. To bind disks at other RAID levels, you must unbind them first. Unbinding a LUN divides it into its constituent disk modules, allowing you to rebind the modules in any valid configuration the customer wants.

Unlike binding, unbinding takes only a few moments. You can unbind a LUN only if it is broken or bound, that is, only if its modules are in the Enabled or Broken state. You can unbind a hot spare only if it is on standby.

To unbind a disk, use the *unbind* parameter with the *raidcli* command. Follow these steps:

1. In an IRIX window, use *raidcli getagent* to get the device name (node number).
2. If necessary, use *raidcli getdisk* to verify the disk positions. See “getdisk” in Appendix B for details.
3. Use *raidcli unbind* to unbind the individual disk units:

   raidcli -d *device* **unbind** *lun-number* **[-o]**

In this syntax, *device* is the device name as returned by *getagent* and *lun-number* is the number of the logical unit to unbind. The **-o** flag specifies that the user is not prompted for permission. The *unbind* parameter has no output.

The following example destroys LUN 3 without prompting for permission:

```
raidcli -d sc4d2l0 unbind 3 -o
```

For more information on the *unbind* parameter, see “unbind” in Appendix B.

4.7 Planning the Bind

This section discusses

- filesystem capacity considerations for RAID
- determining RAID levels
- planning a RAID-3 bind

Note: The information in this section is also included in Chapter 4 of the *CHALLENGE RAID Owner’s Guide.*

4.7.1 Filesystem Capacity Considerations for RAID

You can bind as many as 16 drives in a LUN; for 4 GB drives, this capacity totals 64 GB. However, filesystems have capacity limits that must be taken into account.

- For IRIX 5.3 and earlier, an EFS filesystem is a maximum of 8 GB.
- For IRIX 5.3 and earlier, an XFS filesystem can be over a terabyte; however, the device parameter portion of the volume header is limited. Because values cannot exceed the number of bits allocated for each of the fields in the volume header (*sec/trk*, *trks/cyl*, *#cyls*), decrease the *#cyls* to within the 16-bit value available and increase one of the other fields.

To determine allowed capacity limits for an IRIX release, view the header files *dvh.h* and *dksc.h* in */usr/include/sys*.

4.7.2 Determining RAID Levels

Before you create LUNs, consult with the customer to plan the RAID levels. The LUNs are bound in a particular order, depending on the RAID level for each one:

- first: RAID-1_0
- next: RAID-0, RAID-3, or RAID-5
- next: RAID-1
- last: individual disk units and hot spares

Note the rules outlined in Table 4-1 for binding disk modules into LUNs.

Table 4-1 Binding LUNs

LUN to Bind	Restrictions and Recommendations
Any LUN	You can bind only unbound disk modules. All disk modules in a LUN must have the same capacity to use the modules' storage space fully. For higher availability with any redundant LUN, use slots on separate internal buses. For example, for a RAID-1 LUN, use slots A2 and B2 instead of A2 and A3.
RAID-0	You must bind a minimum of three disk modules, and no more than 16. If possible, use modules on different SCSI buses for highest availability. You can select the modules in any order.
RAID-1	You must bind exactly 2 disk modules, which you can select in any order.
RAID-1_0	You must bind an even number of modules: a minimum of 4 and a maximum of 16. RAID5GUI pairs modules into mirrored images in the order in which you select them. The first and second modules you select are a pair of mirrored images, the third and fourth are another pair of mirrored images, and so on. The first module you select in each pair is the primary image, and the second module is the secondary image. If possible, the modules you select for a pair should be on different SCSI buses for highest availability.
RAID-3 (firmware revision 9.0 and higher, SP model 7305, RAID agent 1.55 and higher)	You must bind exactly 5 disk modules, selected in any order. If possible, use modules on separate SCSI buses for highest availability (for example, A2, B2, C2, D2, E2). Before you bind, you must allocate RAID-3 memory for the SP, then allocate memory for the LUN. For example, if you allocate 12 MB for an SP, you can assign 6 MB to each of 2 LUNs owned by that SP, or 4 MB to each of 3 LUNs owned by that SP. In a failover system (dual-SP), the SPs must have RAID-3 memory allocated the same way, and each SP must have twice as much RAID-3 memory as each one actually uses in a normally functioning system. For example, if each SP has two LUNs with 3 MB each, the RAID-3 memory for each SP must be 12 MB, so that each SP is able to take over the function of the other SP if it fails. That is, each SP must have enough RAID-3 memory for its own LUNs (2×3 MB) plus the other SP's LUNS (2×3 MB).
RAID-5	You can select the modules for the LUN in any order. You must bind a minimum of 3 disk modules and no more than 16. For most efficient use of space, bind 5 modules. If possible, use modules on different SCSI buses for highest availability.
Individual disk	None
Hot spare	You cannot bind the following disk modules as hot spares: A0 through E0 or A3. The capacity of a disk module bound as a hot spare must be at least as great as the capacity of the largest disk module that it might replace.

4.7.3 Planning a RAID-3 Bind

RAID-3 for systems with SPs that have firmware revision level 9.0 and higher (SP model 7305) and RAID agent 1.55 and higher has enhanced performance compared to previous firmware revision levels. This version of RAID-3 ("fast" RAID-3) has specific memory characteristics and requirements, which are explained in this section. (To determine revision level, see "getsp" in Appendix B.)

4.7.3.1 Basic Memory Requirements for RAID-3

An SP's memory is divided into RAID-3 space, storage-system buffer space, write cache space, and read cache space. You allocate RAID-3 memory before binding LUNs; this chapter gives specific steps for each user interface.

Storage-system buffer space in a PowerPC-based SP is always 4 MB. You can allocate all or part of the remaining memory for RAID-3 and other uses. For example, if the SP in a system has 16 MB of memory, memory can be allocated for:

- system buffers (4 MB)
- read or write cache for non-RAID-3 LUNs, if any
- any previously bound RAID-3 LUNs, if any
- RAID-3 LUNs

For RAID-3 LUNs, you allocate RAID-3 memory instead of read or write cache. For example, if a system with one 16-MB SP system is to contain two RAID-3 LUNs only, you can allocate all 12 MB (after the 4 MB required for system buffers) to RAID-3 memory.

You then split this memory between the two LUNs, assigning 6 MB to each. Figure 4-1 diagrams this example.

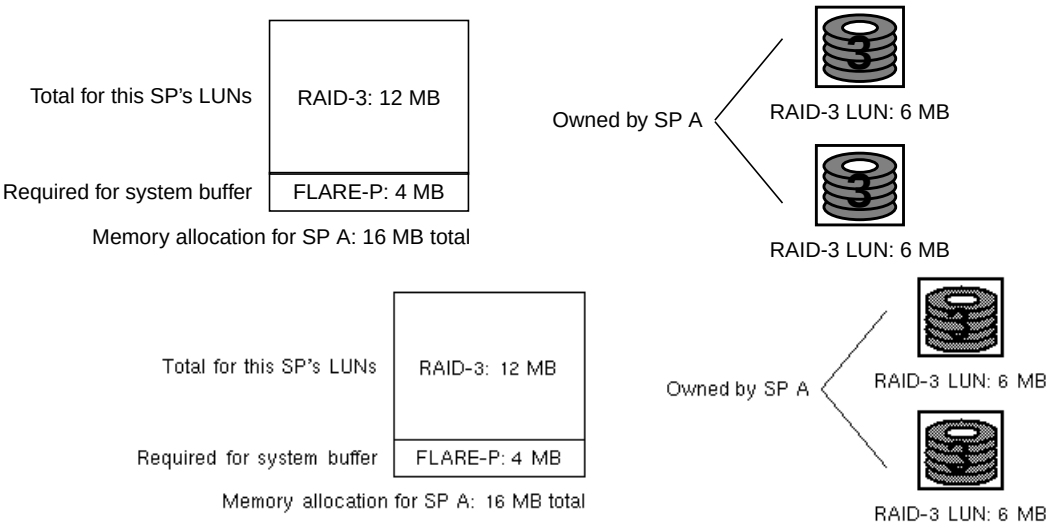

Figure 4-1 Example RAID-3 Memory Allocation: One-SP System With RAID-3 LUNs Only

In systems with two SPs, the amount of SP memory allocated for RAID-3 must be the same for each SP.

4.7.3.2 RAID-3 and Failover

A RAID-3 LUN uses the RAID-3 memory on the SP that owns it. Therefore, if failure occurs and ownership of the failed SP is transferred, the surviving SP must have enough memory for its own LUNs and that of the failed SP. Figure 4-2 diagrams an example.

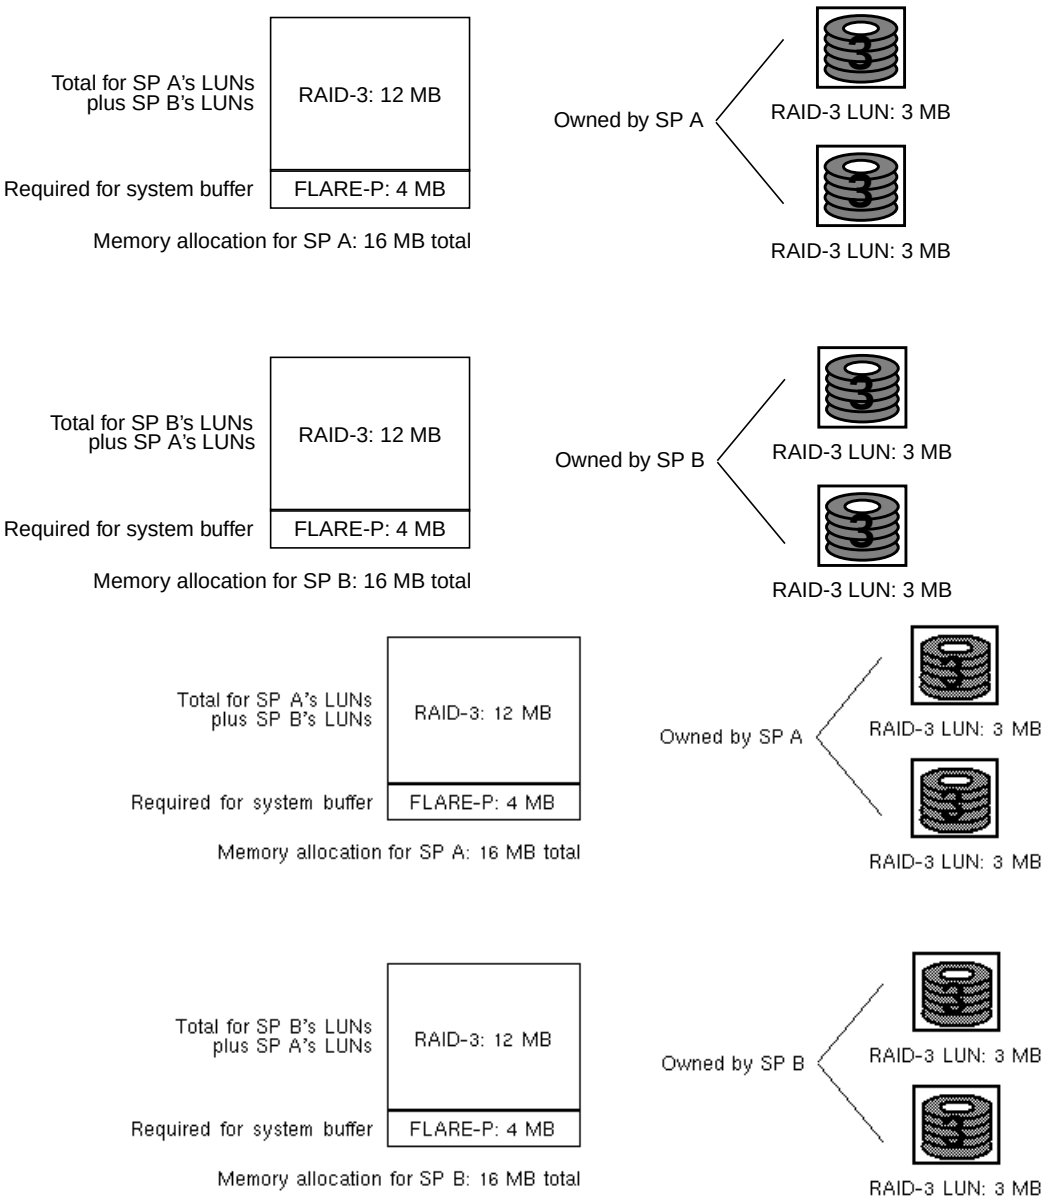

Figure 4-2 Example RAID-3 Memory Allocation for Failover System With RAID-3 LUNs Only

For failover systems, both SPs should have the same amount of physical memory. If the two SPs in a system have different amounts of physical memory, ownership of RAID-3 LUNs transfers from a failed SP to a working SP only if the working SP has

- 4 MB required for system buffers

- memory allocated for read or write cache for non-RAID-3 LUNs, if any
- total RAID-3 memory for its RAID-3 LUNs and the other SP's RAID-3 LUNs
- minimum of 4 MB for each RAID-3 LUN

Figure 4-3 diagrams an example.

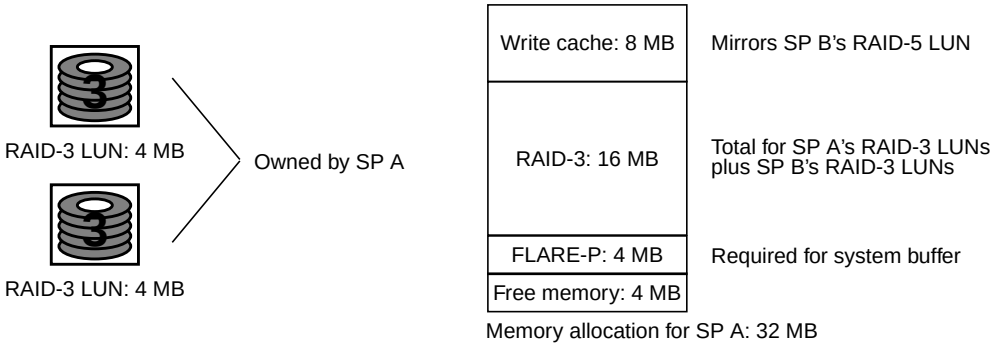
RAID-3 LUN: 4 MB
RAID-3 LUN: 4 MB
Owned by SP A
Write cache: 8 MB
Mirrors SP B's RAID-5 LUN
RAID-3: 16 MB
Total for SP A's RAID-3 LUNs
plus SP B's RAID-3 LUNs
FLARE-P: 4 MB
Required for system buffer
Free memory: 4 MB
Memory allocation for SP A: 32 MB

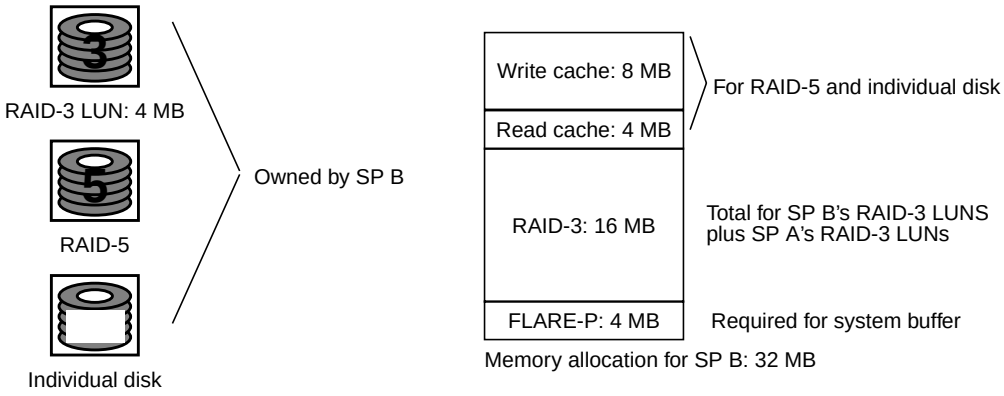
RAID-3 LUN: 4 MB
RAID-5
Individual disk
Owned by SP B
Write cache: 8 MB
For RAID-5 and individual disk
Read cache: 4 MB
RAID-3: 16 MB
Total for SP B's RAID-3 LUNS
plus SP A's RAID-3 LUNs
FLARE-P: 4 MB
Required for system buffer
Memory allocation for SP B: 32 MB

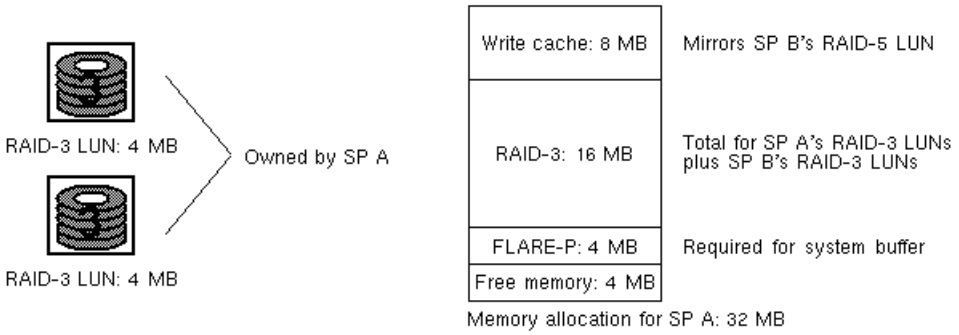
RAID-3 LUN: 4 MB
RAID-3 LUN: 4 MB
Owned by SP A
Write cache: 8 MB
Mirrors SP B's RAID-5 LUN
RAID-3: 16 MB
Total for SP A's RAID-3 LUNs
plus SP B's RAID-3 LUNs
FLARE-P: 4 MB
Required for system buffer
Free memory: 4 MB
Memory allocation for SP A: 32 MB

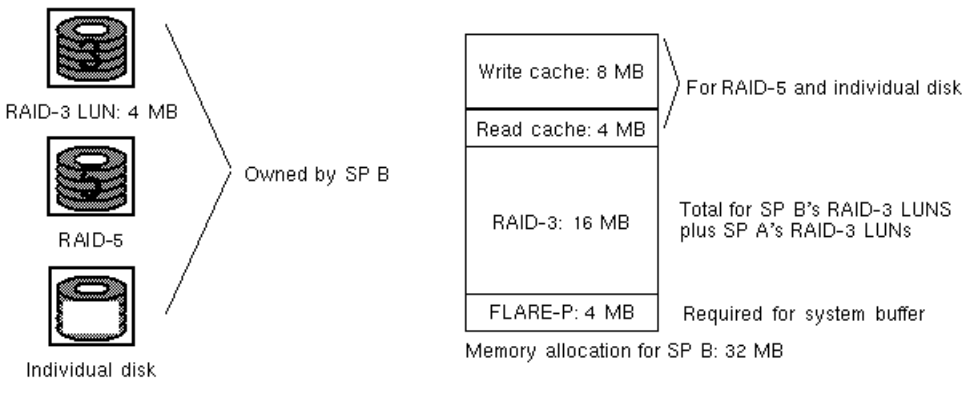
RAID-3 LUN: 4 MB
RAID-5
Individual disk
Owned by SP B
Write cache: 8 MB
For RAID-5 and individual disk
Read cache: 4 MB
RAID-3: 16 MB
Total for SP B's RAID-3 LUNS
plus SP A's RAID-3 LUNs
FLARE-P: 4 MB
Required for system buffer
Memory allocation for SP B: 32 MB

Figure 4-3 Example RAID-3 Memory Allocation for Failover System

4.7.3.3 Maximizing RAID-3 LUN Performance

To get the most out of a storage system with RAID-3, follow these guidelines:

- Bind all LUNs in the system to RAID-3 only. In a storage system with onlyRAID-3 LUNs, no memory is allocated for storage-system read or write caching; all memory is allocated for RAID-3.
- Allocate the maximum of 6 MB of storage-system memory for each RAID-3 LUN. For two RAID-3 LUNs, allocate 12 MB of storage-system memory for the RAID-3 memory partition and 6 MB to each of the two RAID-3 LUNs when you bind them. Since an SP requires 4 MB of storage-system memory for system buffers, each SP needs 16 MB of memory. Failover requires 32 MB**.**

4.8 Allocating Memory for RAID-3 LUNs

If the LUNs you are creating are RAID-3 for SP model 7305/firmware revision 9.0 and higher and RAID agent 1.55 and higher, you must allocate a special amount of memory dedicated to RAID-3 before starting the bind. In this revision, RAID-3 is "fast" RAID-3. Use *raidcli -d <device> getsp* to determine an SP's firmware revision level and model number.

Follow these steps:

Note: The system cannot be used during allocation of RAID-3 memory.

1. Determine the amount of memory you can use for each SP. If necessary, use *raidcli getcache* to determine each SP's total memory. Look for `Raid 3 memory size for spA/spB` in the output.

   Note: For RAID-3, you must allocate at least 2 MB of RAID-3 memory per LUN for each SP. Failover requires a minimum of 4 MB per LUN for each SP.

2. Quiesce the bus: log users off and back up the system. The system cannot be used during RAID-3 memory allocation.

3. Make sure read and write caches are disabled and there are no unassigned dirty pages. For example:

   ```
   raidcli -d sc4d2l0 getcache
   ```

   In the output, look for these lines:

   ```
   Cache State: Disabled
   ...
   Unassigned Cache Pages: 0
   ...
   Prct Dirty Cache Pages = 0
   ...
   SPA Read Cache State enabled
   SPB Read Cache State enabled
   ```

4. Use *raidcli setcache* to set the total amount of caching dedicated to RAID-3. (For full details of *setcache*, see "setcache" in Appendix B.)

   Note these guidelines:

   - Storage-system buffer space in a PowerPC-based SP is always 4 MB. You can allocate all or part of the remaining memory for RAID-3.
   - Allocating write/read cache memory decreases the amount of memory available for RAID-3 memory. If the storage system uses RAID-3 LUNs only, no write and read caches are used; do not allocate memory for this purpose.
   - Each SP must have at least 2 MB allocated for RAID-3. For maximum RAID-3 performance, allocate 6 MB of RAID-3 memory for each RAID-3 LUN you want to bind.
   - For failover, each SP must have the same amount of memory allocated. Furthermore, the total memory you set for each SP must include memory for this SP plus that of the other SP. See Section 4.7.3, "Planning a RAID-3 Bind," earlier in this chapter.

   Caution: Setting RAID-3 cache memory reboots the firmware. Make sure the system is not in use before proceeding.

The following example sets the RAID-3 cache to 6 MB:

```
raidcli -d sc4d2l0 setcache 1 -r3 6
```

For failover, RAID-3 memory assigned to each SP must be the total of the SPs used in the failover system. This example sets the RAID-3 cache for a system in which each SP has 6 MB:

```
raidcli -d sc4d2l0 setcache 1 -r3 12
```

When you press **<Enter>** to complete the command, a message informs you that the SPs will be rebooted.

5. Wait two minutes; then stop and restart the RAID agent:

```
/etc/init.d/raid5 stop
/etc/init.d/raid5 start
```

6. Create LUNs for each SP using the **-r3** option to allocate RAID-3 memory; for example:

```
raidcli -d sc4d2l0 bind r3 2 a2 b2 c2 d2 e2 -r3 4
raidcli -d sc4d2l0 bind r3 3 a3 b3 c3 d3 e3 -r3 4
```

In this example:

- `r3` sets the RAID level to RAID-3
- `2` sets the LUN number to 2; `3` sets the LUN number to 3
- `a2 b2 c2 d2 e2` binds disks to LUN 2 and `a3 b3 c3 d3e3` binds disks to LUN 3; five disks are required for each RAID-3 LUN
- `-r3 4` sets RAID-3 cache memory for each LUN to 4 MB

For failover, you must set a total of the memory for this SP's LUNs plus that of the other SP; each SP must have the same amount of memory allocated. That is, the RAID-3 cache memory you set with *raidcli bind -r3* for each SP must be half the amount you have set with *raidcli setcache -r3* in step 4. See Section 4.7.3, "Planning a RAID-3 Bind," earlier in this chapter.

4.9 Binding Disks Into RAID Units

The physical disk unit number is also known as the logical unit number, or LUN. The LUN is a decimal number between 0 and 7. You can bind disks into RAID levels 0, 1, 1_0 (0+1), 3, or 5.

To group physical disks into a RAID unit, use

`raidcli -d` *device* `bind` *raid-type lun-number disk-names* [*optional-args*]

Variables in this syntax are explained below.

-d *device* Target RAID device, as returned by *raidcli getagent*.

raid-type Choices are

- r0: RAID-0
- r1: RAID-1
- r1_0: RAID-1_0
- r3: RAID-3
- r5: RAID-5
- hs: hot spare

lun-number Logical unit number to assign the unit (0-7).

disk-names Indicates which physical disks to bind, in the format bd, where b is the physical bus name (a through e) and d is the device number on the bus (0 through 3). For example, a0 represents the device 0 on bus A, and e2 represents device 2 on bus E.

A RAID-0 bind requires a minimum of 3 (maximum 16) disks.

A RAID-1 bind requires 2 disks on separate buses.

A RAID-1_0 bind requires an even number of disks (minimum 2; maximum 16). For high availability, a member of each image pair must be on a different bus. Select the disks in this order: first disk on first bus, first disk on second bus, second disk on first bus, second disk on second bus, third disk on first bus, third disk on second bus, and so on.

A RAID-5 bind also requires separate buses, each with five disk modules. Legal RAID-5 bind configurations are

- a0 b0 c0 d0 e0
- a1 b1 c1 d1 e1
- a2 b2 c2 d2 e2
- a3 b3 c3 d3 e3

A RAID-5 bind requires a minimum of 3 disks and a maximum of 15 disks. For high availability, use groups of 5.

A hot spare bind requires 1 disk, which cannot be A0, B0, C0, D0, E0, or A3. The capacity of the hot spare must be at least as great as the capacity of the largest disk module it might replace.

All disks in a bind should have the same capacity, so that disk space is used fully.

Note: If the LUN you are creating is RAID-3, you must allocate SP memory before starting the bind; see Section 4.8, “Allocating Memory for RAID-3 LUNs.”

The optional arguments are as follows.

-r *rebuild-time* Maximum time in hours to rebuild a replacement disk. Default is 4 hours; legal values are any number greater than or equal to 0.

A rebuild time of 2 hours rebuilds the disk more quickly, but degrades response time slightly. A rebuild time of 0 hours rebuilds as quickly as possible but degrades performance significantly.

If the customer requires fast response time and they want to minimize degradation to normal I/O activity, enter a longer period of time, such as 24 hours. The customer can change the rebuild time later without damaging the information stored on the physical disk unit.

-s *stripe-size* Number of blocks per physical disk in a RAID-5 stripe. Default is 128; legal values are any number greater than 0.

The smaller the stripe element size, the more efficient the distribution of data read or written. However, if the stripe size is too small for a single host I/O operation, the operation requires accessing two stripes, thus causing the hardware to read and/or write from two disk modules instead of one. Generally, it is best to use the smallest stripe element size that will rarely force access to another stripe. The default stripe element size is 128 sectors. The size should be an even multiple of 16 sectors (8 KB).

Note: This value is ignored for hot spares. For RAID-3 binds, the only legal value is 1.

-c *cache-flags* Values are:

- none: no caching
- read: read caching
- write: write caching
- rw: read and write caching

The default is none.

-z *stripe-count* Sets the number of stripes in a LUN. For example, if you bind a RAID-5 LUN with a stripe count of 2, you partition the LUN into two stripes, thus preventing access to the remaining available space. This option is useful for short bind operations. Legal values are any number greater than or equal to 0. The default value is 0, which binds the maximum number of stripes available.

-r3 *raid3-memory-size*

Sets amount of RAID-3 memory t use for each LUN. See Section 4.8, "Allocating Memory for RAID-3 LUNs." This option is available only for firmware revision 9.0 and higher (SP 7305) and RAID agent 1.55 and higher.

Note: Although *bind* returns immediate status, the bind itself does not complete for 15 to 60 minutes, depending on system traffic. Use *getlun* to monitor the progress of the bind; *getlun* returns the percent bound. When the bind is complete, each disk is noted as "Bound But Not Assigned" in the *getlun* output.

The following example binds disks A0, B0, C0, D0, and E0 into a RAID-5 logical unit with a logical unit number of 3, a four-hour maximum rebuild time, and a 128-block stripe size per physical disk, with read cache enabled.

```
raidcli -d sc4d2l0 bind r5 3 a0 b0 c0 d0 e0 -r 4 -s 128 -c read
```

The following example binds A2 and B2 into a RAID-1 logical unit with a LUN number of 2 and a four-hour maximum rebuild time, with read cache enabled.

```
raidcli -d sc4d2l0 bind r1 2 a2 b2 -r 4 -c read
```

The following example binds disks A1, B1, C1, and D1 into a RAID-1_0 logical unit with a LUN number of 1, a four-hour maximum rebuild time, and a 128-block stripe size per physical disk, with read cache enabled.

```
raidcli -d sc4d2l0 bind r1_0 1 a1 b1 c1 d1 -r 4 -s 128 -c read
```

The following example binds A1, B1, C1, D1, and E1 into a RAID-3 logical unit with a LUN number of 3.

```
raidcli -d sc4d2l0 bind r3 3 a1 b1 c1 d1 e1
```

The following example binds A3, B3, C3, D3, and E3 into a RAID-0 logical unit with a LUN number of 3, and a 128-block stripe size per physical disk, with read cache enabled.

```
raidcli -d sc4d2l0 bind r0 3 a3 b3 c3 d3 e3 -s 128 -c read
```

The following example binds disk E3 as a hot spare with a LUN number of 7.

```
raidcli -d sc4d2l0 bind hs 7 e3
```

There is no output for *raidcli* with the *bind* parameter.

Note: For complete messages, it is recommended that you use the `-v` option.

To display information on a logical unit and the components in it, use *getlun*. To change parameters for a logical unit, use *chglun*. These and all other *raidcli* parameters are explained in Appendix B, “The raidcli Command-Line Interface.”

4.10 Enabling Command-Tagged Queuing

Command-tagged queueing (CTQ) allows multiple outstanding commands to a single SCSI target, that is, a LUN in a Challenge RAID storage system, resulting in increased I/O performance.

The Sauna SP supports SCSI-2 queuing of requests for its LUNs. The requests are first-come, first-served in that all requests could be sent to one LUN from one initiator and cause a Queue Full status (unexpected SCSI status byte 0x28) to be returned for all other I_T_L SCSI selections. This condition would continue until one of the outstanding requests completes and thus frees queue space. The Sauna SP can handle up to 96 CTQs.

For IRIX versions before 6.1, if Queue Full status is returned for a given I/O request, that request is retried. If the requests cannot be sent to the SP after four retries, the request is aborted, which, in the case of a write request, can have unfortunate consequences. The default CTQ depth (*dkmaxq*) for the device driver is 16 (maximum 96); this value is stored in */var/sysgen/master.d/dksc*. Any changes made to the *dkmaxq* value requests an *lboot*, followed by a reboot.

When CTQ is enabled (with *fx*) for a given LUN, the default CTQ depth for the LUN is 2; this value is stored in the LUN's volume header. You must use *fx* to change this value. Table 4-2 and Table 4-3 define the maximum CTQ depth values per LUN for single-hosted SPs and dual-hosted SPs, respectively.

Table 4-2 Maximum CTQ Depths per LUN: Single-Hosted SPs

Number of LUNs	Maximum CTQ Depth per LUN	S1 Driver dkmaxq Value	SCIP driver dkmaxq Value
1	96	96	96
2	48	96	48
3	32	96	32
4	24	96	24

Table 4-3 Maximum CTQ Depths per LUN: Dual-Hosted SPs

Number of LUNs	Maximum CTQ Depth per LUN	S1 Driver dkmaxq Value	SCIP Driver dkmaxq Value
1	48	48	48
2	24	48	24
3	16	48	16
4	12	48	12

For optimum system performance, enable command-tagged queuing. Table 4-4 shows performance benefits of CTQ.

Table 4-4 CTQ Performance Benefits for 2 K Random Read, 16 Threads

CTQ	avque	r+w/s	blks/s	w/s	wblks/s	avwait (μs)	avserv (ms)
Disabled	16.0	69	137	0	0	217.6	14.5
Enabled	16.0	304	607	0	0	48.8	3.3

The *fx* program syntax is as follows.

`fx -x "`*controllertype*`(`*controller_number*`,`*drive_number*`,`*lun_number*`)"`

Follow these steps:

1. Enter the *fx* command with appropriate parameters; for example:

   ```
   fx -x "dksc(6,2,2)"
   ```

   Output such as the following appears:

   ```
   fx version 5.3, Jan 3, 1995
   ...opening dksc(6,2,2)
   ...controller test...OK
   Scsi drive type == SGI RAID 5 0757
   fx: Warning: bad sgilabel on disk

   creating new sgilabel


   ----- please choose one (? for help, .. to quit this menu)-----
   [exi]t             [d]ebug/           [l]abel/           [a]uto
   [b]adblock/        [exe]rcise/        [r]epartition/     [f]ormat
   ```

2. Update parameters; at the `fx>` prompt, type the following:

   `fx> /label/set/param`

   Output such as the following appears:

   ```
   fx/label/set/parameters: Error correction = (enabled)
   fx/label/set/parameters: Data transfer on error = (enabled)
   fx/label/set/parameters: Report recovered errors = (enabled)
   fx/label/set/parameters: Delay for error recovery = (enabled)
   fx/label/set/parameters: Err retry count = (0)
   fx/label/set/parameters: Transfer of bad data blocks = (enabled)
   fx/label/set/parameters: Auto bad block reallocation (write) =
   (enabled)
   fx/label/set/parameters: Auto bad block reallocation (read) =
   (enabled)
   fx/label/set/parameters: Read ahead caching = (disabled)
   ```

3. At the `Enable CTQ` prompt, type **`enable`**:

   ```
   fx/label/set/parameters: Enable CTQ = (disabled) enable
   ```

4. At the `CTQ depth` prompt, type **`10`**:

   ```
   fx/label/set/parameters: CTQ depth = (2) 10
   ```

   Output such as the following appears.

   ```
   fx/label/set/parameters: Read buffer ratio = (0/256)
   fx/label/set/parameters: Write buffer ratio = (0/256)
    * * * * * W A R N I N G * * * * *
   ```

5. At the following prompt in the last line above, type **yes**.

   ```
   about to modify drive parameters on disk dksc(6,2,2)! ok? yes
   ```

   The following output appears:

   ```
   ----- please choose one (? for help, .. to quit this menu)-----
   [exi]t             [d]ebug/           [l]abel/           [a]uto
   [b]adblock/        [exe]rcise/        [r]epartition/     [f]ormat
   ```

6. Type **exit** to exit *fx*. The following message appears:

   ```
   label info has changed for disk dksc(6,2,2). write out changes? (yes)
   ```

7. Type **y** to write the changes to disk.

4.11 Understanding Dual Processors, Load Balancing, and Device Names

If the Challenge RAID storage system has two SPs (split-bus, dual-bus / dual-initiator, or dual-adapter / dual-processor configuration), the customer can choose which RAID-5 disk units (LUNs) to bind on each processor. Thus, the load on the SCSI-2 adapters and processors can be balanced.

The processor on which you bind a LUN is the default owner of the physical disk unit. The route through the processor that owns a LUN is the primary route to the LUN, and determines the device name of the LUN. In a dual-processor system, the route through the other processor is the secondary route to the LUN.

4.12 Setting Up Caching

Storage-system caching requires

- two storage-control processors
- fully charged storage-system backup battery
- minimum of 8 MB of memory for each processor
- cache enabling using the *raidcli setcache* command, as explained in this chapter
- disk modules in slots A0, B0, C0, D0, and E0 as a fast repository for cached data

Caching cannot occur unless all these conditions are met.

This section explains

- setting cache parameters
- upgrading Challenge RAID to support caching
- changing cache unit parameters

Note: If a disk module in location A0, B0, C0, D0, or E0 fails, caching is disabled until the failed module is replaced. If an SP fails, caching is disabled until the failed SP is replaced.

4.12.1 Setting Cache Parameters

The cache parameters you specify for the entire storage system are the cache size of 8 or 64 MB, depending on the amount of memory the storage-control processors have, and the cache page size, of 2, 4, 8, or 16 KB.

To set up caching, use the *raidcli setcache* command:

`raidcli -d` *device* `setcache` *enable* | *disable* `[-u` *usable*`] [-p` *page*`] [-l` *low*`] [-h` *high*`]`

In this syntax, variables mean the following.

enable | *disable* Enter 1 to enable caching or 0 to disable caching.

Note: When enable / disable is the only option used with *setcache*, 1 enables system read and write cache and 0 disables system read and write cache.

Using *setcache* 1 with options takes 45 seconds while cache states are disabled; cache dumping might take some time.

-u *usable* Size in megabytes to use for caching, not greater than the SP memory size as displayed by *raidcli getagent* (see "getagent" in Appendix B). Valid values are 0 through 64 in increments of 1 MB.

The command-line interface does not let you specify more memory than you have. If you specify less than you have, the remaining memory is unused.

Note: For mirrored caching, both SPs must have the same amount of cache memory in order for caching to be preserved in the event of shutdown or other power loss.

-p *page* Size in KB of pages into which to partition the cache. Valid sizes are 2, 4, 8, and 16. The default is 2, regardless of whether caching is enabled or disabled.

-l *low* Percentage of cache full that discontinues flushing. Valid values are 0 through 100; the default is 50, regardless of whether caching is enabled or disabled.

-h *high* Percentage of cache full that initiates flushing. Valid values are 0 through 100; the setting must be greater than the low watermark. The default is 75, regardless of whether caching is enabled or disabled.

-sma *system-memory-SPA*, -smb *system-memory-SPB*
Set the system memory on SP A or B, respectively.

Note: This option reboots each SP in the system. Rebooting takes several minutes, during which *raidcli getagent* is the only command accepted. When the reboot completes, stop and restart the agent before any *raidcli* commands are accepted.

-sta *read-cache-state-SPA*, stb *read-cache-state-SPB*
Set the read cache state for SP A or B, respectively: 0 to disable, 1 to enable.

-rca *read-cache-size-SPA*, -rcb *read-cache-size-SPB*
Set the read cache size in MB for SP A and B, respectively.

-r3 *raid3-memory-size*

Sets the RAID-3 memory size. For details on allocating memory for RAID-3, see Section 4.7.3, "Planning a RAID-3 Bind" and Section 4.8, "Allocating Memory for RAID-3 LUNs" earlier in this chapter.

Note: The sum of an SP's write cache size, system memory size, and read cache size must be less than or equal to the SP's physical memory size.

You can change the cache size, the cache page size values, or the type of caching for any physical disk unit, without affecting the information stored on it. Follow these steps:

1. To disable the cache, enter

   `raidcli -d` *device* **`setcache 0`**

2. Wait for the cache memory to be written to disk, which may take several minutes.
3. To reassign cache size and reenable caching, enter

   `raidcli -d` *device* **`setcache 1 -u 64 [`***parameters***`]`**

Note: Before changing a cache parameter, you must always disable cache.

To view cache statistics, use the *getcache* parameter, as explained in Appendix B.

Table 4-5 gives suggested sizes for read and write partitions for PowerPC-based SPs (firmware revision 8.50 and higher; SP 7305).

Table 4-5 Read and Write Partition Sizes (in MB): Firmware Revision 8.50 and Higher (No RAID-3 Memory Allocated)

Total SP Memory	FLARE(P)	Available Cache Memory	Suggested Read Partition	Suggested Write Partition
8	4	4	2	2
16	4	12	8	4
32	4	28	18	10
64	4	60	40	20

In a system with RAID-3 LUNs (firmware revision 9.0 and higher), for best performance you should allocate at least 6 MB of storage-system memory for each RAID-3 LUN. For two RAID-3 LUNs, allocate 12 MB of storage-system memory for the RAID-3 memory partition and 6 MB to each of the two RAID-3 LUNs when you bind them. Since an SP requires 4 MB of storage-system memory for system buffers, each SP needs 16 MB of memory. Failover requires 32 MB.

To maximize RAID-3 performance, bind all LUNs in the system to RAID-3 only. In a storage system with only fast RAID-3 LUNs, no memory is allocated for cache memory, read partitions, or write partitions. For information on RAID-3 memory requirements, see Section 4.7.3, "Planning a RAID-3 Bind," earlier in this chapter.

Table 4-6 gives suggested sizes for read and write partitions for AMD-based SPs (model 7264).

Table 4-6 Read and Write Partition Sizes in MB: Firmware Revision 8.00-8.49

Total SP Memory	FLARE(P)	Available Cache Memory	Suggested Read Partition	Suggested Write Partition
8	2	6	3	3
16	2	14	8	6
32	2	30	18	12
64	2	62	40	22

4.12.2 Upgrading Challenge RAID to Support Caching

Once you have installed the necessary hardware components, follow these steps to set up caching. See Section 4.12.1, “Setting Cache Parameters,” for instructions.

1. Enable the read and / or write caches for the physical disk units that you want to use for caching.
2. Specify the cache page size and cache size parameters.
3. Enable storage-system caching with *raidcli setcache*.

4.12.3 Changing Caching Parameters

To change the caching parameter for a LUN, follow these steps:

1. Run *raidcli getcache* to determine if caching is enabled. If it is enabled, make sure both storage-control processors are powered on; check the status lights or run *raidcli getcrus* (see Appendix B for information).
2. If caching is enabled, disable it by typing

   `raidcli -d` *device* **`setcache 0`**
3. Wait for the cache memory to be written to disk, which may take several minutes.
4. Follow instructions in Section 4.12.1, “Setting Cache Parameters,” earlier in this chapter.

Chapter 5

Maintaining and Upgrading Disk Modules

This chapter explains replacing and installing disk modules:

- handling FRUs
- identifying a failed disk module
- unbinding the disk
- replacing a disk module
- installing an add-on disk module array

Before you follow procedures explained in this chapter, note the following points:

Caution: To prevent thermal shutdown of the storage system, never operate it for more than two minutes with the fan module open.

- Never remove more than one disk module or disk filler module at a time.
- Disk modules A0, B0, C0, D0, and E0 contain the cache vault for the Challenge RAID storage system. Also, disk modules A0, B0, C0, and, if installed, A3 contain the licensed internal code. Do not remove these disk modules for use in other disk module positions.
- Although you can remove a disk module (other than those in bus 0) without damaging the disk data, do this only when the disk module has actually failed. Never remove a disk module unless you are installing its replacement.
- Removing the wrong drive module can introduce an additional fault that shuts down the physical disk containing the failed module. Before removing a disk module, verify that the suspected module has actually failed.
- Use only Challenge RAID disk modules as replacements; only they contain the correct device firmware. For replacement drive marketing codes, see Table 1-2 in Chapter 1. Other disk modules, even those from other Silicon Graphics equipment, will not work. Do not mix disk modules of different capacities within one array.
- When removed from the chassis, the disk modules are extremely sensitive to shock and vibration. Even a slight jar can severely damage them.
- Do not operate the system with an open SP slot.
- SCSI terminators must be in place at all times during operation. Removing them causes the storage system to crash.

You upgrade Challenge RAID by adding optional modules that are field-replaceable units (FRUs). You repair Challenge RAID by replacing faulty FRUs. You can add or replace the following FRUs while the storage system is powered on:

- disk modules (the customer can also replace these)
- storage control processors (SPs) and SP memory modules (SIMMs)
- power supply modules (voltage semi-regulated converters, or VSCs)
- fan module
- optional battery backup unit (BBU)

You can also replace the external SCSI bus cables, power cord, and SCSI terminator plugs, as shown in Figure 6-3 and Figure 6-4 in Chapter 6. See Chapter 3, “Installing a Challenge RAID Storage System,” for information on where and how to connect them.

Note: For information on replacing SCSI-2 adapters in the Challenge server, see the Challenge server documentation.

Figure 5-1 diagrams location of disks, which can be replaced by owners or SSEs.

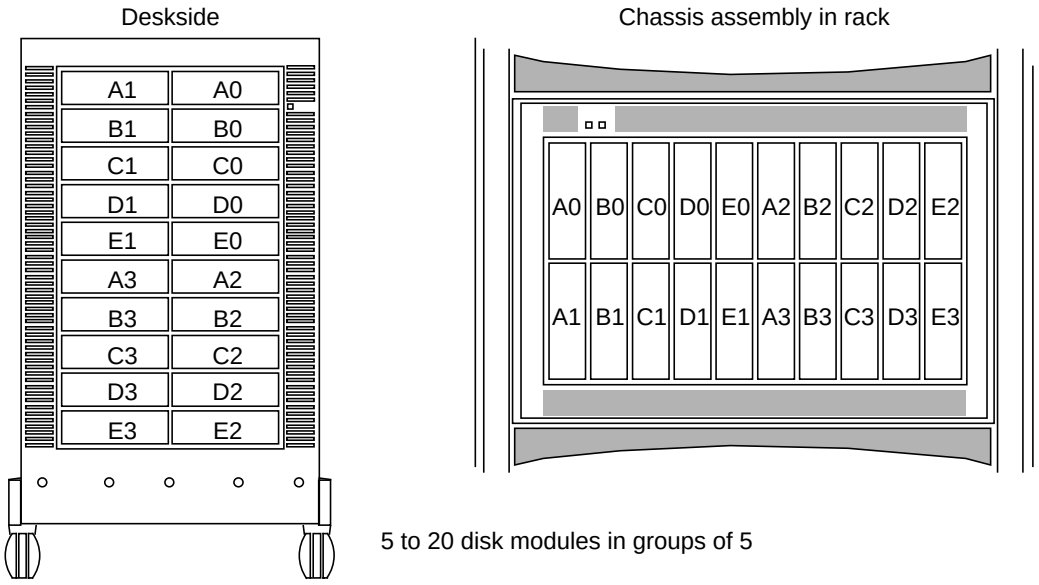

Figure 5-1 Location of Disks (Front of Challenge RAID)

5.1 Handling FRUs

This section describes the precautions that you must take and the general procedures you must follow when removing, installing, and storing FRUs.

5.1.1 Avoiding Electrostatic Discharge (ESD) Damage

The cover(s) and filler panel(s) installed on the Challenge RAID storage system protect the electronic circuits inside the equipment from electrostatic discharge (ESD) damage. However, when you remove these covers and filler panels to replace or install subassemblies, you can inadvertently damage the sensitive electronic circuits in the equipment by simply touching them. Electrostatic charge that has accumulated on your body discharges through the circuits. If the air in the work area is very dry, running a humidifier in the work area will help decrease the risk of ESD damage. You must follow the procedures below to prevent damage to the equipment.

Caution: Read and understand the following instructions before you remove the cover(s) or panel(s) from the equipment.

- Provide enough room to work on the equipment. Clear the work site of any unnecessary materials or materials that naturally build up electrostatic charge, such as foam packaging, foam cups, cellophane wrappers, and similar materials.
- Do not remove replacement or upgrade subassemblies from their antistatic packaging until the exact moment that you are ready to install them.
- Gather the tools, manuals, a wrist strap, and all other materials you will need before you remove covers and panels from the equipment. After you remove a cover or panel, avoid moving away from the work site; otherwise, you may build up an electrostatic charge.
- Use a wrist strap when handling circuit boards or when touching the electronic circuits inside the equipment.
- Replace the cover(s) or panel(s) on the equipment as soon as possible so that the electronic circuits are protected.

5.1.2 Precautions for Removing, Installing, or Storing FRUs

When removing or installing a FRU, never use excessive force. If you have difficulty removing or installing a FRU, read the procedures again. Once you have removed a FRU, handle it gently, because a sudden jar or drop could permanently damage it.

The replacement or add-on FRU is shipped in a specially designed shipping container. Store the FRU in this container, and use this container if you need to return the FRU for repair. The storage location of a disk module, SP, power supply module, fan module, or battery backup unit must be maintained within the nonoperating limits specified in Appendix A.

5.2 Identifying and Verifying a Failed Disk Module

If you have determined that a module has failed by examining the cabinet fault light or by using the *raidcli getdisk* or *raidcli getcrus* command, as explained in Chapter 3 in this guide, you can replace the defective module and rebuild your data without powering off the Challenge RAID storage system or interrupting user applications.

Caution: Removing the wrong drive module can introduce an additional fault that shuts down the physical disk containing the failed module. Before removing a disk module, verify that the suspected module has actually failed.

The fault indicator on a disk module does not necessarily mean that the drive itself has failed. Failure of a SCSI bus, for example, lights the fault indicator on each disk module on that bus.

Note: If a disk module in location A0, B0, C0, D0, or E0 fails, caching is disabled until the failed module is replaced.

To verify a suspected disk module failure, follow these steps:

1. Look for the module with its amber fault light on. Figure 5-2 shows the fault indicator light and other lights on a disk module.

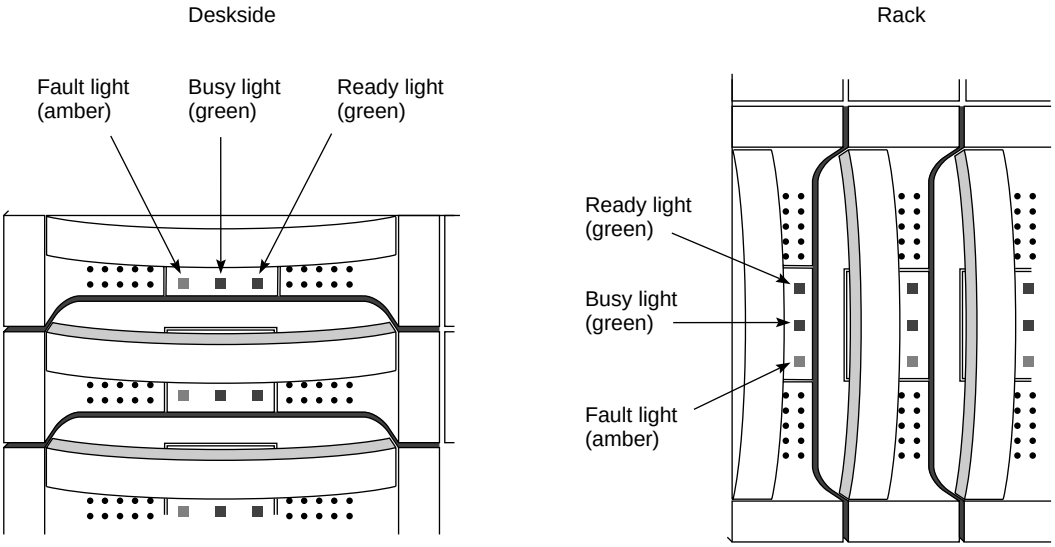

Figure 5-2 Disk Module Status Lights

2. Determine the failed module's ID; see Figure 5-1.

Caution: Use only Challenge RAID disk modules to replace failed disk modules. Order them from the Silicon Graphics hotline. Challenge RAID disk modules contain proprietary firmware that the storage system requires for correct functioning. Using any other disks, including those from other Silicon Graphics systems, can cause failure of the storage system.

3. If you have not already checked the module status with *raidcli getdisk*, do so now; see Appendix B.

4. If you have not already checked the unsolicited error log with *raidcli getlog* for a message about the disk module, as explained in Appendix B, do so now.

   A message about the disk module contains its module ID (such as A0 or B3). Check for any other messages that indicate a related failure, such as failure of a SCSI bus or a general shutdown of a chassis, that might mean the disk module itself has not failed.

   Note: If you are using storage system caching, the system uses modules A0, B0, C0, D0, and E0 for its cache vault. If one of these modules fails, the storage system dumps its cache image to the remaining modules in the vault; then it writes all dirty (modified) pages to disk and disables caching. The cache status changes, as indicated in the output of the *raidcli getcache* command. Caching remains disabled until you insert a replacement module and the storage system rebuilds the module into the physical disk unit. For information on caching, see Section 4.12, "Setting Up Caching," in Chapter 4.

Caution: Although you can remove a disk module without damaging the disk data, do this only when the disk module has actually failed. Never remove a disk module unless absolutely necessary, and only when you have its replacement available. Never replace more than one disk module at a time; use only correct disk modules available from Silicon Graphics, Inc.

5.3 Unbinding the Disk

When you change a physical disk configuration, you change the bound configuration of a physical disk unit. Physical disk unit configuration changes when you add or remove a disk module, or physically move one or more disk modules to different slots in the chassis.

Caution: Unbinding destroys all the data on a physical disk unit. Before unbinding any physical disk unit, make a backup copy of any data on the unit that you want to retain.

To unbind a disk, use the *unbind* parameter with the *raidcli* command. See Section 4.6, "Unbinding Disks," in Chapter 4. Use *raidcli bind* to configure disks, as explained in Chapter 4.

5.4 Replacing a Disk Module

This section explains

- removing the failed disk module
- installing a replacement disk module
- updating the disk module firmware

Caution: Use only Challenge RAID disk modules as replacements; only they contain the correct device firmware. For replacement drive marketing codes, see Table 1-2 in Chapter 1. Other disk modules, even those from other Silicon Graphics equipment, will not work. Do not mix disk modules of different capacities within one array. If you replace an unbound disk module, you must update the firmware as explained in Section 5.4.3, "Updating the Disk Module Firmware."

5.4.1 Removing a Failed Disk Module

You can replace a failed disk module while the storage system is powered on. If necessary, you can also replace a disk module that has not failed, such as a module that has reported many "soft" errors. When replacing a module that *has not* failed, you must do so while the storage system is powered on so that the SP knows the module is being replaced.

Caution: To maintain proper cooling in the storage system, never remove a disk module until you are ready to install a replacement. Never remove more than one disk module at a time.

To remove a disk module, follow these steps:

1. Verify that the suspected module has actually failed.

   Caution: If you remove the wrong disk module, you introduce an additional fault that shuts down the physical disk containing the failed module. In this situation, the operating system software cannot access the physical disk until you initialize it again.

2. When a disk fails, the SP automatically unbinds it. To check that the disk has been unbound, use the following *raidcli* command for the disk module position in question:

   `raidcli -d` *device* `getdisk [`*diskposition*`]`

   If necessary, get the device name first with *raidcli getagent*. See Appendix B for information on these parameters.

3. Read Section 5.1, "Handling FRUs," earlier in this chapter.
4. Locate the disk module that you want to remove; see Figure 5-1 if necessary.
5. Position the new disk module in its antistatic packaging within reach of the storage system.

6. If you are using an ESD wrist strap, attach its clip to the ESD bracket at the bottom of the storage system, as shown in Figure 5-3. Put the wrist band around your wrist with the metal button against your skin.

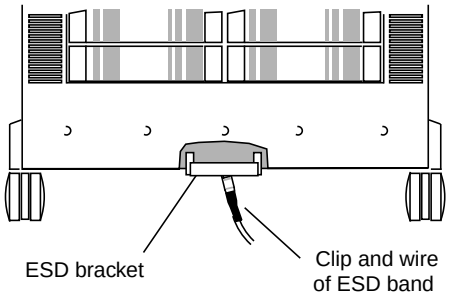

Figure 5-3 Attaching the ESD Clip to the ESD Bracket on the Deskside Storage System

Figure 5-4 shows where to attach the clip on a rack storage system.

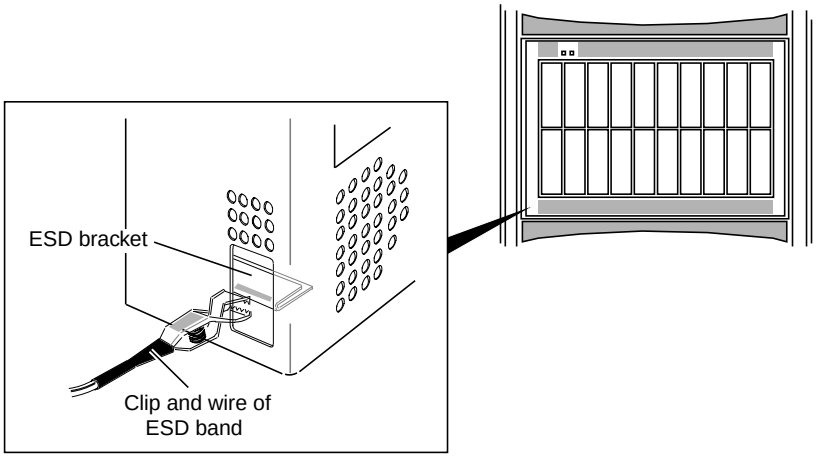

Figure 5-4 Attaching the ESD Clip to the ESD Bracket on a Rack Storage System

7. Make sure the disk has stopped spinning and the heads have unloaded.

8. Grasp the disk module by its handle and pull it partway out of the cabinet, as shown in Figure 5-5.

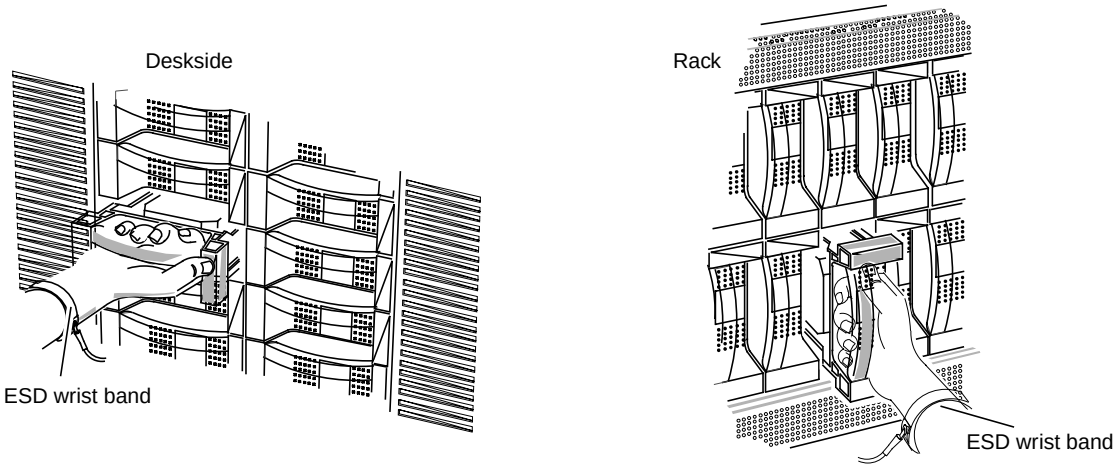

Figure 5-5 Pulling Out a Disk Module

Caution: Never remove more than one disk module at a time.

Warning: **When removing a disk module from an upper chassis assembly in a Challenge RAID rack system, make sure that you adequately balance the weight of the disk module.**

9. Supporting the disk module with your free hand, pull it all the way out of the cabinet, as shown in Figure 5-6.

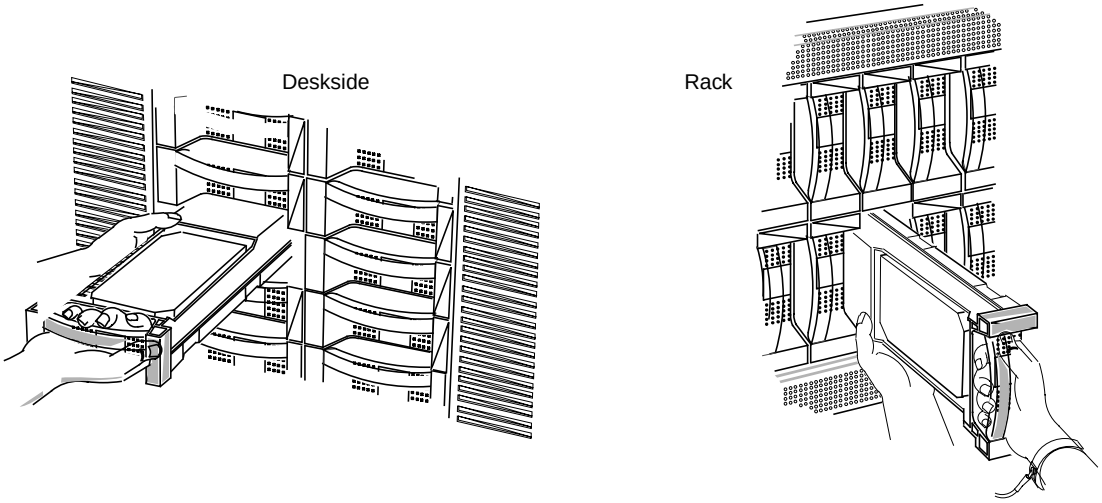

Figure 5-6 Removing a Disk Module

Caution: When removed from the chassis, the disk modules are extremely sensitive to shock and vibration. Even a slight jar can severely damage them.

10. If the ID number for the compartment from which you removed the drive has not been written on the label on the side of the disk module, write it or check the appropriate box; for example, A2.

    For the compartment ID numbers, refer to Figure 5-1.

11. Put the failed disk module in an antistatic bag and store it in a place where it will not be damaged.

    Caution: Before installing a replacement module, wait at least 15 seconds after removing the failed module to allow the SP time to recognize that the module has been removed. If you insert the replacement module too soon, the SP may report the replacement module as defective.

5.4.2 Installing a Replacement Disk Module

To install the replacement disk module, follow these steps:

1. Touch the new disk module's antistatic packaging to discharge it and the drive module. Remove the new disk module from its packaging.

   Caution: The disk module is extremely sensitive to shock and vibration. Even a slight jar can severely damage it.

2. Before installing a replacement module, wait at least 15 seconds after removing the failed module to allow the SP time to recognize that the module has been removed. If you insert the replacement module too soon, the SP may report the replacement module as defective.
3. Position the disk module in its antistatic packaging within reach of the storage system.
4. Locate the slot where you will install the disk module; see Figure 5-1.
5. Engage the disk module's rail in the chassis rail slot, as shown in Figure 5-7.

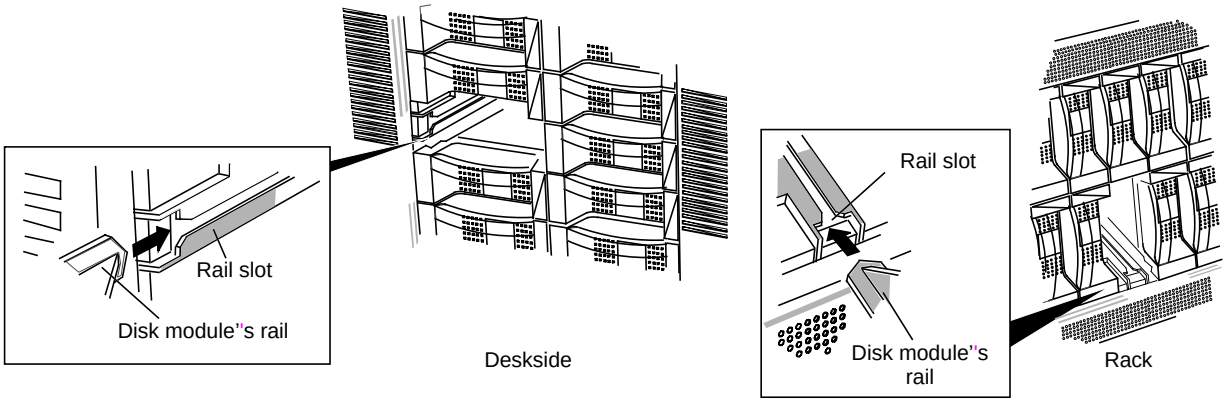

Figure 5-7 Engaging the Disk Module Rail

6. Engage the disk module's guide in the chassis guide slot, as shown in Figure 5-8.

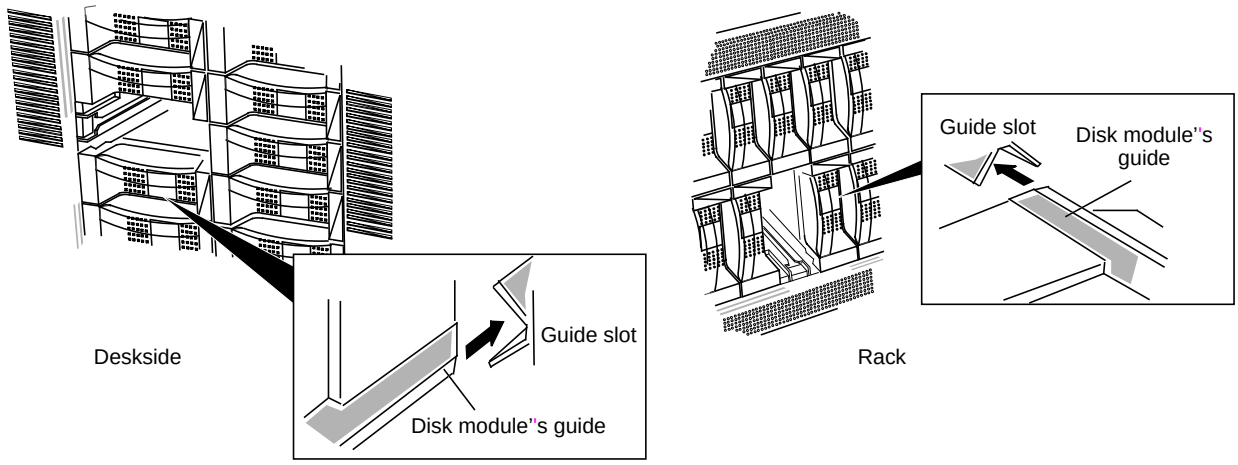

Figure 5-8 Engaging the Disk Module Guide

7. Insert the disk module, as shown in Figure 5-9. Make sure it is completely seated in the slot.

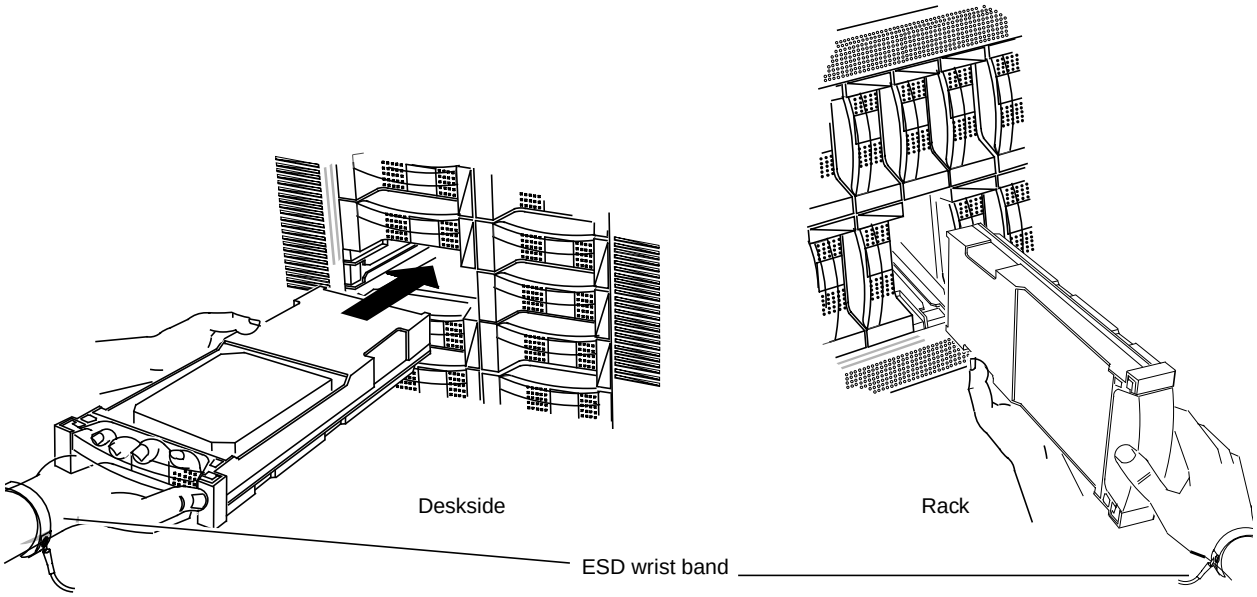

Figure 5-9 Inserting the Replacement Disk Module

8. Remove and store the ESD wrist band, if you are using one.
9. The SP formats and checks the new module, and then begins to reconstruct the data. While rebuilding occurs, you have uninterrupted access to information on the physical disk unit.

The default rebuild period is 4 hours.

Note: During the rebuild period, performance might degrade slightly, depending on the rebuild time specified and on I/O bus activity.

For more information on changing the default rebuild period, use *raidcli bind*, as explained in Appendix B.

5.4.3 Updating the Disk Module Firmware

After replacing a failed unbound disk module (A0, B0, C0, or A3), update the firmware on the Challenge RAID SP. The firmware (FLARE code) is stored in directories in */usr/raid5/flare*:

- */usr/raid5/flare/sauna7/flarecode.bin*: use this for AMD-based SPs with FLARE code rev. level 7.xx
- */usr/raid5/flare/sauna8/flarecode.bin*: use this for AMD-based SPs with FLARE code rev. level 8.00 through 8.49
- release 2.0 and earlier: */usr/raid5/flare/phoenix8/flarecode.bin*: use this for PowerPC-based SPs with FLARE code rev. level 8.50 through 8.99
- release 2.1 and later: */usr/raid5/flare/phoenix9/flarecode.bin*: use this for PowerPC-based SPs with FLARE code rev. level 9.00 through 9.99

Follow these steps:

1. Quiesce the bus, disabling all applications. Make sure that only the RAID agent is running.
2. Enter as root:

   raidcli -d *device* **firmware /usr/raid5/flare/***directory***/flarecode.bin**

In this string, the new firmware image */usr/raid5/flare/<directory>/flarecode.bin* contains microcode that runs on the SP and also a microcode image destined for the SP PROM, which runs the power-on diagnostics.

You must use this command every time you replace a failed unbound disk module (A0, B0, C0, or A3).

The image in the file given in the command contains microcode that runs on the storage-control processor and possibly also a microcode image destined for the storage-control processor PROM, which runs the power-on diagnostics.

Note: Once the microcode has been downloaded, each SP in the cabinet reboots. The reboot process takes several minutes, during which time no command-line interface commands are accepted other than *getagent*. When the code is downloaded, the system reboots the SPs. After the download is complete, stop and restart the agent.

This command has no output. Use *getagent* or *getsp* to return the new firmware version.

5.5 Installing an Add-On Disk Module Array

Table 5-1 gives marketing codes for add-on disk module arrays. Add-on disk modules are available only in arrays of five.

Table 5-1 Ordering Add-On Disk Module Sets

Unit	Marketing Code
Add-on five 2 GB drives	P-S-RAID-5X2
Add-on five 4.3 GB drives	P-S-RAID-5X4
Base array with five 2 GB drives	P-S-RAID-B5X2
Base array with five 4 GB drives	P-S-RAID-B5X4

Caution: Use only Challenge RAID disk modules; only they contain the correct device firmware. Other disk modules, even those from other Silicon Graphics equipment, will not work. Do not mix disk modules of different capacities within one array. Do not remove disk modules from bus 0 (slots A0, B0, C0, D0, and E0) or for A3 (if installed) for use in other disk module positions.

Installing an add-on disk module array consists of two procedures:

- inserting the new disk module array
- creating device nodes and binding the disks

5.5.1 Inserting the New Disk Module Array

Follow these steps:

Caution: Never remove more than one disk module or disk filler module at a time.

1. Read Section 5.1.1, “Avoiding Electrostatic Discharge (ESD) Damage,” earlier in this chapter.
2. Position the new disk modules in their antistatic packaging within reach of the storage system.
3. If you are using a wrist band, attach its clip to the ESD bracket on the bottom of the storage system or chassis assembly. Put the wrist band around your wrist with the metal button against your skin.

4. Locate the slots where you will install the add-on disk modules; consult Figure 5-1 if necessary.

Warning: **In a rack, you need not complete each chassis assembly that is partially filled before installing more disk modules in the next chassis assembly. However, avoid making the rack top-heavy.**

Fill the slots in this order:

- first, in this order: A0, B0, C0, D0, E0
- next, in this order: A1, B1, C1, D1, E1
- next, in this order: A2, B2, C2, D2, E2
- next, in this order: A3, B3, C3, D3, E3

5. Grasp the filler module for the first slot and pull it out of the cabinet; set it aside. If you cannot grasp the module, use a medium-size flat-blade screwdriver to pry it out gently.
6. Touch the new disk module's antistatic packaging to discharge it and the drive module. Remove the new disk module from its packaging.
7. On the label on the side of the disk module, write the ID number for the compartment into which the drive is going. You can either write the slot position on the label in the corresponding place on the matrix or make a check mark in the position to indicate the slot that the disk module occupies. Figure 5-10 shows these two ways of labeling disk module B0.

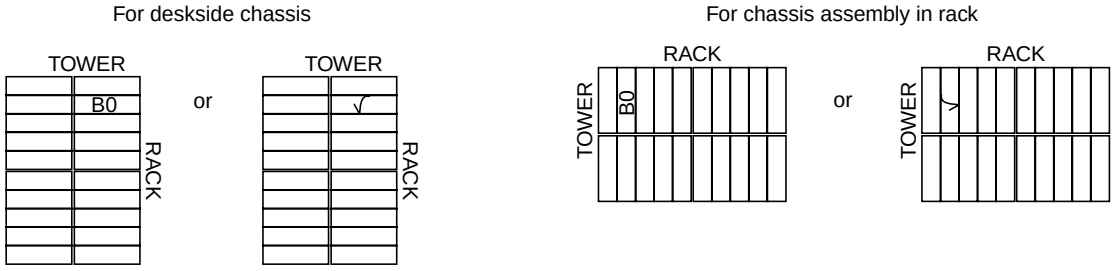

Figure 5-10 Marking the Label for Disk Module B0

For reference, Figure 5-11 diagrams all disk module locations.

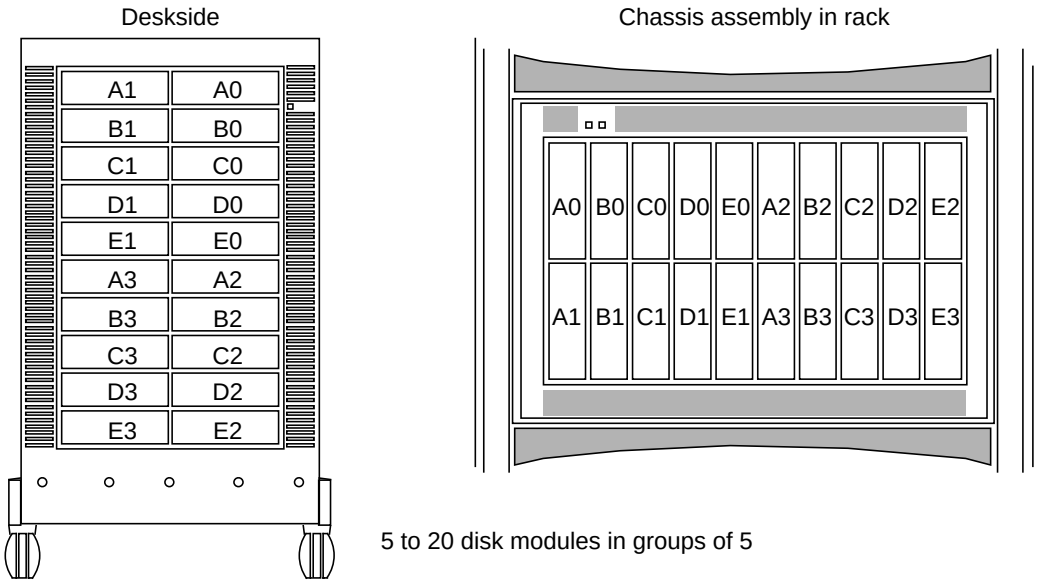

Figure 5-11 Disk Drive Locations

8. Engage the disk module's rail in the chassis rail slot, as shown in Figure 5-12.

Caution: Disk modules are extremely sensitive to shock and vibration. Even a slight jar can severely damage them.

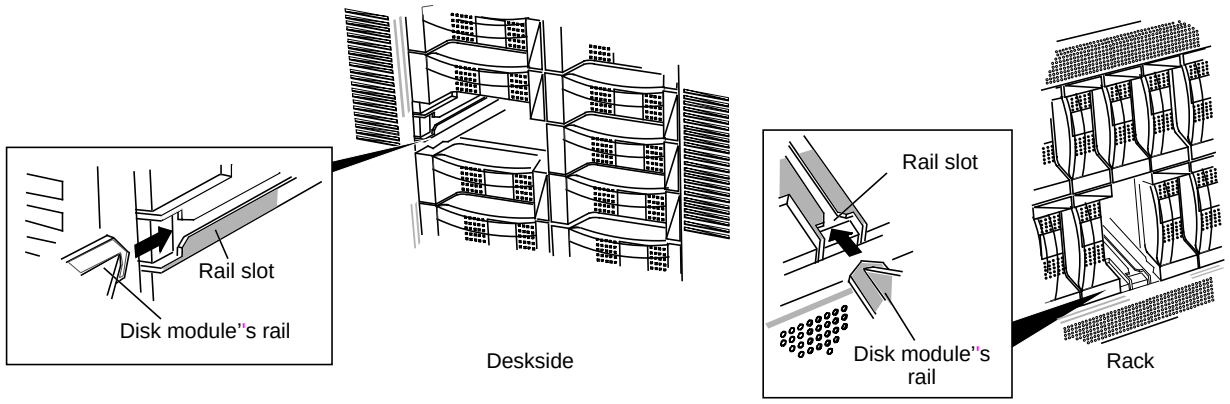

Figure 5-12 Engaging the Disk Module Rail

9. Engage the disk module's guide in the chassis guide slot, as shown in Figure 5-13.

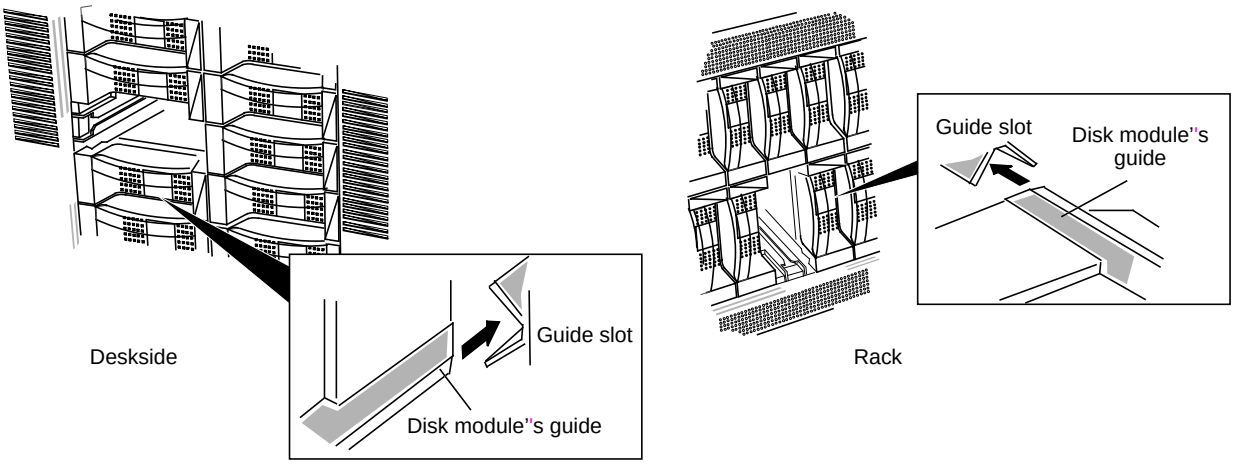

Figure 5-13 Engaging the Disk Module Guide

10. Insert the disk module, as shown in Figure 5-13. Make sure it is completely seated in the slot.
11. Repeat steps 4 through 10 until all add-on modules are installed.
12. When you are finished installing add-on modules, remove and store the ESD wrist band, if you are using one.

5.5.2 Creating Device Nodes and Binding the Disks

If you are adding disk arrays to a storage system that already has at least one LUN configured, the SPs must be made aware of the new disks. This section explains how to accomplish this without rebooting. Also, in a system with two SPs which are used for primary and secondary paths, both SPs must be made aware of the new disks. Also, the new disks must be bound into LUNs.

Follow these steps:

1. Change to the */dev* directory:

   `cd /dev`

2. Type

   `./MAKE_VLUNS` *controller-number target-number*

   This command creates the device nodes for the new disks.

3. Bind the newly installed modules into one or more physical disk units, as described in Section 4.9, "Binding Disks Into RAID Units" in Chapter 4 in this guide.

Chapter 6

Maintaining and Upgrading SPs

This chapter explains

- opening and closing the fan module
- locating the SPs
- replacing or adding an SP
- installing or removing SP memory modules (SIMMs)
- upgrading FLARE 7.xx to FLARE 8.xx for Sauna
- upgrading to FLARE 9.xx in a storage system with RAID-3 LUNs

Before you follow procedures explained in this chapter, note the following points:

Caution: To prevent thermal shutdown of the storage system, never operate it for more than two minutes with the fan module open.

- Do not operate the system with an open SP slot.
- SCSI terminators must be in place at all times during operation. Removing them causes the storage system to crash.
- Read “Handling FRUs” in Chapter 5.

6.1 Opening and Closing the Fan Module

To remove or install an SP, power supply, or battery backup unit, you must unlock and swing open the fan module:

1. On the back of the storage system, move the fan module's latch to the **UNLOCK** position, as shown in Figure 6-1.

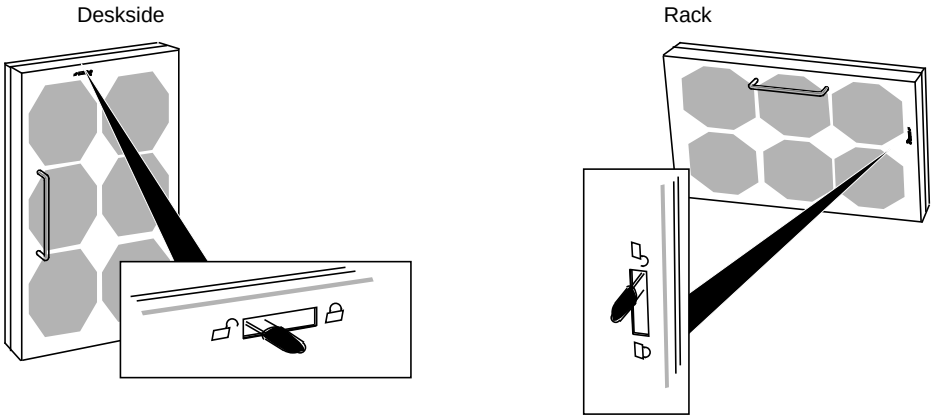

Figure 6-1 Unlocking the Fan Module

2. Swing open the fan module, as shown in Figure 6-2.

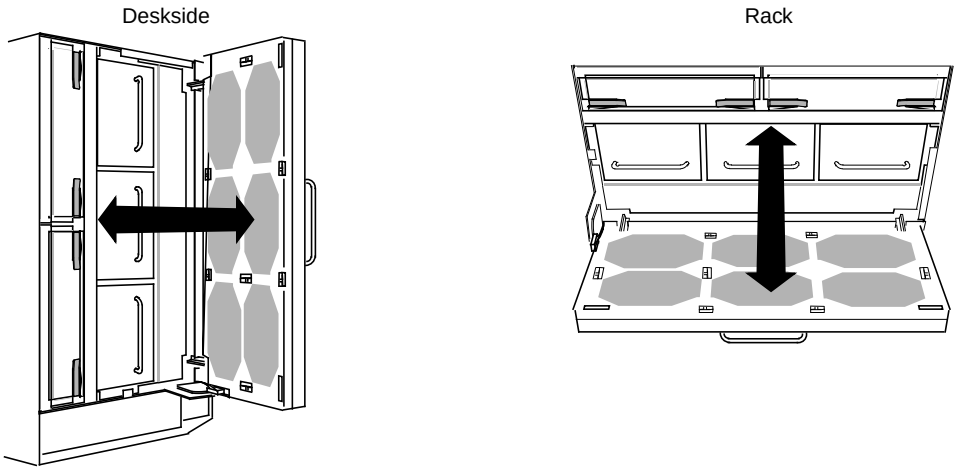

Figure 6-2 Opening the Fan Module

Close the fan module by reversing these steps.

Caution: To prevent thermal shutdown of the storage system, never operate the storage system for more than two minutes with the fan module in the open position.

6.2 Locating the SPs

Figure 6-3 shows the location of SPs in the deskside system.

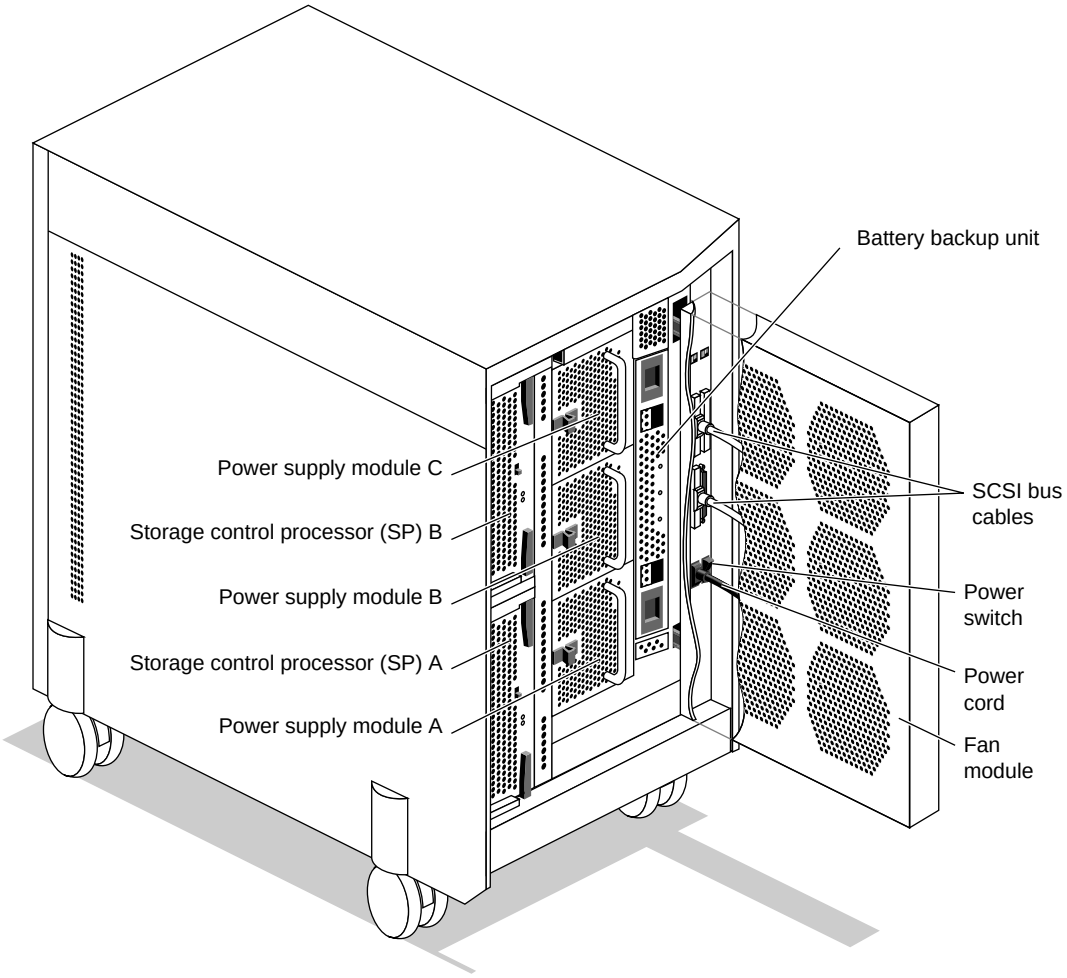

Figure 6-3 Field-Replaceable Units (FRUs): Back of Challenge RAID Deskside System

Figure 6-4 shows the SPs in a RAID chassis assembly of a Challenge RAID rack storage system.

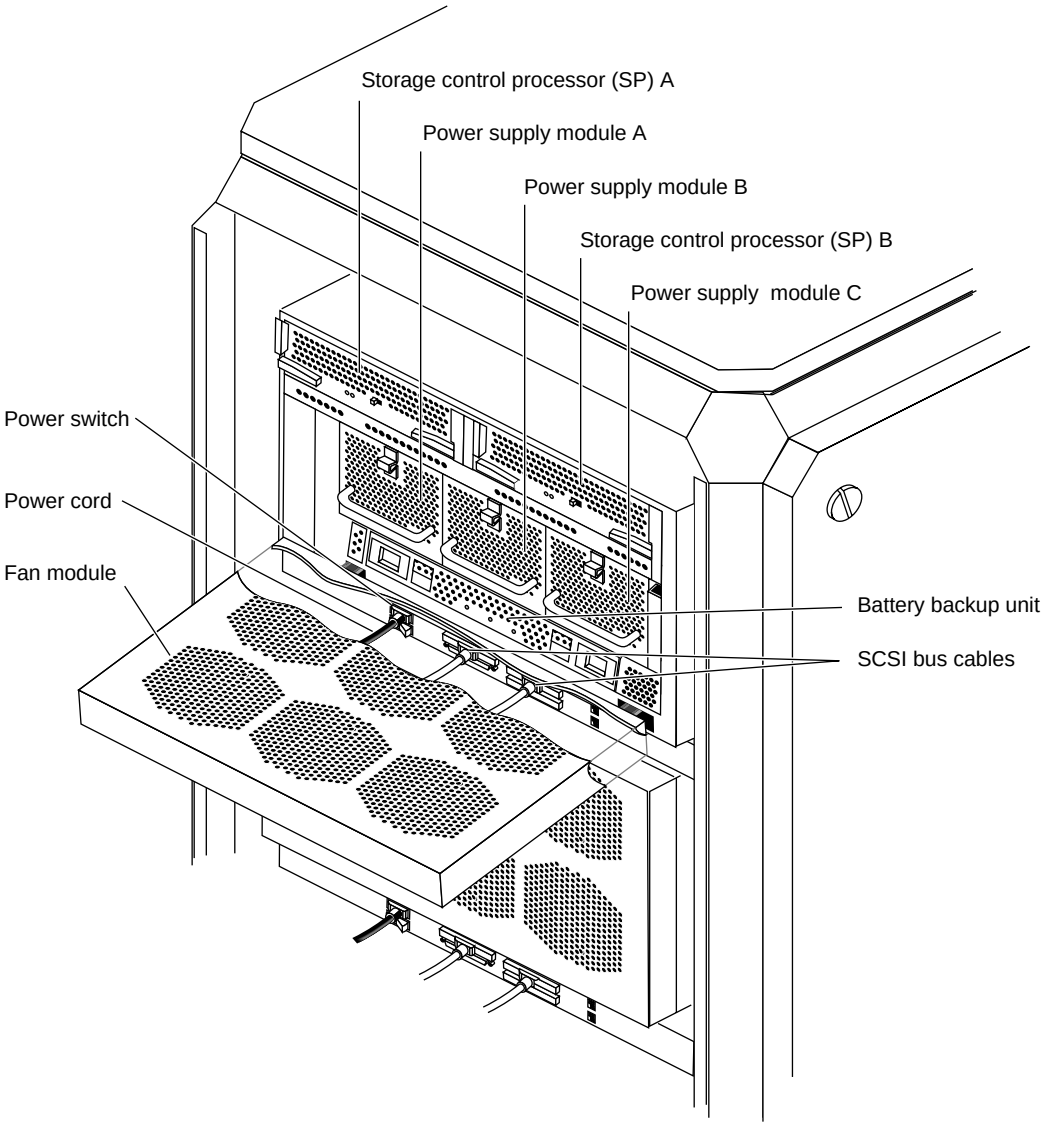

Figure 6-4 Field-Replaceable Units: Rack System

6.3 Replacing or Adding an SP

If the storage system has *one* SP, you cannot replace a failed SP without powering down the storage system You must reboot the operating system software once you install the replacement SP.

Note: If an SP fails, caching is disabled until the failed SP is replaced.

If the storage system has two SPs, you can replace a failed SP without powering down the storage system.

If necessary, use the auto-assign capability to transfer control of the physical disk units from the failed SP to the working SP, as described in Section 6.3.1, "Using the Auto-Assign Capability in the Case of a Failed SP." If the storage-system cache is enabled, disable it as described in Section 4.12.1, "Setting Cache Parameters," in Chapter 4.

Note: Many systems (such as IRIS FailSafe systems) have failover software, which make using the auto-assign capability unnecessary. The recommendation is to leave auto-assign disabled. Consult with the customer before changing LUN ownership.

If you are installing a second SP in a system that has only one, you must remove the storage-control SP filler board in the slot for the second SP board.

This section explains

- using the auto-assign capability in the case of a failed SP
- removing an SP
- installing an SP

6.3.1 Using the Auto-Assign Capability in the Case of a Failed SP

If one of two SPs becomes physically unavailable to the system, that is, if an SP fails or becomes disconnected, the Challenge RAID storage system can automatically reassign the LUNs it controls to the remaining SP.

Auto-assign controls the ownership of the LUN when an SP fails in a storage system with two SPs. With auto-assign enabled, if the SP that owns a LUN fails and the server tries to access that LUN through the second SP, the second SP assumes ownership of the LUN so the access can occur. The second SP continues to own the LUN until the SP's power is cycled (turned off and on again). When the power is cycled, ownership of each LUN returns to its default SP. If auto-assign is disabled in the previous situation, the other SP does not assume ownership of the LUN, so the access to the LUN does not occur.

Note: Consult with the customer before changing LUN ownership. Many systems (such as IRIS FailSafe systems) have failover software, which make using the auto-assign capability unnecessary.

The process consists of

- using RAIDGUI to enable auto-assign
- powering off the failed SP

6.3.1.1 Using RAIDGUI to Enable Auto-Assign

In the command-line interface, auto-assign is disabled by default; in RAIDGUI, it is enabled by default. Normally, disabling auto-assign is recommended for most systems, particularly those with failover software.

To enable auto-assign to reassign LUNs, follow these steps:

1. To start RAIDGUI, enter the following on the host to which the Challenge RAID storage system is cabled:

   ```
   /usr/raid5/raid5gui
   ```

2. After introductory screens, the Select Hosts window appears, Select a host, or enter one by clicking *Add* and entering the name or IP address of the host you want to add.

   Note: For more information, see Chapters 3 and 4 of the *CHALLENGE RAID Owner's Guide*, revision 007-2532-007.

3. Click *Select* at the bottom of the window. If RAIDGUI was able to communicate with the agent on the server, the Select Chassis window appears. Select the name of the chassis connection that you want to manage and click *Select*.

   The Select Chassis and Select Hosts windows close, and the Equipment View for the storage system appears. This window depicts the storage system; Figure 6-5 shows an example.

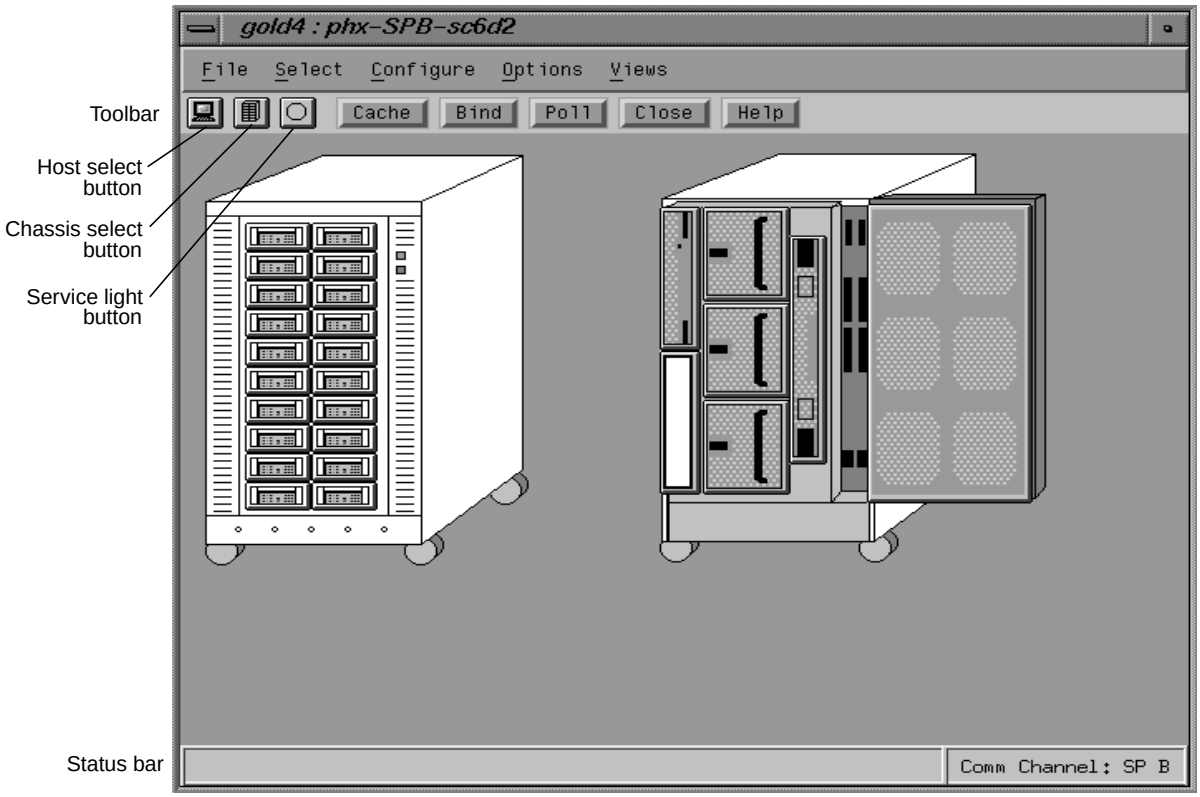

Figure 6-5 Equipment View

4. Click the *Bind* button. Figure 6-6 shows an example.

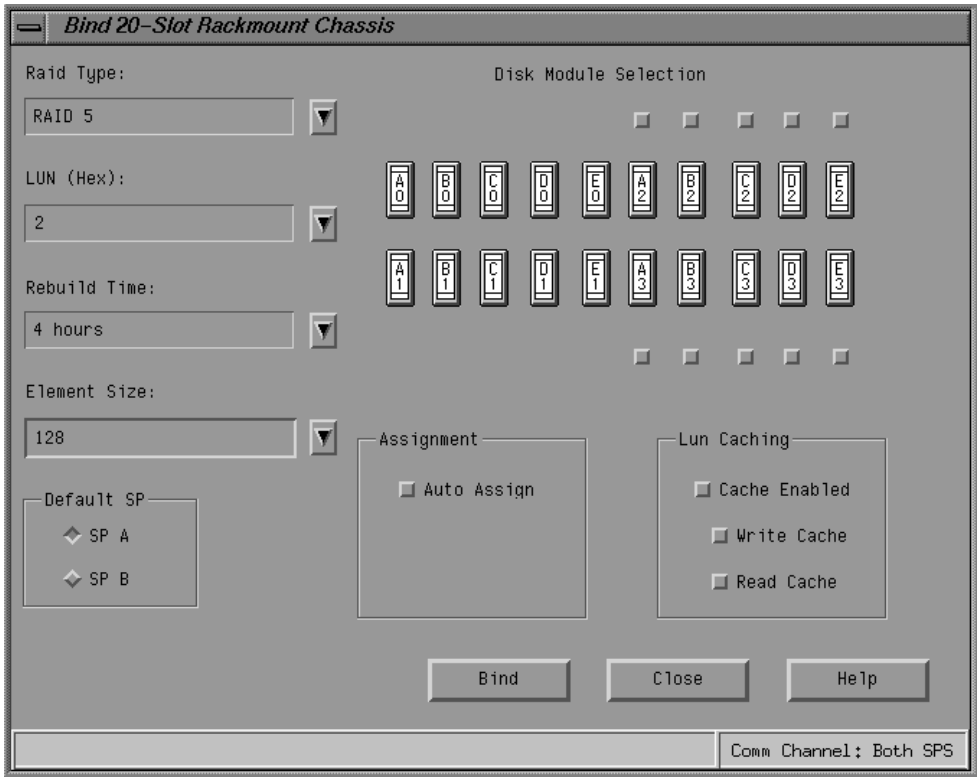

Figure 6-6 Summary View

5. Select the SP that is to take over LUN ownership (the surviving SP).
6. If the Auto Assign select button in the Assignment field is not checked, select it to enable auto-assign.
7. Exit RAIDGUI.

6.3.1.2 Powering Off the SP

Note the following points:

- Auto-assign capability cannot be used in the case of problems caused by a defective Challenge to Challenge RAID connection, such as a damaged cable.
- To force auto-assign, power off the failed SP. You do not need to power off the Challenge RAID storage system. To power off a Sauna SP, use its switch; to power off a Phoenix SP, remove it entirely from the chassis.
- Powering off an SP requires opening the fan module at the back of the storage system. Because the fan module must not be left in the open position for more than two minutes, make sure you can complete this process within that time limit.

Caution: To prevent thermal shutdown of the storage system, never operate the storage system for more than two minutes with the fan module in the open position.

Follow these steps:

1. Make sure a static-free work surface is available near the Challenge RAID storage system for the SP, if it is Power-PC based.
2. On the back of the Challenge RAID storage system, move the fan module's latch to the UNLOCK position, as shown in Figure 6-1. Swing open the fan module, as shown in Figure 6-2.
3. For a failed Sauna SP, disable it by sliding the switch to the Disable position, as shown in Figure 6-7.

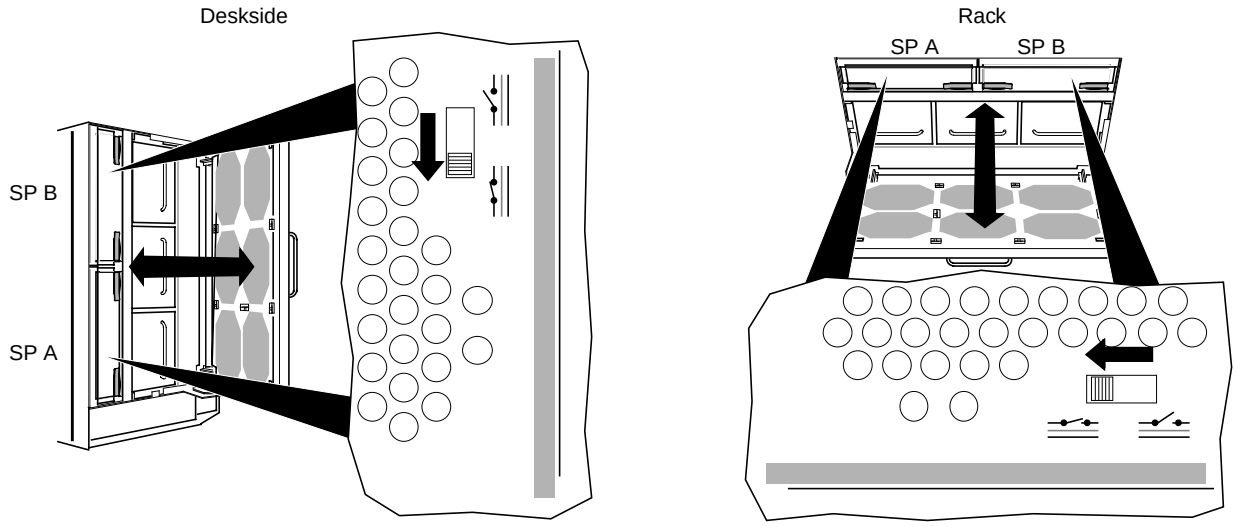

Figure 6-7 Disabling a Sauna SP

The Phoenix SP does not have power switches; disable it by removing it from the chassis, as explained in Section 6.3.2, "Removing an SP."

4. Immediately close and lock the fan module.

6.3.2 Removing an SP

In a system with two SPs, you can replace one while the storage system is powered on. However, to avoid automatic thermal shutdown of the storage system, you must complete Steps 1 and 2 within two minutes and Steps 4 through 8 within two minutes.

If you are installing a second SP in a system that has only one, follow the procedure in this section to remove the SP filler board in the slot for the second SP board.

Before you begin to remove an SP or filler board, make sure that you understand all the steps in the following procedure.

Note: If you must replace a failed SP and you do not have a replacement for it, leave the failed SP installed until the replacement is available. Do not operate the system with an open SP slot.

To remove an SP or filler board, follow these steps:

1. Unlock and open the fan module, as shown in Figure 6-1 and Figure 6-2.

Warning: **Some components on the SP become hot during operation. To prevent personal injury, disable the SP's power and do not attempt to remove the SP until you have waited at least five minutes for these components to cool.**

2. If the SP that you want to remove is a Sauna, disable its power by sliding the switch to the Disable position, as shown in Figure 6-7. Immediately close and lock the fan module to let the SP cool.

   If the SP is a Phoenix, removing it from the chassis disables its power.

3. Read "Avoiding Electrostatic Discharge (ESD) Damage" in Chapter 5. Put the ESD wrist band (strap) around your wrist with the metal button against your skin.

4. After waiting five minutes, reopen the fan module, and attach the wrist band's clip to the ESD bracket on the bottom of the storage system, as shown in Figure 5-3 and Figure 5-4 in Chapter 5.

5. Pull both ejectors until they unlock and the SP or filler board starts coming out of chassis, as shown in Figure 6-8.

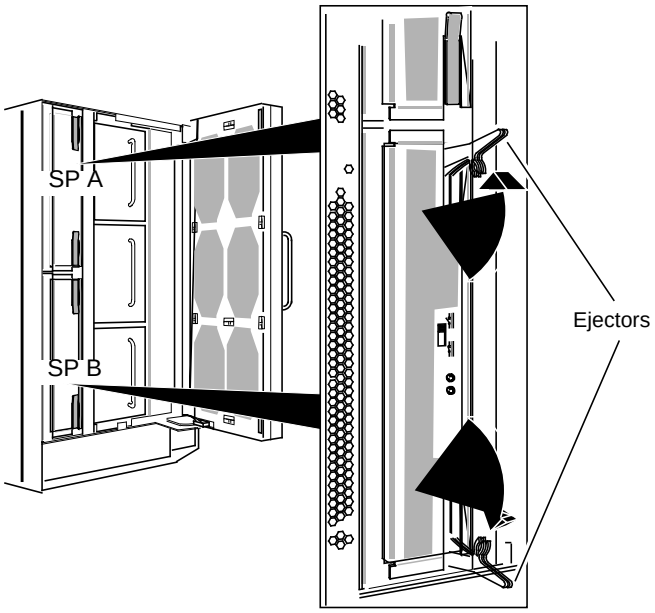

Figure 6-8 Ejectors for SP

6. Rest the edge of the SP or filler board on the palm of your hand. With your other hand, pull it out of the chassis, making sure not to any touch any electronic components or circuits.

7. Place the SP component-side up on a static-free work surface.
8. Immediately close and lock the fan module.

Note: Before shipping a failed SP, remove its memory modules as described in Section 6.4.2, “Removing SP Memory Modules.” Write down the part numbers that are on the two part number labels on the SP board.

6.3.3 Installing an SP

You can install an SP while the storage system is powered on. However, you must complete steps 5 through 9 within two minutes to avoid automatic thermal shutdown of the storage system. Before you begin to install an SP, make sure that you understand all the steps in the following procedure.

Caution: For caching, both SPs must have the same amount of cache memory in order for caching to be preserved in the event of shutdown or other power loss.

To install an SP, follow these steps:

1. Install the memory modules that shipped with the SP (see Section 6.4.1, “Installing SP Memory Modules”), and position the SP within easy reach of the storage system.
2. If you are installing an add-on SP, power off the Challenge RAID and the Challenge server.
3. If the SP is a Sauna, make sure its power switch is set to the disabled position.
4. Unlock and open the fan module as shown in Figure 6-1 and Figure 6-2.
5. If you are installing an add-on SP, remove the SP filler board from the SP slot, as described in Section 6.3.2, “Removing an SP.”

Caution: To maintain proper cooling, never operate the storage system with an empty SP slot. Save the filler board to fill the slot in case you need to operate the storage system with an SP removed temporarily.

6. Engage the edges of the SP in the board guides and slide the SP into chassis until the ejectors engage with the chassis edge, as shown in Figure 6-9.

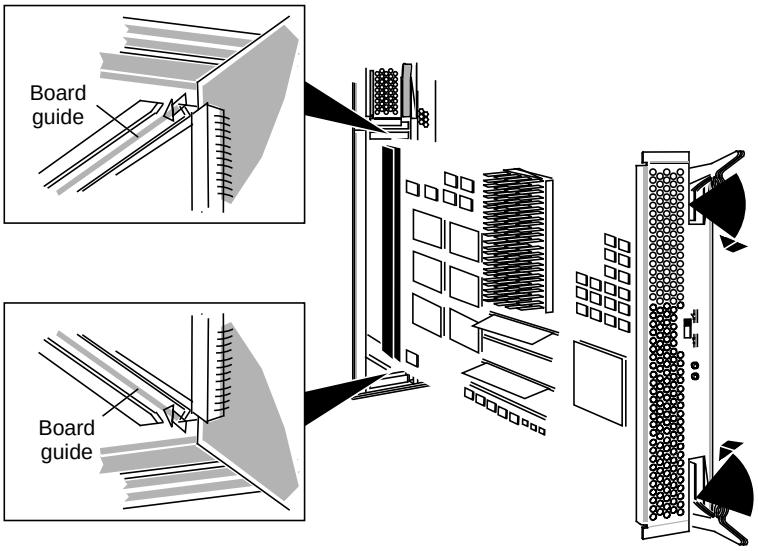

Figure 6-9 Installing an SP

7. Press both ejectors at the same time until they lock and the SP moves further into the chassis.
8. If the SP is a Sauna, enable its power by sliding the switch to the Enable position.
9. Immediately close and lock the fan module as shown in Figure 6-1 and Figure 6-2.
10. Remove and store the ESD wrist band.

If you installed a replacement SP, what you do next depends whether the storage system contains one or two SPs, and whether you powered off the storage system and host.

- If the storage system has one SP, reboot IRIX on the Challenge server.
- If you powered off the Challenge RAID and Challenge server, power them on and reboot IRIX.
- If you installed an add-on SP, connect the SCSI cable and SCSI terminator plug, as described in Chapter 3 in this guide.
- If you upgraded an SP with cache memory, set up storage-system disk caching as described in Chapter 4 in this guide.

  Note: For caching, both SPs must have the same amount of cache memory in order for caching to be preserved in the event of shutdown or other power loss.

6.4 Installing or Removing SP Memory Modules (SIMMs)

When you replace or add an SP, you must install the memory modules in the SIMM connectors on the replacement or add-on SP's printed-circuit board.

Silicon Graphics supports only 8 MB and 64 MB memory configurations:

- in the 8 MB configuration, two SIMMs of 4 MB each are installed on each SP PC board
- in the 64 MB configuration, four SIMMs of 16 MB each are installed on each SP PC board

Caution: For caching, both SPs must have the same amount of cache memory in order for caching to be preserved in the event of shutdown or other power loss. In storage systems with two SPs, both SPs must be of the same type (both Sauna or both Phoenix).

If the Challenge RAID has one SP, you can install a second one while the storage system is running and configure it into the RAID array system. When both SPs are installed and configured, you can replace either SP while the storage system is running.

To upgrade an existing 8 MB memory configuration to a 64 MB configuration (required for caching), remove the existing SIMMs on the existing SP board, populate it fully with four 16 MB SIMMs, and reinstall the board; insert four 16 MB SIMMs onto the new SP board and install the board.

Note: You may find two 1 MB SIMMs on the Sauna SP board instead of two 4 MB SIMMs.

The SP must have one or two pairs of memory modules, and the modules in each pair must have the same capacity. If the SP has only one pair of modules, they must be in the SIMM1 and SIMM3 connectors on the Sauna SP printed-circuit board, as shown in Figure 6-10.

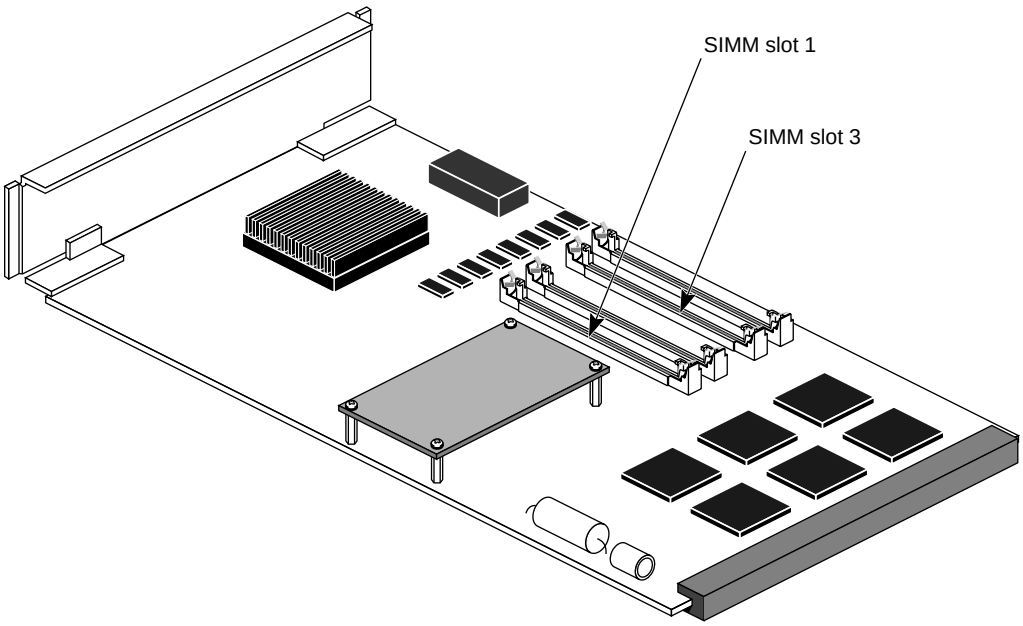

Figure 6-10 SIMM 1 and 3 Connectors on the Sauna SP Board

On the Phoenix board, the modules are in connectors 1 and 2 (numbered from the bottom up), as shown in Figure 6-11.

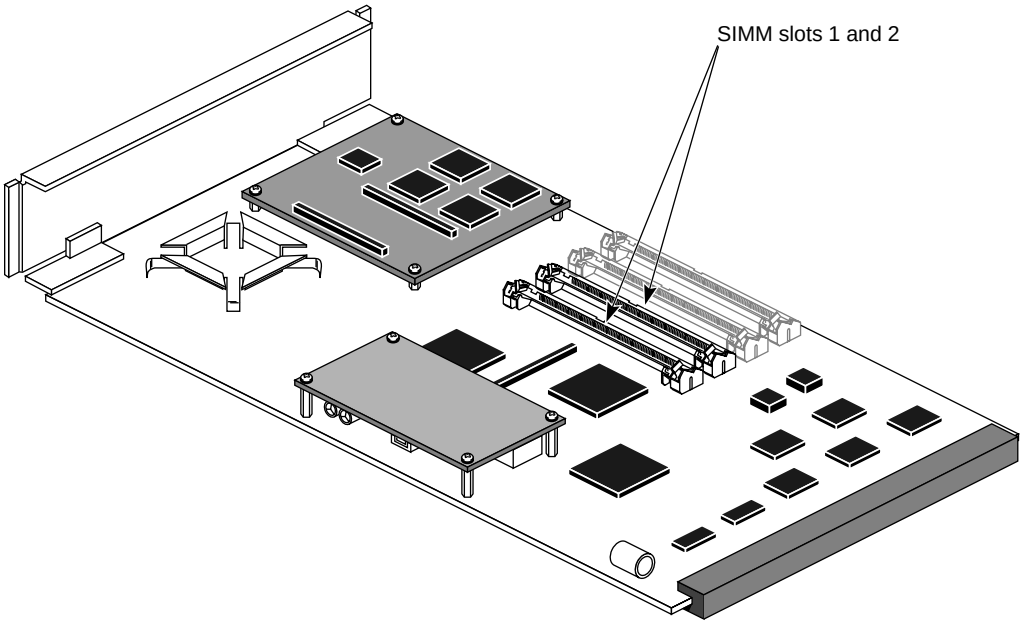

Figure 6-11 SIMM 1 and 2 Connectors on the Phoenix SP Board

6.4.1 Installing SP Memory Modules

To install SP memory modules, follow these steps:

1. Run *raid5 getcache* to determine if caching is enabled. If it is enabled, make sure both SPs are powered on; check the status lights or run *raid5 getcrus* (see Appendix B for information).
2. If caching is enabled, disable it by entering

   `raid5 -d` *device* **`setcache 0`**
3. To make sure that cache state is set to enabled, enter

   `raid5 -d` *device* **`getcache`**

   Look for `Cache State: Enabled` in the output of this command.
4. Read Section 5.1.1, "Avoiding Electrostatic Discharge (ESD) Damage," in Chapter 5.

Warning: **Some components on the SP become hot during operation. To prevent personal injury, disable the SP's power and do not attempt to remove the SP until you have waited at least five minutes for these components to cool.**

5. If you are installing memory on both SPs in a system, power off the Challenge RAID storage system.
6. Disable and remove the SP:
   - To take a Sauna SP offline, move the power switch to the disable position, as shown in Figure 6-7**.** Immediately close the fan module. Wait five minutes for the SP to cool in the chassis before removing it and placing it on a static-free work surface.

     Attach the clip of the ESD wrist band (strap) to the ESD bracket on the bottom of the storage system (see Figure 5-3 and Figure 5-4 in Chapter 5), and put the wrist band around your wrist with the metal button against your skin. Pull both ejectors until they unlock and the SP or filler board starts coming out of chassis, as shown in Figure 6-8.
   - To take a Phoenix SP offline, remove it from the chassis and place it on a static-free work surface. Wait for it to cool before handing it further.

   Note: You do not have to power off the system if you are installing a second SP in a system that currently has only one, or if you are replacing only one SP in a system that currently has two configured.
7. Remove the memory modules (and, if the customer has ordered it, the SP board) from their packaging, and place them on a static-free work surface.
8. If you are upgrading 4 MB SIMMs (or 1 MB SIMMs) to 16 MB SIMMs, remove them following instructions in Section 6.4.2, "Removing SP Memory Modules."

9. Starting with the memory module in the lowest numbered connector (see Figure 6-10), insert the memory modules:

   - For Sauna, position the memory module in the appropriate SIMM connector at a 20- to 30-degree angle, as shown in Figure 6-12.

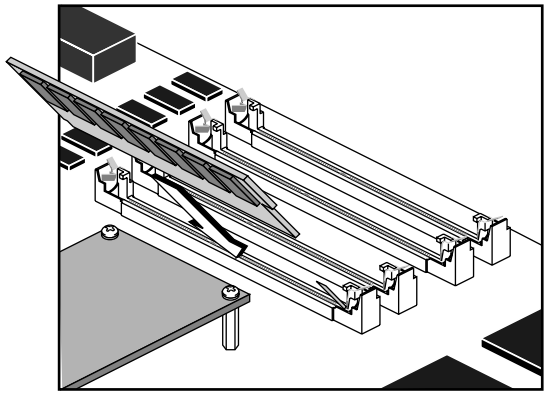

Figure 6-12 Inserting an SP Memory Module: Sauna

Gently push the top of the module up until it locks into the Sauna SIMM connector, as shown in Figure 6-13.

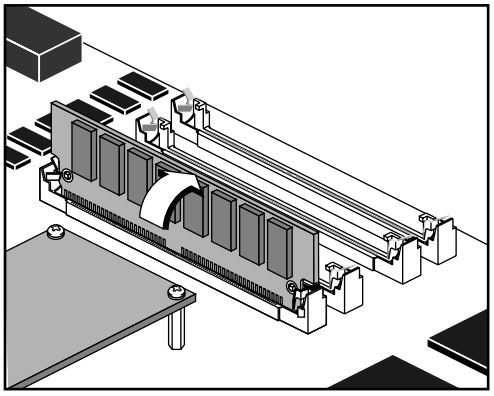

Figure 6-13 Locking In an SP Memory Module: Sauna

- For Phoenix, position the memory module in the appropriate SIMM connector straight down, as shown in Figure 6-14.

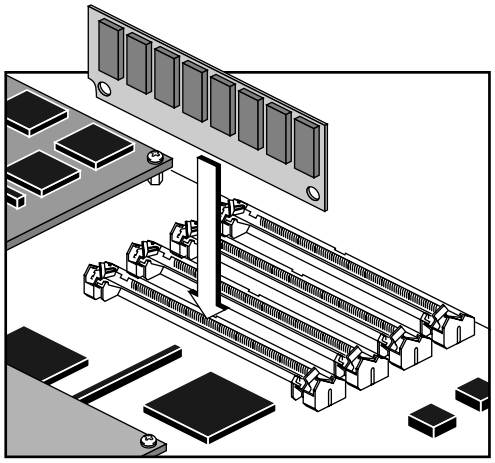

Figure 6-14 Installing an SP Memory Module: Phoenix

Gently push the module to a 20- to 30-degree angle until it locks into the connector, as shown in Figure 6-15.

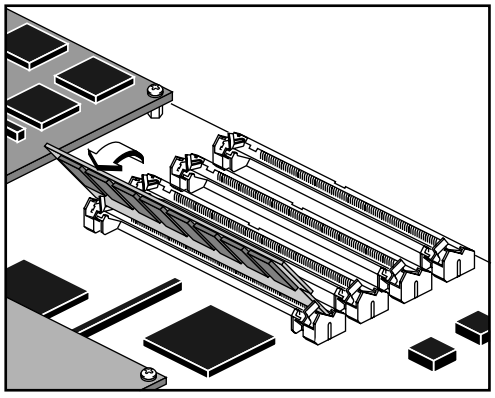

Figure 6-15 Locking In an SP Memory Module: Phoenix

10. Install the remainder of the memory modules.

    Caution: Do not force the memory modules into the connectors or against the locking tabs. A notch on one end of the module prevents you from inserting it the wrong way.

11. Reinstall the SP board into the SP. If the SP is a Sauna, enable it by pushing the switch to the on position.
12. If you have powered off the Challenge RAID storage system because you installed memory in both SPs, power it on.
13. Enable caching in both SPs: see Section 4.12, “Setting Up Caching,” in Chapter 4.

Reinstall the SP following instructions in Section 6.3.3, “Installing an SP.”

6.4.2 Removing SP Memory Modules

To ship an SP, remove the memory modules from the printed-circuit board as follows:

1. Read Section 5.1.1, "Avoiding Electrostatic Discharge (ESD) Damage," in Chapter 5.
2. Attach the clip of the ESD wrist band (strap) to the ESD bracket on the bottom of the storage system (as shown in Figure 5-3 and Figure 5-4 in Chapter 5), and put the wrist band around your wrist with the metal button against your skin.
3. With your thumbs, push out on the memory module's locking tabs. The tabs are indicated in Figure 6-16.

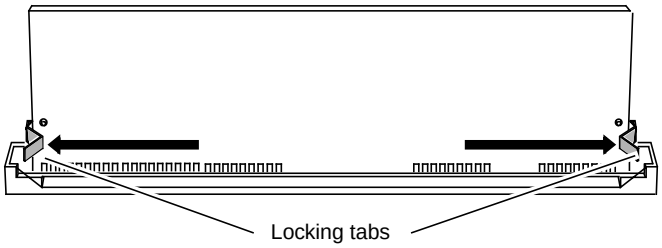

Figure 6-16 Removing a Memory Module From the SIMM Connector

4. With your fingers, push the module forward until it is released from the tabs. Lift the module out of the connector and place it in its own antistatic package.
5. Remove and store the ESD wrist band.
6. If necessary, reboot the operating system software, or use the auto-reassign capability, as explained in Section 6.3.1, "Using the Auto-Assign Capability in the Case of a Failed SP."

Note: Do not ship an SP with its memory modules installed. To return a failed or replaced SP, remove the memory modules from the printed-circuit board and place the board and memory modules in shipping containers.

6.5 Upgrading FLARE 7.xx to FLARE 8.xx for Sauna

Caution: The storage system is automatically rebooted after the code is downloaded to the SP(s). Do not follow the steps below until you have unmounted any filesystems or partitions on the storage system.

Upgrade scripts are included on the Challenge RAID software CD and are installed along with the software. To upgrade the FLARE code for a Sauna SP, run

```
/usr/raid5/sauna7_to_sauna8
```

This script upgrades the FLARE code on the Sauna SP(s) in the system from 7.xx to 8.xx after first disabling caching, if it is enabled. It also loads code to all database disks (A0, B0, C0, and A3).

Note: When you upgrade a Challenge RAID storage system to FLARE code 8.xx, all prefetch parameters are set to 0, except fort the retention parameter, which is set to equal priority for each existing LUN. You must change these values if you want prefetching to occur for a LUN.

6.6 Upgrading Sauna to Phoenix

This section explains how to upgrade the board in an SP from Sauna to Phoenix and how to upgrade the FLARE code to version 8.50 through 8.99, as required by the Phoenix SP. The process consists of

- upgrading and testing the FLARE code
- moving Sauna SIMMs to the Phoenix SP board
- testing the upgrade

Note: For upgrading from Sauna to Phoenix, it is highly recommended that you use an ASCII terminal or PC connected to the Challenge RAID serial port. When the Phoenix firmware comes up, it shows information for the database drives (A0, B0, C0, and A3).

6.6.1 Upgrading the FLARE Code

Caution: The storage system is automatically rebooted after the code is downloaded to the SP(s). Do not follow the steps below until you have unmounted any filesystems or partitions on the storage system.

Upgrade scripts are included on the Challenge RAID software CD and are installed along with the software. To change the FLARE code when upgrading from Sauna to Phoenix, follow these steps:

1. Determine the FLARE code currently on the SPs:

   ```
   /usr/raid5/raidcli getagent
   ```

   Note: For full information on this command, see Section B.6, “getagent,” in Appendix B.

2. If the FLARE code is 7.xx, run the following command to upgrade it to FLARE 8.xx for Sauna (8.00 through 8.49):

   `/usr/raid5/sauna7_to_sauna8`

3. If the FLARE code is 8.xx (or after you complete step 2), run the following command to upgrade it to FLARE 8.xx for Phoenix (8.50 through 8.99):

   `/usr/raid5/sauna8_to_phoenix8`

   This script disables caching, if it is enabled, and then attempts to upgrade the FLARE code from Sauna 8.xx to Phoenix 8.xx on a database disk, in this order: B0, C0, or A3). The script does not update the FLARE code on disk A0, so that it has usable code in case the upgrade fails. Note that this upgrade script works differently from that for the Sauna 7 -to-Sauna 8 upgrade.

4. Wait five to ten minutes. Halt the system, and then restart the RAID agent:

   ```
   /etc/init.d/raid5 stop
   /etc/init.d/raid5 start
   ```

5. Test the upgrade by using *raidcli getagent* to make sure the communication path still works and the device to which the code was downloaded can still be seen.

6.6.2 Moving Sauna SIMMs to the Phoenix SP Board

Follow these steps to install the SIMMs from the existing Sauna board onto the new replacement Phoenix board:

1. After the FLARE code is upgraded and tested as successful, stop the RAID agent:

   ```
   /etc/init.d/raid5 stop
   ```

2. Power off the Challenge RAID storage system.
3. Remove the system's Sauna SP(s) following instructions in Section 6.3.2, "Removing an SP," earlier in this chapter.
4. Remove the Phoenix SP board and, if the customer has ordered them, the additional memory modules from their packaging, and place them on a static-free work surface.
5. On the Sauna board you are replacing, locate the first SIMM module; see Figure 6-10 if necessary.
6. Use your thumbs to push out on the memory module's locking tabs.
7. With your fingers, push the first SIMM module forward until it is released from the tabs, which are shown in Figure 6-17.

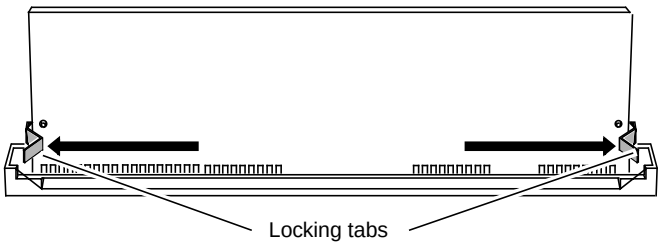

Figure 6-17 Removing a Memory Module From the SIMM Connector

8. Lift the module out of the connector.

9. Insert the SIMM into the first SIMM connector on the Phoenix board. The modules go into connectors 1 and 2 (numbered from the bottom up) on the Phoenix board, as shown in Figure 6-18.

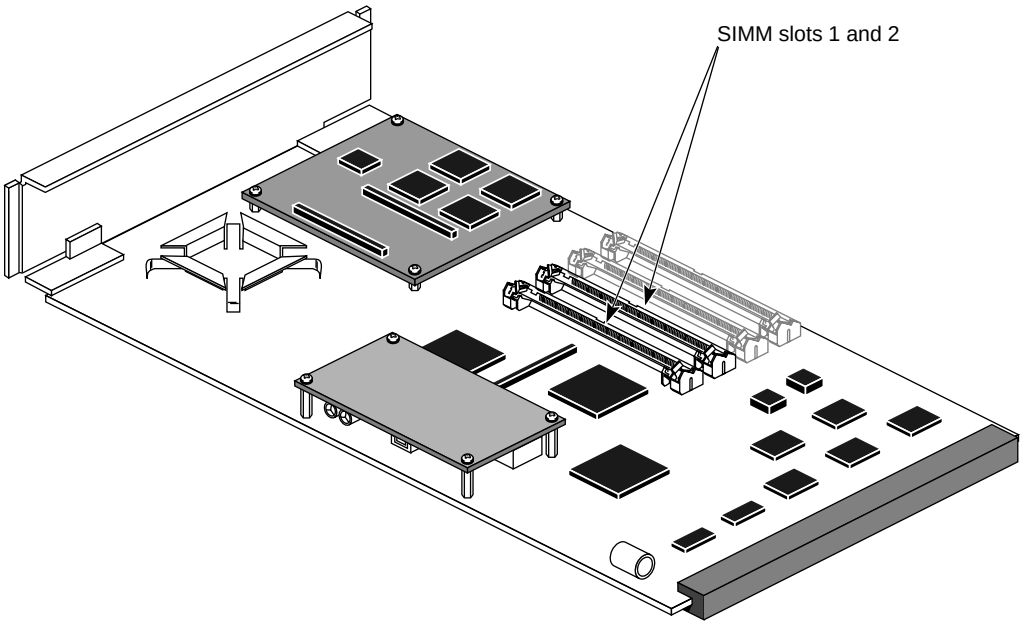

Figure 6-18 SIMM 1 and2 Connectors on the Phoenix SP Board

10. On the Phoenix board, position the Sauna's first SIMM in the first SIMM connector straight down, as shown in Figure 6-19.

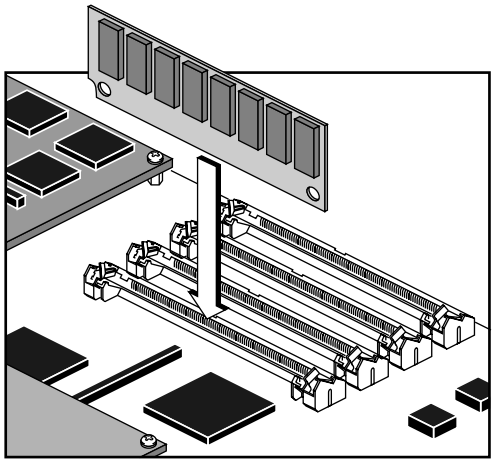

Figure 6-19 Inserting an SP Memory Module on the Phoenix SP Board

Note: A memory module has a notch on one end so that you can insert it only one way.

11. Gently push the module to a 20- to 30-degree angle until it locks into the connector, as shown in Figure 6-20.

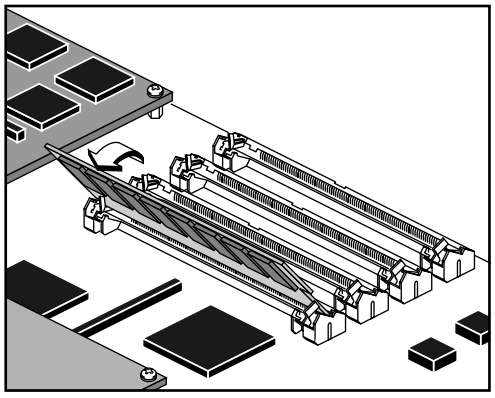

Figure 6-20 Locking the SIMM into the Phoenix SIMM Connector

12. Remove the second SIMM from connector 2 on the Sauna board and install it in connector 2 on the Phoenix board following procedures in the preceding steps.
13. Complete the process for the remainder of the memory modules.
14. Reinstall the SP into the controller.
15. Power on the Challenge RAID storage system.

Note: When you upgrade a Challenge RAID storage system to FLARE code 8.xx, all prefetch parameters are set to 0, except fort the retention parameter, which is set to equal priority for each existing LUN. You must change these values if you want prefetching to occur for a LUN.

6.6.3 Testing the Upgrade

After the Challenge RAID storage system is powered on, restart the RAID agent:

```
/etc/init.d/raid5 start
```

To check whether the revision level is correct, enter

```
/etc/usr/raid5/raidcli getagent
```

In the output of this command, look for the `Revision:` line, which should show a FLARE code number between 8.50 and 8.99.

Note: If the database drive that has the original Sauna code fails, replace the failed disk in order to complete the upgrade process.

If the test is not successful, reverse the upgrade procedure. Follow these steps:

1. If the database drive that has the Phoenix FLARE code fails, replace the failed disk in order to complete the rest of the steps in this section.
2. Power off the Challenge RAID storage system.
3. Remove the Phoenix board from the chassis.
4. Remove each SIMM and replace it to its original slot on the Sauna SP.
5. Replace the Sauna in the chassis.
6. Power on the storage system.

6.7 Upgrading to FLARE 9.xx in a Storage System With RAID-3 LUNs

If you upgrade to FLARE 9.xx in a storage system with existing RAID-3 LUNs, the RAID-3 LUNs will work, but will not be FLARE 9.xx RAID-3 LUNs. To make them FLARE 9.xx RAID-3 LUNs, follow these steps:

1. Back up the data on the RAID-3LUNs.
2. Allocate storage-system memory to RAID-3, as explained in Section 4.8, "Allocating Memory for RAID-3 LUNs," in Chapter 4.
3. Unbind the RAID-3 LUNs.
4. Rebind the RAID-3 LUNs.

   These LUNs are now FLARE 9.xx RAID-3 LUNs.
5. Reload data onto the RAID-3 LUNs.

Note: Once you have loaded FLARE 9.xx and have bound RAID-3 LUNs, do not load an earlier FLARE revision. If you do so, these FLARE 9.xx RAID-3 LUNs appear as Unbound in the output of *raidcli getdisk*.

If you then reload FLARE 9.xx, the status of these RAID-3 LUNs changes to Bound and Not Assigned (in the GUI, LUNs not owned by an SP) and the storage-system RAID-3 memory is 0 MB in size. To use these RAID-3 LUNs, you must resize the storage-system RAID-3 memory to the size equal to the total of the RAID-3 memory assigned to each RAID-3 LUN. (To determine RAID-3 memory, use *raidcli getlun*.)

Chapter 7

Maintaining and Upgrading Other Components

This chapter explains replacing field-replaceable units:

- opening and closing the fan module
- replacing or adding a power-supply module
- replacing a fan module
- replacing or adding a battery backup unit
- replacing the power distribution unit in a Challenge RAID rack
- installing an additional chassis assembly in a Challenge RAID rack

Before you follow procedures explained in this chapter, note the following points:

Caution: To prevent thermal shutdown of the storage system, never operate it for more than two minutes with the fan module open.

- Do not operate the system with an open SP slot.
- SCSI terminators must be in place at all times during operation. Removing them causes the storage system to crash.
- Read “Handling FRUs” in Chapter 5.

Note: For instructions on installing an add-on disk module array, upgrading an SP, replacing or adding an SP module, and installing an additional chassis assembly in a Challenge RAID rack, see Chapter 6, “Maintaining and Upgrading SPs.”

Figure 7-1 shows the location of FRUs. Only SSEs replace these.

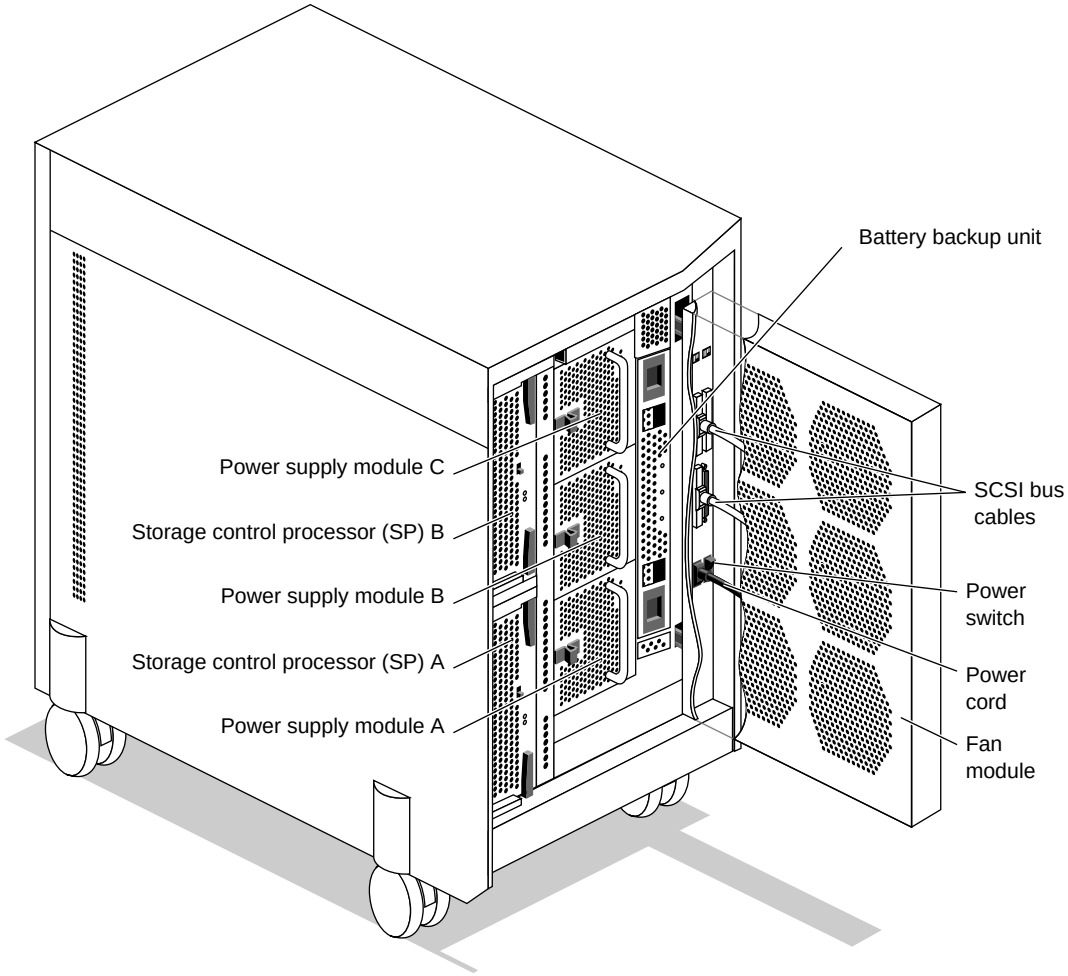

Figure 7-1 Field-Replaceable Units (FRUs): Back of Challenge RAID

Figure 7-2 shows the Challenge RAID field-replaceable units (FRUs) in one RAID chassis assembly of a Challenge RAID rack storage system. Only SSEs replace these.

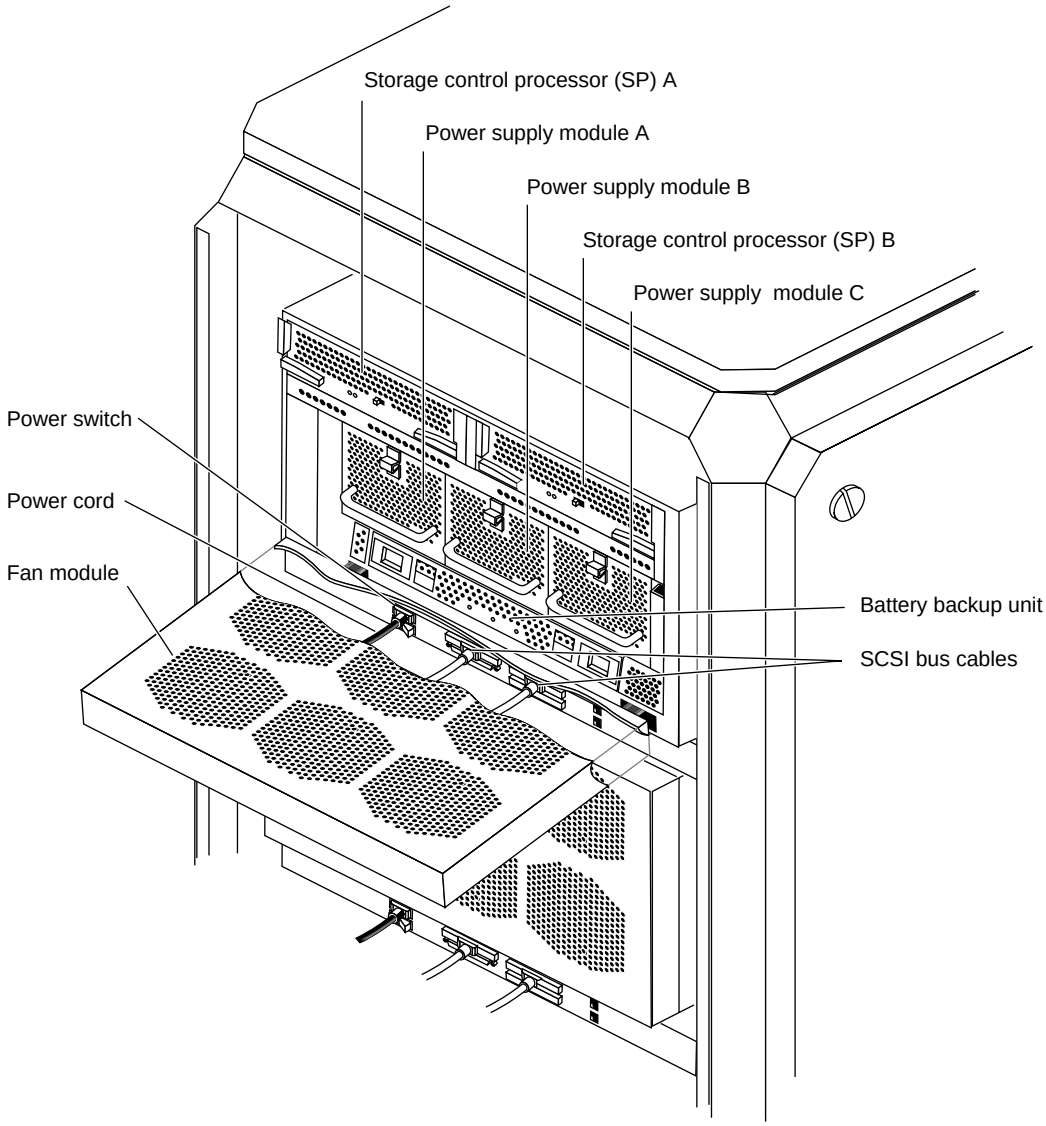

Figure 7-2 Field-Replaceable Units: Rack System

7.1 Opening and Closing the Fan Module

To remove or install an SP, power supply, or battery backup unit, you must unlock and swing open the fan module:

1. On the back of the storage system, move the fan module's latch to the **UNLOCK** position, as shown in Figure 7-3.

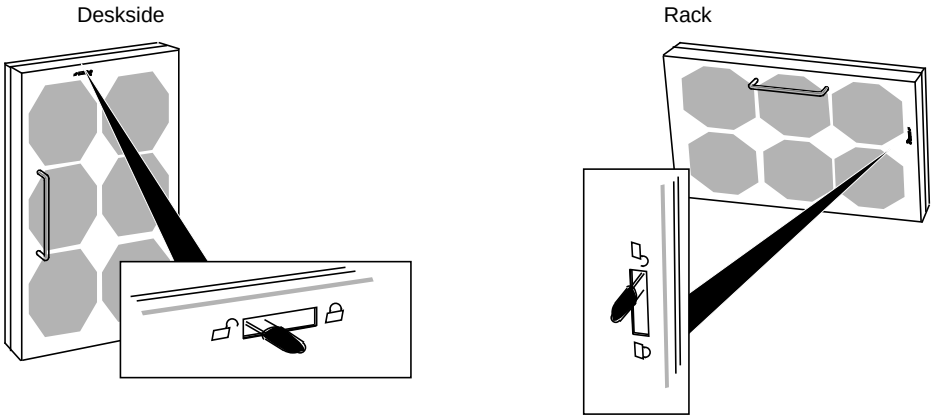

Figure 7-3 Unlocking the Fan Module

2. Swing open the fan module, as shown in Figure 7-4.

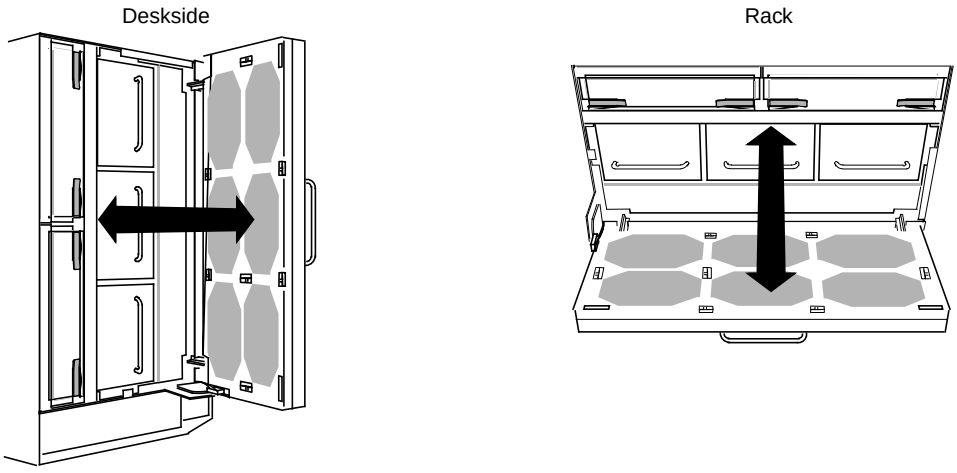

Figure 7-4 Opening the Fan Module

Close the fan module by reversing these steps.

Caution: To prevent thermal shutdown of the storage system, never operate the storage system for more than two minutes with the fan module in the open position.

7.2 Replacing or Adding a Power Supply Module

You can replace a failed power supply module (voltage semi-regulated converter, or VSC) or add a new one while the storage system is powered on. However you must complete steps 2 through 6 within two minutes to avoid automatic thermal shutdown of the storage system. Before you begin to replace or add a power supply module, make sure that you understand all the steps in these procedures.

This section explains

- replacing a power supply module
- adding a power supply module

7.2.1 Replacing a Power Supply Module

To replace or add a power supply module, follow these steps:

1. Unpack the power supply module and place it within reach of the storage system.
2. Unlock and open the fan module, as shown in Figure 7-3 and Figure 7-4.
3. If you are adding a power supply module, remove the power supply module filler module by pushing down on its latch and pulling it out of the slot by its latch, as shown in Figure 7-5.

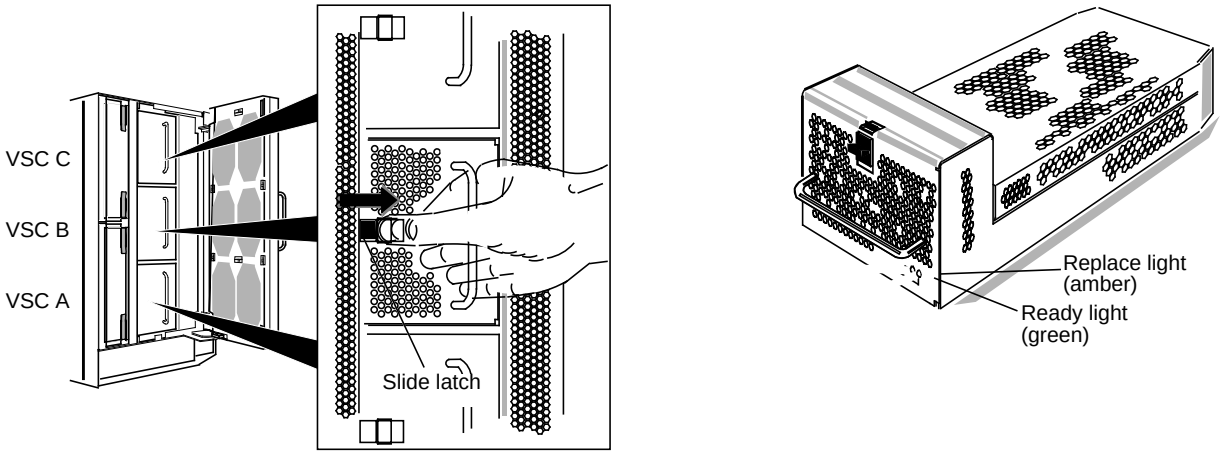

Figure 7-5 Removing a Power Supply Module (VSC)

As indicated in Figure 7-5, a power supply module filler module looks similar to a power supply module without its handle.

Caution: If you are replacing a power supply in a system with more than one installed, be sure to remove the one(s) with the replace light activated. Removing an active power supply introduces an additional fault condition that causes the Challenge RAID to shut down.

Caution: To maintain proper cooling, never operate the storage system with an empty power supply module slot. Reserve the filler module.

4. Install the replacement or add-on power supply module, as shown in Figure 7-6.

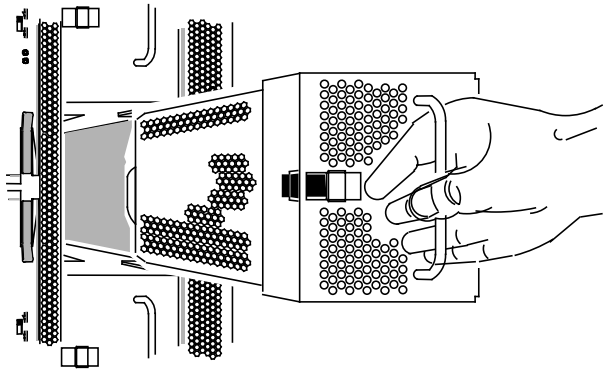

Figure 7-6 Adding a Power Supply Module

5. Immediately close and lock the fan module, as shown in Figure 7-3.
6. If the storage system is not running, power it on, as explained in Section 3.9, "Powering On Challenge RAID," at the end of Chapter 3.

7.2.2 Adding a Power Supply Module

To add a power supply module, follow these steps:

1. Unpack the power supply module and place it within reach of the storage system.
2. Unlock and open the fan module, as shown in Figure 7-3 and Figure 7-4.
3. Remove the power supply module filler module by pushing down on its latch and pulling it out of the slot by its latch, as shown in Figure 7-5.
4. Grasp the power supply module by its handle and push it into the slot until it latches into place, as shown in Figure 7-6.

7.3 Replacing a Fan Module

You can replace a failed fan module while the storage system is powered on. However, you must complete Steps 2 through 7 within two minutes to avoid thermal shutdown of the storage system. Before you remove the fan module, make sure that you understand all the steps in this procedure.

Note: To avoid thermal shutdown of the storage system, never operate it for more than two minutes with the fan module removed or in the open position.

To replace the fan module, follow these steps:

1. Unpack the replacement fan module and place it within reach of the storage system.
2. Unlock and open the failed fan module, as shown in Figure 7-3 and Figure 7-4.

3. Disconnect the fan module's power, as shown in Figure 7-7.

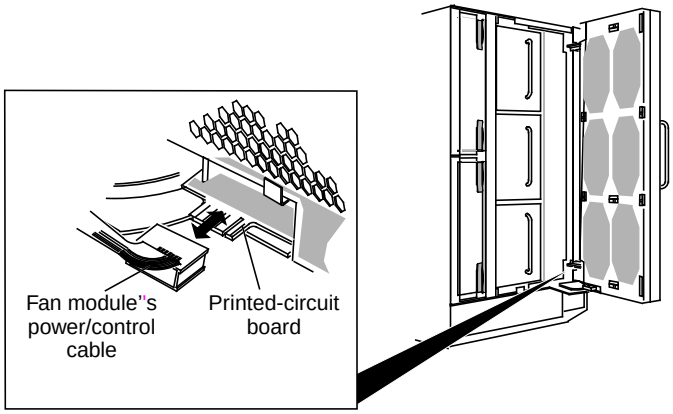

Figure 7-7 Disconnecting the Fan Module Power

4. Press the fan module's latch and guide the tab through the cutout, as shown in Figure 7-8.

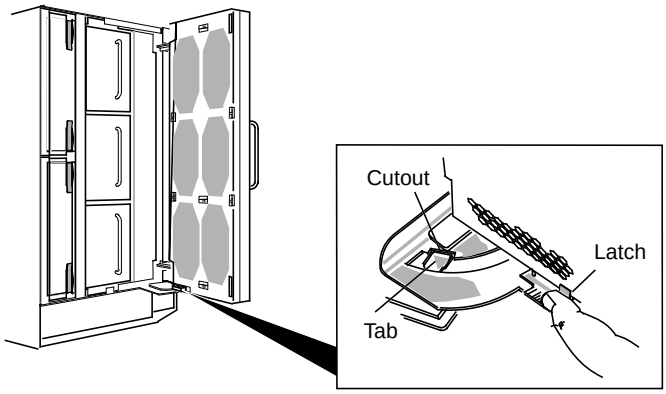

Figure 7-8 Unlatching the Fan Module

5. Lift the fan module free of the hinge pins and release the latch, as shown in Figure 7-9.

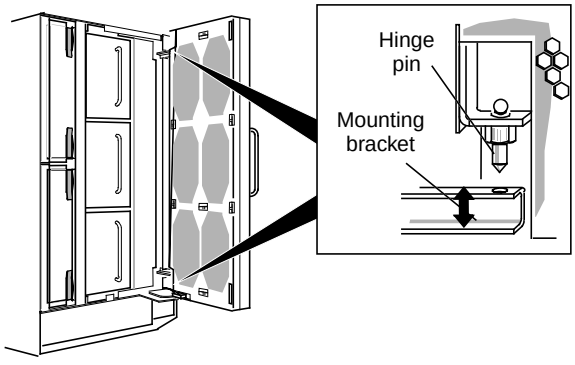

Figure 7-9 Lifting Out the Fan Module

6. To install the replacement fan module, lift it onto the hinge pins and latch it, as shown in Figure 7-9.
7. Press the replacement fan module's latch and guide the tab through the cutout, as shown in Figure 7-8.
8. Connect the fan module's power, as shown in Figure 7-7.

   All six fans begin to rotate.
9. Immediately close and lock the fan module as shown in Figure 7-3 and Figure 7-4.

7.4 Replacing or Adding a Battery Backup Unit

You can replace a failed battery backup unit or install one (the battery backup unit is optional) while the storage system is powered on. However, you must complete steps 2 through 6 within two minutes to avoid thermal shutdown of the storage system. Before you remove or add the battery backup unit, make sure that you understand all the steps in the following procedure.

1. Unpack the battery backup unit and place it within reach of the storage system.
2. Unlock and open the fan module, as shown in Figure 7-3 and Figure 7-4.
3. If you are replacing a battery backup unit, remove the failed battery backup unit, slide both latches in the direction shown in Figure 7-10, and pull the battery backup unit or filler panel from the chassis.

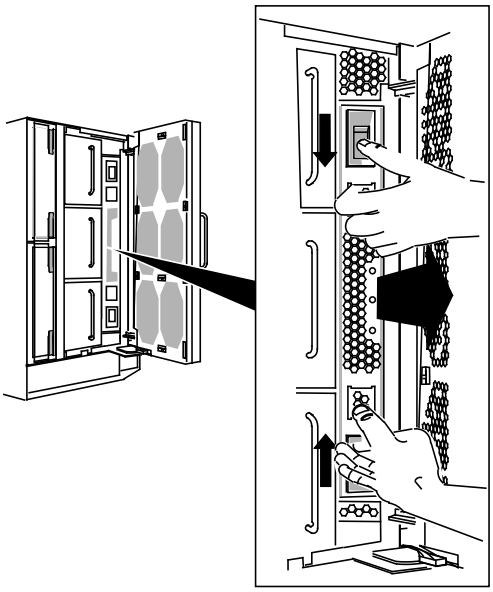

Figure 7-10 Removing a Battery Backup Unit or Filler Panel

4. If you are adding a battery backup unit, remove the filler module as shown in Figure 7-10.

5. Push the battery backup unit or filler panel into the slot until it latches into place, as shown in Figure 7-11.

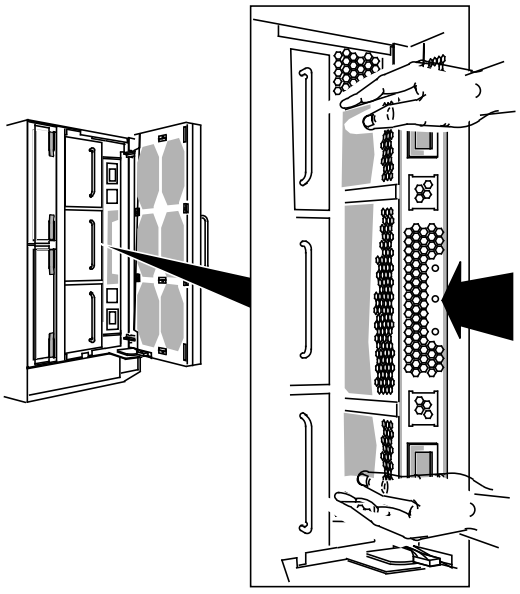

Figure 7-11 Removing or installing a Battery Backup Unit or Filler Panel

Caution: To maintain proper cooling, never operate the storage system without a battery backup unit or filler panel installed. Save the filler panel to fill the slot in case you need to operate the storage system with the battery backup unit removed temporarily.

6. Immediately close and lock the fan module, as shown in Figure 7-3 and Figure 7-4.

7.5 Replacing the Power Distribution Unit in a Challenge RAID Rack

This section explains how to swap out a failed power distribution unit from the Challenge RAID rack. This unit is attached to the base of the rack.

1. Work with the system administrator at the site to power down the rack. Unplug the power cable.
2. Remove the side panels from the Challenge RAID rack and open the rear door. It should not be necessary to remove the rear door.

   Note: Rehanging the rear door might require a second person because of the difficulty of lining up the three pinlocks.

3. Remove the six screws that hold the lowest chassis assembly in place and push the chassis assembly halfway out of the rack to clear access to the power distribution unit. See Figure 7-12.

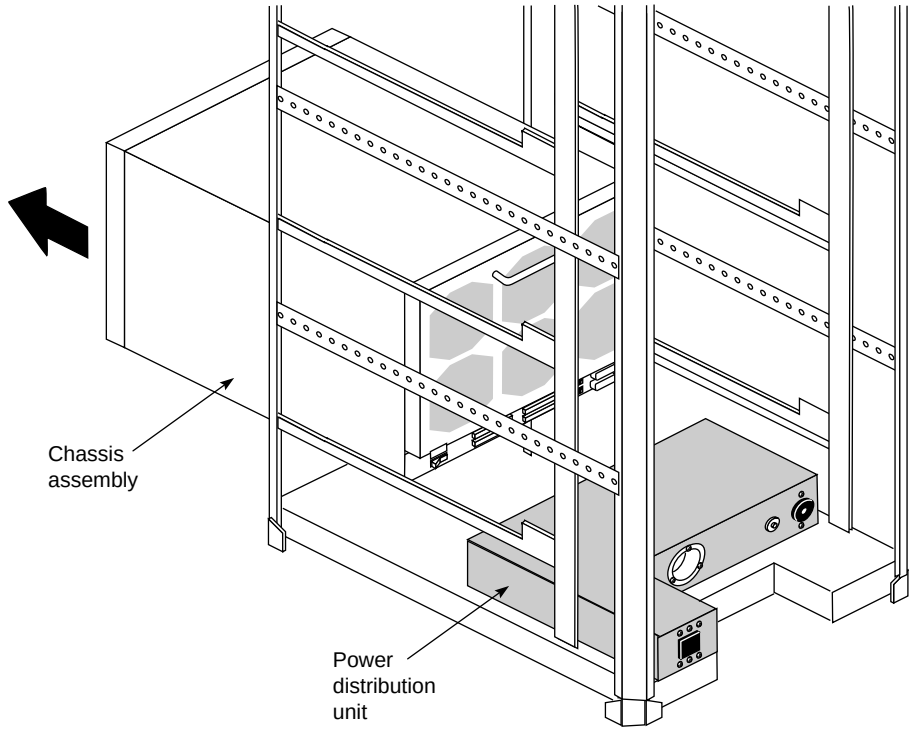

Figure 7-12 Clearing Access to the Power Distribution Unit in the Challenge RAID Rack

4. Remove the five screws that attach the lid of the failed power distribution unit and remove the lid, as shown in Figure 7-13.

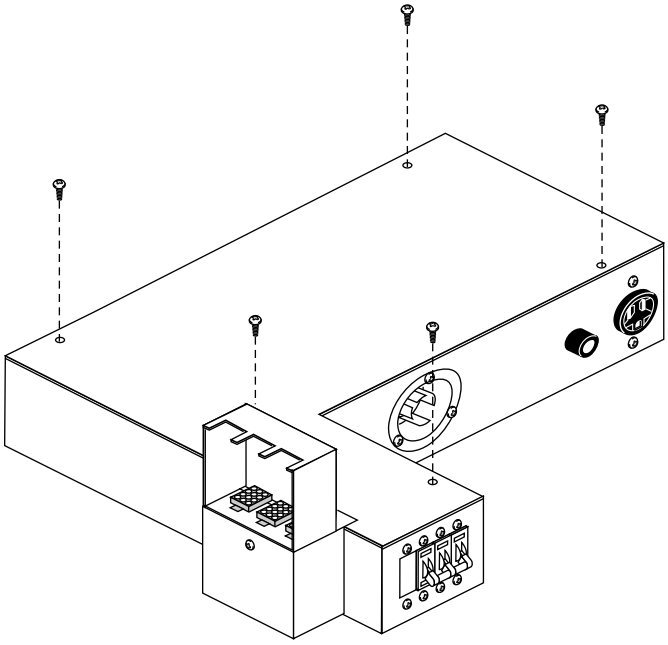

Figure 7-13 Removing the Power Distribution Unit Cover

5. Inside the unit, remove the four screws that attach it to the rack frame, as shown in Figure 7-14.

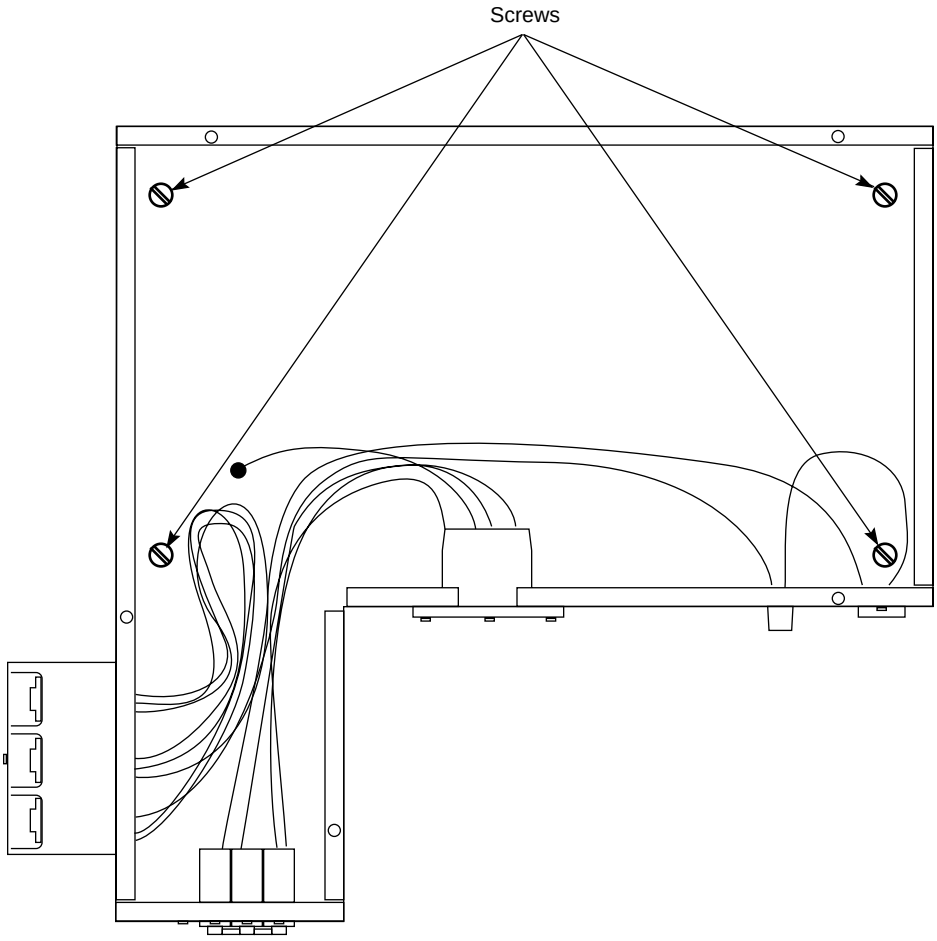

Figure 7-14 Removing the Power Distribution Unit (Bird's-Eye View)

6. Attach the replacement power distribution unit and its cover.
7. Reseat the chassis assembly and replace the screws. Replace the side panels. Reattach the power cable. Work with the system administrator at the site to power on the system.

7.6 Installing an Additional Chassis Assembly in a Challenge RAID Rack

Warning: **Two persons are required to install additional chassis assemblies in a Challenge RAID rack. Do not undertake this procedure unless another qualified Silicon Graphics SSE is available for it.**

To install an additional chassis assembly in a Challenge RAID rack, follow these steps:

1. Unpack the unit and verify that all parts are present. If they are not, notify the shipper of the missing parts.
2. To reduce the weight of the chassis assembly for insertion into the rack, remove FRUS from the new chassis assembly:
   - disk modules and front bezel (see Figure 7-15)
   - SPs, VSCs, BBUs
   - fan assemblies
3. Remove the filler panel at the next open chassis assembly location, starting at the bottom of the rack. (You fill the rack with chassis assemblies starting at the bottom, not at the top.)
4. With another SSE, lift the empty chassis assembly you want to install.
5. Lifting the empty chassis assembly, engage the rails at the rear of the open chassis location.
6. Slide the chassis assembly into the rack as shown in Figure 7-15.

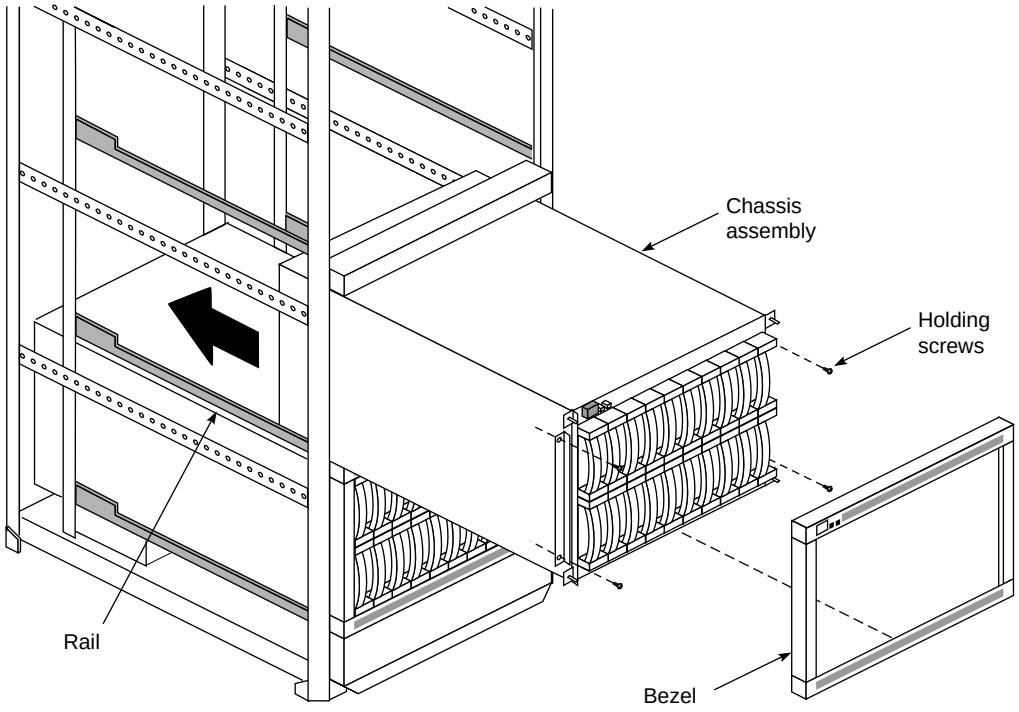

Figure 7-15 Installing an Additional Chassis Assembly in the Challenge RAID Rack

7. Secure the chassis assembly with six screws, as shown in Figure 7-16.

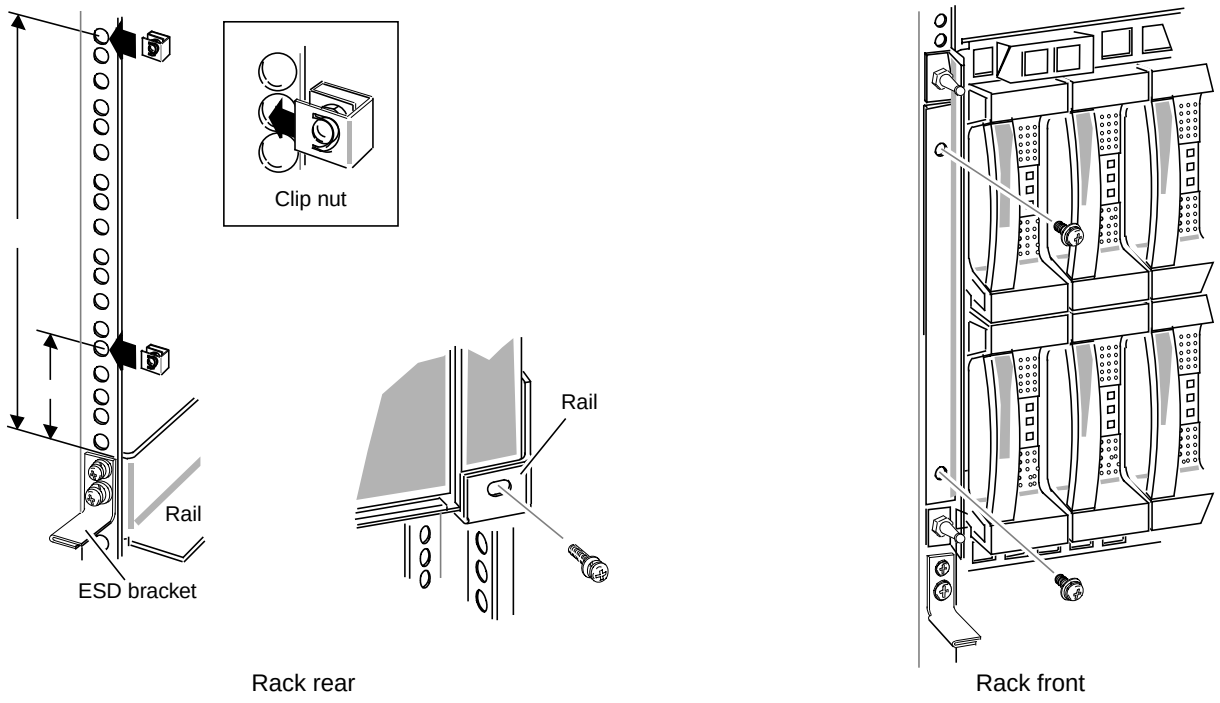

Figure 7-16 Securing the Chassis Assembly to the Rack

8. Position the front bezel and push it onto the ball studs, as shown in Figure 7-17.

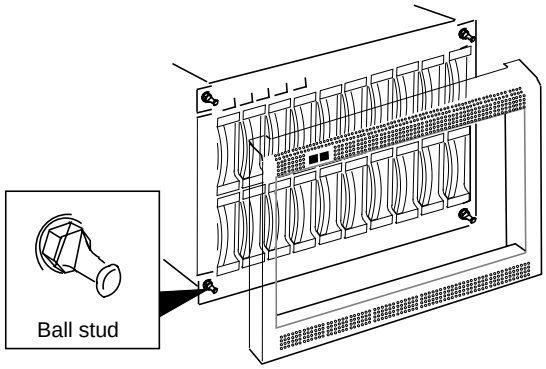

Figure 7-17 Attaching the Front Bezel

9. Install the FRUs that were removed.
10. Connect the power cables. Power the unit on and check for error conditions.
11. Repeat steps 3 through 10 for additional chassis assemblies to be installed.
12. Shut down the CHALLENGE RAID rack and the host; connect the SCSI cables between the host and the newly installed chassis assembly or assemblies.
13. Bring up the storage system and the host; configure the disks and perform other software tasks as explained in Chapter 4.

Appendix A

Technical Specifications

This appendix lists the technical specifications for the Challenge RAID deskside disk-array storage system. It also lists marketing codes and part numbers for FRUs (field-replaceable units).

Table A-1 gives specifications for the Challenge RAID deskside storage system.

Table A-1 Deskside Challenge RAID Specifications

Classification	Specification	Value
AC power requirements	Voltage	100 VAC to 200 VAC, single-phase autoranging 47 Hz to 63 Hz
	Current draw	9.0 A max. at 100 VAC input
	Power consumption	Apparent power: 900 VA max
	True power	880 W max
Connector	Type	L5-15P / L5-15R
Operating limits	Ambient temperature	10 degrees C to 38 degrees C (50 degrees F to 100 degrees F)
	Relative humidity	20% to 80% noncondensing
	Elevation	2439 m (8000 ft)
	Heat dissipation	2675-3000 BTU / hour maximum
	Shock	3 g @ 11 ms
	Vibration	0.25 g peak @ 5 Hz to 500 Hz
Nonoperating limits (shipping and storing)	Ambient temperature	-40 degrees C to 65 degrees C (-40 degrees F to 149 degrees F)
	Gradient, maximum	24 degrees C / hr (43.2 degrees F / hr)
	Relative humidity	10% to 90% noncondensing
	Elevation	7625 m (25,000 ft)

Table A-1 (continued) Deskside Challenge RAID Specifications

Classification	Specification	Value
Dimensions	Height	62.9 cm (24.75 in.)
	Width	35.6 cm (14.0 in.)
	Depth	76.2 cm (30.0 in.)
Weight	Minimum	Chassis with 5 disk modules, 1 SP, 2 VSCs, without packaging: 51.6 kg (114.0 lb)
	Maximum	Chassis with 20 disk modules, 2 SPs, 3 VSCs, BBU, without packaging: 83.9 kg (184.5 lb)
	Disk-drive module	1.6 kg (3.5 lb)
	Third power supply	2.4 kg (5.3 lb)
	Second storage-control processor	1.2 kg (2.6 lb)
	Battery backup unit	5.4 kg (12 lb)
Service clearance	Front	81.3 cm (32.0 in.)
	Rear	81.3 cm (32.0 in.)
Buses	External host bus	Differential fast and wide SCSI-2 (synchronous)
	internal storage system buses	Five single-ended SCSI buses

Table A-2 summarizes specifications for the rack storage system.

Table A-2 Challenge RAID Rack Specifications

Classification	Specification	Value
AC power requirements	Cabinet voltage	200 VAC to 240 VAC, single-phase 47 Hz to 63 Hz
	Current draw per chassis assembly	5.0 A max at 200 VAC input (20.0 A max for 4 chassis assemblies)
	Power consumption	Apparent power: 1000 VAC max True power: 980 W max
Connector	Type	Domestic: L6-30P / L6-30R International: IEC Plug 309 2-P

Table A-2 (continued) Challenge RAID Rack Specifications

Classification	Specification	Value
Operating limits	Ambient temperature	10 degrees C to 38 degrees C (50 degrees F to 100 degrees F)
	Relative humidity	20 to 80% noncondensing
	Elevation	2439 m (8000 ft)
	Heat dissipation	3168 x 103 J / hr (3000 BTU / hr) max per chassis assembly
	Shock	3 g @ 11 ms
	Vibration	0.25 g peak @ 5 Hz to 500 Hz
Nonoperating limits (shipping and storing)	Ambient temperature	-40 degrees C to 65 degrees C (-40 degrees F to 149 degrees F)
	Gradient, maximum	24 C / hr (43.2 degrees F / hr)
	Relative humidity	10 to 90% noncondensing
	Elevation	7625 m (25,010 ft)
Physical	Cabinet dimensions	Height: 180.34 cm (71.0 in)
		Width: 58.42 cm (23.0 in)
		Depth: 78.74 cm (31.0 in)
	Maximum cabinet weight	407.96 kg (899.4 lb): 4 chassis assemblies, each with 20 disk modules, 2 SPs, 3 power supplies, battery backup unit, without packaging
	See Table A-1 for weights of FRUs.	
Service clearance	Front access	81.3 cm (32.0 in)
	Rear and side access	76.2 cm (30.0 in)
Buses	External host bus	Differential fast and wide SCSI-2 (synchronous)
	Internal storage system buses	Five single-ended SCSI buses

Appendix B

The raidcli Command-Line Interface

This appendix is a complete guide to the */usr/raid5/raidcli* command (command-line interface, or CLI). It explains the *raidcli* parameters, which are summarized in Table B-1.

Table B-1 *raidcli* Parameters

Parameter	Use
bind	Bind (group) physical disks
chglun	Change parameters for a currently bound group
clearlog	Destroy the contents of the Challenge RAID storage controller's error log
clearstats	Reset statistics logging on the Challenge RAID storage-control processor (SP), setting all log counters to 0
firmware	Update the firmware on the Challenge RAID SP
getagent	Get names and descriptions of devices controlled by the SP
getcache	Get information about the storage-system caching environment
getcontrol	Get general system information
getcrus	Display status information on all system components, such as the fan module, battery backup unit, and so on
getdisk	Display status information on one disk module or all disk modules in the system
getlog	Display the SP log, in which errors are stored, or display the head and tail of the log
getlun	Display information about a disk group
getsp	Display firmware revision level and controller model number
setcache	Change information about the Challenge RAID controller's caching environment
setstats	Set statistics logging
trespass	Restore original status of an SP in a dual-SP RAID system when the SP is turned off and then turned back on: available only with RAID agent revision level 1.55 and higher
unbind	Deconfigure physical disks from their current logical configuration, destroying all data on the logical unit (group)

The *raidcli* command sends Challenge RAID storage management and configuration requests to an application programming interface (API) on the Challenge server. For the *raidcli* command to function, the agent—an interpreter between the command-line interface and the Challenge RAID storage system—must be running.

Note: Although the directory and command are *raidcli*, the command is valid for all or the original Sauna codeRAID levels.

The synopsis of the *raidcli* command is

`raidcli [-vp] [-d` *device*`]` *parameter* `[`*optional_arguments*`]`

In this syntax, options and variables have the following meanings:

-v	Enables verbose return.
-p	Parses the *raidcli* command without calling the API. If the string does not parse correctly, an error message is printed to *stderr*; otherwise there is no output.
-d *device*	Target RAID device. Use *raidcli getagent* for a list of RAID devices. This switch must be present for all *raidcli* management and configuration commands unless the environment variable indicates otherwise. This switch overrides an environment variable.

The environment variable RaidAgentDevice is the default value for the device when none is specified with the **-d** flag. If RaidAgentDevice is not set and no **-d** switch is present, an error is generated on all commands that need device information. For example:

```
setenv RaidAgentDevice sc4d2l0
setenv RaidAgentDevice /dev/scsi/sc4d2l0
```

The */dev/scsi* prefix is optional.

B.1 bind

The physical disk unit number is also known as the logical unit number, or LUN. The unit is a logical concept but is recognized as a physical disk unit by the operating system. The LUN is a hexadecimal number between 0 and 7.

Unlike standard disks, physical disk unit numbers (LUNs) lack a standard *geometry.* Disk capacity is not a fixed quantity between disk-array LUNs. The effective geometry of a disk-array LUN depends on the type of physical disks in the array and the number of physical disks in the LUN.

Note: Although *bind* returns immediate status for a RAID device, the bind does not complete for 15 to 60 minutes, depending on system traffic. Use *getlun* to monitor the progress of the bind; *getlun* returns the percent bound. When the bind is complete, each disk is noted as “Enabled” in the *getlun* output.

The command *raidcli bind* is used to group physical disks into RAID-0, RAID-1, RAID-1_0, RAID-3, or RAID-5 units or to create a hot spare:

`raidcli -d` *device* `bind` *raid-type lun-number disk-names* [*optional-args*]

Variables in this syntax are explained below.

-d *device* — Target RAID device, as returned by *raidcli getagent*.

raid-type — Choices are

- r0: RAID-0
- r1: RAID-1
- r1_0: RAID-1_0
- r3: RAID-3
- r5: RAID-5
- id: individual disk unit
- hs: hot spare

lun-number — Logical unit number to assign the unit (a hexadecimal number between 0 and 7).

disk-names — Indicates which physical disks to bind, in the format *bd*, where *b* is the physical bus name (a through e; be sure to use lowercase) and *d* is the device number on the bus (0 through 3). For example, a0 represents the device 0 on bus A, and e2 represents device 2 on bus E.

A RAID-0 bind requires a minimum of 3 (maximum 16) disks.

A RAID-1 bind requires 2 disks on separate buses.

A RAID-1_0 bind requires an even number of disks (minimum 2; maximum 16). For high availability, a member of each image pair must be on a different bus. Select the disks in this order: first disk on first bus, first disk on second bus, second disk on first bus, second disk on second bus, third disk on first bus, third disk on second bus, and so on.

A RAID-3 or RAID-5 bind also requires separate buses, each with five disk modules. Legal RAID-3 or RAID-5 bind configurations are

- a0 b0 c0 d0 e0
- a1 b1 c1 d1 e1
- a2 b2 c2 d2 e2
- a3 b3 c3 d3 e3

A RAID-5 bind requires a minimum of 3 disks and a maximum of 15 disks. For high availability, use groups of 5.

A hot spare bind requires one disk, which cannot be A0, B0, C0, D0, E0, or A3. The capacity of the hot spare must be at least as great as the capacity of the largest disk module it might replace.

All disk modules in a bind should have the same capacity to fully use the modules' disk space.

The optional arguments are as follows:

Note: You can enable auto-assign in RAIDGUI, but not in the command-line interface.

-r *rebuild-time* Maximum time in hours to rebuild a replacement disk. Default is four hours; legal values are any number greater than or equal to zero.

A rebuild time of two hours rebuilds the disk more quickly but degrades response time slightly. A rebuild time of zero hours rebuilds as quickly as possible but degrades performance significantly.

If your site requires fast response time and you want to minimize degradation to normal I/O activity, you can extend the rebuilding process over a longer period of time, such as 24 hours. You can change the rebuild time later without damaging the information stored on the physical disk unit.

Caution: The LUN is not fault-tolerant during the rebuild; selecting a short rebuild time minimizes the chance of data loss.

-s *stripe-size* Number of blocks per physical disk in a RAID-5 stripe. Default is 128; legal values are any number greater than zero.

The smaller the stripe element size, the more efficient the distribution of data read or written. However, if the stripe size is too small for a single host I/O operation, the operation requires accessing two stripes, thus causing the hardware to read and/or write from two disk modules instead of one. Generally, it is best to use the smallest stripe element size that will rarely force access to another stripe. The default stripe element size is 128 sectors. The size should be an even multiple of 16 sectors (8 KB).

Note: This value is ignored for hot spares.

-c *cache-flags* Values are

- none: no caching
- read: read caching
- write: write caching
- rw: read and write caching

The default is none. Caching is not specified for RAID-3.

-z *stripe-count* Sets the number of stripes in a LUN. For example, if you bind a RAID-5 LUN with a stripe count of 2, you partition the LUN into two stripes, thus preventing access to the remaining available space. This option is useful for short bind operations. Legal values are any number greater than or equal to 0. The default value is 0, which binds the maximum number of stripes available.

-r3 *raid3-memory-size*

Sets amount of RAID-3 memory to use for each LUN. This option is available only for firmware revision 9.0/SP model 7305 and higher and RAID agent 1.55 and higher.

Note: Before binding fast RAID-3 LUNs, you must set cache memory; see Section B.14, “setcache,” later in this appendix.

The following example binds disks A1, B1, C1, D1, and E1 into a RAID-5 logical unit with a logical unit number of 3, a four-hour maximum rebuild time, and a 128-block stripe size per physical disk, with read cache enabled:

```
raidcli -d sc4d2l0 bind r5 3 a1 b1 c1 d1 e1 -r 4 -s 128 -c read
```

The following example binds A2 and B2 into a RAID-1 logical unit with a LUN number of 2 and a four-hour maximum rebuild time, with read cache enabled:

```
raidcli -d sc4d2l0 bind r1 2 a2 b2 -r 4 -c read
```

The following example binds disks A1, B1, C1, and D1 into a RAID-1_0 logical unit with a LUN number of 1, a four-hour maximum rebuild time, and a 128-block stripe size per physical disk, with read cache enabled:

```
raidcli -d sc4d2l0 bind r1_0 1 a1 b1 c1 d1 -r 4 -s 128 -c read
```

The following example binds A1, B1, C1, D1, and E1 into a RAID-3 logical unit with a LUN number of 3.

```
raidcli -d sc4d2l0 bind r3 3 a1 b1 c1 d1 e1
```

The following example binds A2, B2, C2, D2, and E2 into a RAID-0 logical unit with a LUN number of 3, and a 128-block stripe size per physical disk, with read cache enabled:

```
raidcli -d sc4d2l0 bind r0 3 a2 b2 c2 d2 e2 -s 128 -c read
```

The following example binds disk E3 as a hot spare with a LUN number of 7:

```
raidcli -d sc4d2l0 bind hs 7 e3
```

There is no output for *raidcli* with the *bind* parameter. Errors are printed to *stderr.*

Note: For complete messages, it is recommended that you use the `-v` option.

B.2 chglun

The command *raidcli chglun* is used to change parameters for a LUN:

`raidcli -d` *device* `chglun [-l` *lun*`] [-c` *cache-flags*`] [-d` *default-owner*`]`
`[-r` *rebuild-time*`] [-i` *idle-thresh*`] [-t` *idle-delay-time*`] [-w` *write-aside*`] [-a` *autotrespass*`]`
`[-pf` *prefetch type*`] [-rt` *retention type*`] [-sc` *configure-read-cache*`]`
`[-sm` *segment-size/multiplier*`] [-pm` *prefetch-size/multiplier*`] [-mp` *max-prefetch*`]`
`[-pd` *min-prefetch*`] [-pc` *max-IO*`]`

Note: Only root can use this parameter.

In this syntax, variables mean the following:

-l *lun* Logical unit number to be changed

-c *cache-flags* Values are

- none: no caching
- read: read caching
- write: write caching
- rw: read and write caching

The default is none.

-d *default-owner*

Values are

- 1: change SP ownership of LUN
- 0: don't change ownership

If your storage system has dual SPs, see Section 4.11, "Understanding Dual Processors, Load Balancing, and Device Names," in Chapter 4.

-r *rebuild-time* Maximum time in hours to rebuild a replacement disk. Default is four hours; legal values are any number greater than or equal to 0.

-i *idle-thresh* Maximum number of I/Os that can be outstanding to a LUN and have the LUN still be considered idle. Used to determine cache flush start time. Legal values are 0 to 255; the default is 0. A value of 255 leaves the previous setting unchanged.

-t *idle-delay-time*

Amount of time in 100-ms intervals that a unit must be below *idle-thresh* to be considered idle. Once a unit is considered idle, any dirty pages in the cache can begin idle time flushing. Legal values are 0 to 255; the default is 20. A value of 255 leaves the previous setting unchanged.

-w *write-aside* The smallest write request size, in blocks, that will bypass the cache and go directly to the disks. Legal values are 16 to 65535; the default is 2048. A value of 65535 leaves the previous setting unchanged.

-a *autotrespass*

Sets autotrespass on or off. If autotrespass is enabled, the other SP in the system takes over ownership of a failed SP's LUNs. When the failed SP is replaced, ownership of its LUNs must be restored with *raidcli trespass*. Values are

- 0: disable autotrespass for the specified LUN; this value is recommended
- 1: enable autotrespass for the specified LUN

-pf *prefetch-type* Values are

- 0: disabled
- 1: constant length
- 2: variable length (default)

Note: See Table B-7 for an explanation of this and the remaining options for this command.

-rt *retention-type*

Retention determines whether or not prefetched data has equal or favored priority over host-requested data when the read cache becomes full. Values are

- 0: all data treated equally
- 1: prefetched data retained over requested data (default)

-sc *configure-read-cache*

Values are

- 0: configure read cache according to parameters specified in the command string
- 1: configure read cache according to default settings

-sm *segment-size/multiplier*

Segment size or segment multiplier determines the size of the segments that make up a prefetch operation.

Legal values for segment size are 0 to 255; a value of 255 leaves the previous setting unchanged. Legal values for the segment multiplier are 1 to 32; the default is 4.

- For constant-length prefetching, specify the segment size in blocks of data to prefetch in one read operation from the LUN.
- For variable-length prefetching, specify the segment multiplier to calculate the size of the segment. For example:

  ```
  sgmtblocks = readrequestsize * sgmtmultiplier
  ```

You can specify a segment size that is equal to or less (but not greater) than the prefetch size. You can specify a segment multiplier that is equal to or less (but not greater) than the prefetch multiplier. If you specify a segment multiplier that is equal to the prefetch multiplier, prefetch operations are not divided into segments.

-pm *prefetch-size/multiplier*

Prefetch size or prefetch multiplier determines the size of the segments that make up a prefetch operation. Legal values for prefetch size are 0 to 2048. Legal values for the prefetch multiplier are 1 to 32; the default is 4.

- For constant-length prefetching, specify the size of the prefetch data in blocks.
- For variable-length prefetching, specify the multiplier to calculate the prefetch length to fetch. For example:

  ```
  prefblocks = readrequestsize * prefmultiplier
  ```

-mp *max-prefetch*

Maximum number of blocks to prefetch for variable-length prefetching.

- For constant-length prefetching, all prefetches are of the constant length specified by the prefetch size and this field is ignored. Legal values are 0 to 2048; the default is 512.

- For variable-length prefetching, specify the maximum number of blocks to prefetch in response to a read request.

-pd *min-prefetch*
: Prefetch Disable specifies the minimum read request size at which the system disables prefetching. Legal values are 0 to 65535; the default is 129. Any value between 65280 and 65535 leaves the previous setting unchanged.

-pc *max-IO*
: The maximum number of outstanding I/Os at which the LUN stops prefetching. Legal values are 0 to 100; the default is 40.

The following example changes LUN 3 to perform write caching and rebuild in four hours; it does not change the default owner:

```
raidcli -d sc4d2l0 chglun -l 3 -c write -d 0 -r 4
```

There is no output for *raidcli* with the *chglun* parameter. Errors are printed to *stderr*.

B.3 clearlog

The *clearlog* parameter destroys the contents of the Challenge RAID SP's error log. Its syntax is

```
raidcli -d device clearlog
```

Note: You must be root to use this parameter.

This command has no output.

B.4 clearstats

The *clearstats* parameter resets statistics logging on the Challenge RAID SP, resetting all log counters to 0. The command *raidcli clearstats* is used to reset statistics logging:

```
raidcli -d device clearstats
```

Note: You must be root to use this parameter.

This command has no output.

B.5 firmware

The firmware (FLARE code) is stored in directories in */usr/raid5/flare*:

- */usr/raid5/flare/sauna7/flarecode.bin*: use this for AMD-based SPs with FLARE code rev. level 7.xx
- */usr/raid5/flare/sauna8/flarecode.bin*: use this for AMD-based SPs with FLARE code rev. level 8.00 through 8.49
- release 2.0 and earlier: */usr/raid5/flare/phoenix8/flarecode.bin*: use this for PowerPC-based SPs with FLARE code rev. level 8.50 through 8.99
- release 2.1 and later: */usr/raid5/flare/phoenix9/flarecode.bin*: use this for PowerPC-based SPs with FLARE code rev. level 9.00 through 9.99

Note: The bus must be quiesced, with all applications disabled and only the RAID agent running. You must use this command every time you replace a failed unbound disk module (A0, B0, C0, or A3).

The command *raidcli firmware* is used to update the firmware on the Challenge RAID SP:

`raidcli -d` *device* `firmware /usr/raid5/flare/`*directory*`/flarecode.bin [-o]`

You must be root to use this parameter.

In this string, the new firmware image */usr/raid5/flare/<directory>/flarecode.bin* contains microcode that runs on the SP and also a microcode image destined for the SP PROM, which runs the power-on diagnostics. If you enter a relative path instead of an absolute path, the current directory name is prepended to *filename* before the download starts.

The flag -o instructs the command not to prompt for permission (confirmation).

This command has no output.

Note: Once the microcode has been downloaded, each SP in the cabinet reboots. The reboot process takes several minutes, during which time no command-line interface commands are accepted other than *getagent*. When the code is downloaded, the system reboots the SPs. After the download is complete, stop and restart the agent.

Use *getagent* or *getsp* to return the new firmware version.

B.6 getagent

The command *raidcli getagent* is used to display information on devices controlled by the API:

```
raidcli getagent
```

The following is sample output for one device; normally, the output would give information on all devices:

```
Agent Rev: 1.55.0
Name: Disk Array
Desc: RAID5 Disk Array
Node: sc4d2l0
Signature:0xf3b51700
Peer Signature: 0x657e0a00
Revision: 9.3.0, Model: 7305
SCSI ID: 1
Prom Rev: 1.55.0
SP Memory: 64
Serial No: 96-7240-808
```

Table B-2 summarizes entries in the *raidcli getagent* output.

Table B-2 Output of *raidcli getagent*

Entry	Meaning
Agent Rev	Revision number of RAID agent
Name	ASCII string found in the agent configuration file that assigns a name to the node being accessed (see node description below)
Desc	ASCII string found in the agent configuration file that describes the node being accessed (see node description below)
Node	The */dev/scsi* entry that the agent uses as a path to the actual SCSI device; this value must be entered by the user for every CLI command (except *getagent*)
Signature	Unique 32-bit identifier for the SP being accessed through Node
Peer Signature	Unique 32-bit identifier for the other SP in the chassis; 0 if no additional SP is present
Revision	Revision of firmware currently running on the SP: 7305: PowerPC-based SP 7624: AMD-based SP
SCSI ID	SCSI ID number
Prom Rev	Current PROM revision on the SP
SP Memory	Amount of DDRAM present on the SP
Serial No	12-digit ASCII string that uniquely identifies this subsystem

B.7 getcache

If you use storage system caching, you can use the *raidcli getcache* command to get information on cache activity. The information in this command's output, particularly the percentage of cache hits, may help you decide on the most efficient cache page size and whether a physical disk unit really benefits from caching.

The command *raidcli getcache* is used to display cache information:

`raidcli -d` *device* `getcache`

The following is sample output:

```
Usable Cache: 8
Page size: 2
Cache State: Enabled
Low Watermark; 75
High Watermark: 90
SP A Cache Pages: 2048
SP B Cache Pages: 2047
Unassigned Cache Pages: 0
Read Hit Ratio: 82
Write Hit Ratio: 74
Prct Dirty Cache Pages = 63
Prct Cache Pages Owned = 50
Prct Read Flushes = 10
Prct Write Flushes = 12
SPA Read Cache State enabled
SPB Read Cache State enabled
Minimum system memory allowed 2 MB
Physical memory size of spA = 8 MB
Physical memory size of spB =16 MB
System memory size of spA = 8 MB
System memory size of spB = 8 MB
Read Cache Size of spA = 3 MB
Read Cache Size of spB = 3 MB
Raid 3 memory size for spA = 2 MB
Raid 3 memory size for spB = 2 MB
Free Memory size of spA = 5 MB
Free Memory size of spB = 5 MB
```

Table B-3 summarizes entries in the *raidcli getcache* output.

Table B-3 Output of *raidcli getcache*

Entry	Meaning
Usable Cache	Size of cache in megabytes to use for caching, not greater than the SP memory size as displayed by raidcli getagent (see Section B.6, "getagent," earlier in this appendix). Valid values are 0, 8, and 64. Defaults are 0 (user has selected to disable the cache) and 8 (user has selected to enable the cache). The command-line interface does not let you specify more memory than you have. If you specify less than you have, the remaining memory is unused.
Page Size	User-specified page size in KB for caching; for example, 2 means a 2 KB page size.
Cache State	Enabled: SP is fully functional. Disabled: SP not capable of or configured for caching. Synching: SP is synchronizing its cache with the peer SP. Enabling: cache is in the process of becoming enabled. Quiescing: cache was enabled but system experienced a fault or was disabled by user. The system is now finishing processing outstanding I/O before saving the contents of the cache. Frozen: SP has finished quiescing and is now waiting for peer SP to finish quiescing or dumping. Dumping: SP is in the process of dumping contents of the cache to a private location on the disk. Initing: SP is initializing.
Low Watermark	Percentage of dirty cache pages that, when reached during flush operations, will cause the SP to cease flushing the cache. Valid values are 0 through 100; the default is 50, regardless of whether caching is enabled or disabled.
High Watermark	Percentage of dirty cache pages which, when reached during flush operations, will cause the SP to cease flushing the cache. Valid values are 0 through 100; the setting must be greater than the low watermark. The default is 75, regardless of whether caching is enabled or disabled.
SP A Cache Pages SP B Cache Pages	Total number of cache pages currently owned by SP A; each page has the cache page size selected when storage-system caching was set up. This number equals the cache size divided by the cache page size, minus space for checksum tables. If the storage system has two working SPs, they divide the total number of pages between them. If an SP is idle for a long period or fails, the active SP may increase its share of pages.

Table B-3 (continued) Output of *raidcli getcache*

Entry	Meaning
Unassigned Cache Pages	Number of dirty cache pages not owned by A or B. These can occur when a unit is broken and there are no disks to which to flush the dirty pages.
Read Hit Ratio	Percent of read requests to the SP that can be satisfied from the cache without requiring disk access.
Write Hit Ratio	Percent of write requests to the controller that can be satisfied with the cache without requiring disk access. A write cache hit occurs when the SP finds and modifies data in the write cache memory, which usually saves a write operation. For example, a hit that occurs when a page is sought in a RAID 5 LUN eliminates the need to read, modify, and write the data. High ratios are desirable because each hit indicates at least one disk access that was not needed.
Prct Dirty Cache Pages	Percentage of pages that have been modified in the SP's write cache, but that have not yet been written to disk. A high percentage of dirty pages means the cache is handling many write requests.
Prct Cache Pages Owned	Percentage of cache pages owned by this controller when pages from other SP are taken into account.
Prct Read Flushes	Percentage of reads to SP that caused the cache to flush a page(s).
Prct Write Flushes	Percentage of writes to SP that caused the cache to flush a page(s).
SP Read Cache State	Enabled or Disabled.
Minimum system memory allowed	AMD-based SP: 2 MB; PowerPC-based SP: 4 MB.
Physical memory size of SP	Number of MB in SP's physical memory.
System memory size of SP	Number of MB in SP's system memory.
Read Cache Size of SP	Number of MB in SP's read cache.
Raid 3 memory size for SP	Number of MB allocated to RAID-3 memory.
Free Memory size of SP	Unassigned memory.

B.8 getcontrol

The command *raidcli getcontrol* is used to get general system information:

`raidcli -d` *device* `getcontrol`

A sample output of this command follows:

```
System Fault LED: OFF
Statistics Logging: ON
System Cache: ON
Max Requests: 23
Average Requests: 5
Hard errors: 0
Total Reads: 18345
Total Writes: 1304
Prct Busy: 25
Prct Idle: 75
System Date: 1/5/1996
Day of the week: Friday
System Time: 12:43:54
read_requests: 145
write_requests: 34642049
blocks_read: 1633
blocks_written: 363696414
sum_queue_lengths_seen_by_arrivals: 401
arrivals_to_non_zero_queue: 400
hw_flush_on: 0
idle_flush_on: 5314
lw_flush_off: 0
write_Cache_flushes: 98
write_cache_blocks_flushed: 1150
Internal bus 1 busy ticks: 429
Internal bus 1 busy ticks: 429
Internal bus 2 busy ticks: 429
Internal bus 2 busy ticks: 429
Internal bus 3 busy ticks: 429
Internal bus 3 busy ticks: 429
Internal bus 4 busy ticks: 429
Internal bus 4 busy ticks: 429
Internal bus 5 busy ticks: 429
Internal bus 5 busy ticks: 429
Controller busy ticks: 1403792
Controler idle ticks: 3775334
```

In the *getcontrol* output,

- statistics logging is always turned on by the agent
- hard errors are those returned to the host
- total reads and writes are the totals as seen by the SP
- system time is in 24-hour format

B.9 getcrus

The command *raidcli getcrus* is used to display state information on every field-replaceable unit in the Challenge RAID storage system except disks in the disk modules:

```
raidcli -d device getcrus
```

Sample output of this command follows.

```
FANA State: Present
FANB State: Present
VSCA State: Present
VSCB State: Present
VSCC State: Present
SPA State: Present
SPB State: Present
BBU State: Present
```

Table B-4 interprets items in this output. Values for all entries of the output are Present or Not Present, except as noted.

Table B-4 Output of *raidcli getcrus*

Output	Meaning
FANA, FANB	Fan banks A and B.
VSCA	Power supply (voltage semi-regulated converter).
VSCB	Optional second power supply.
SPA	Storage control processor.
SPB	Optional second storage-control processor.
BBU	Battery backup unit, which has three states: Present (fully charged) and Not Present (removed or charging) If the battery backup unit takes longer than an hour to charge, it shuts itself off and transitions to the “Faulted” state.

B.10 getdisk

The command *raidcli getdisk* is used to display information on all disks in the system, whether bound or not:

```
raidcli -d device getdisk
```

For information on a particular disk, use

```
raidcli -d device getdisk [diskposition]
```

In this command, *diskposition* has the format *bd*, where *b* is the bus on which the disk is located (a through e) and *d* is the device number (0 through 3). Figure B-1 diagrams disk module locations.

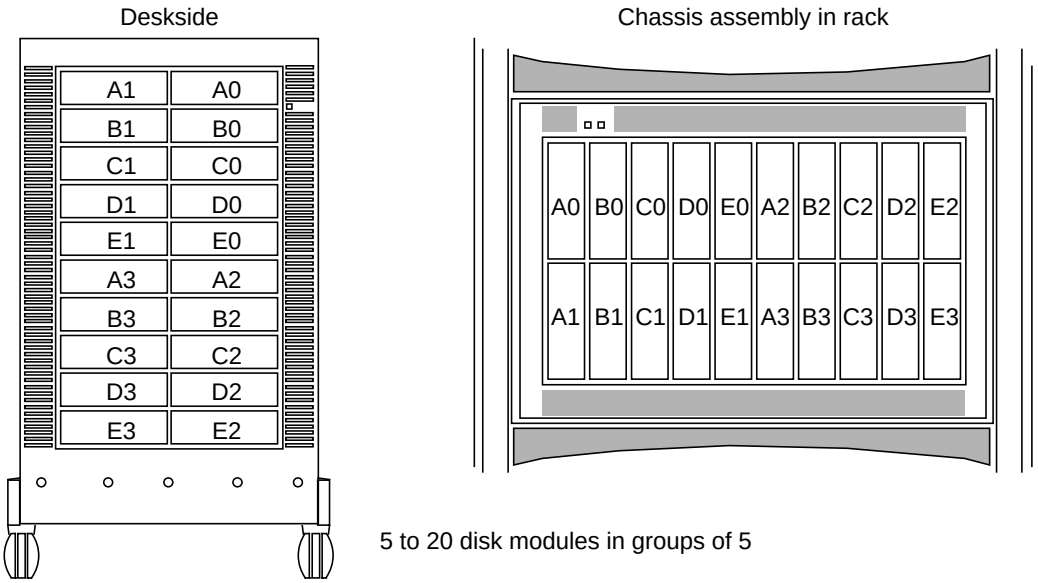

Figure B-1 Disk Module Locations

For example, the following command retrieves information for disk A0:

```
raidcli -d sc4d2l0 getdisk a0
```

A sample output of this command follows.

```
A0 Vendor Id: <manufacturer>
Ao Product Id: <part number>
A0 Lun: 0
A0 State: Bound and Not Assigned
A0 Hot Spare: NO
A0 Prct Rebuilt: 100
A0 Prct Bound: 100
A0 Serial Number: 00019699
A0 Capacity: 0x000f42a8
A0 Private: 0x00009000
A0 Bind Signature: 0x1c4eb2bc
A0 Hard Read Errors: 0
A0 Hard Write Errors: 0
A0 Soft Read Errors: 0
A0 Soft Write Errors: 0
A0 Read Retries: 0
A0 Write Retries: 0
A0 Remapped Sectors: 0
A0 Number of Reads: 1007602
A0 Number of Writes: 1152057
```

Table B-5 interprets items in this output.

Table B-5 Output of *raidcli getdisk*

Output	Meaning
Vendor Id	Manufacturer of disk drive
Product Id	Part number of disk
Lun	Logical unit number to which this disk is bound
State	Removed: disk is physically not present in the chassis or has been powered off Off: disk is physically present in the chassis but is not spinning Powering Up: disk is spinning and diagnostics are being run on it Unbound: disk is healthy but is not part of a LUN Bound and Not Assigned: disk is healthy, part of a LUN, but not being used by this SP Rebuilding: disk is being rebuilt Enabled: disk is healthy, bound, and being used by this SP Binding: disk is in the process of being bound to a LUN Formatting: disk is being formatted
Hot Spare	YES or NO
Prct Rebuilt	Percentage of disk that has been rebuilt
Prct Bound	Percentage of disk that has been bound
Serial Number	Serial number from disk inquiry command
Capacity	Actual disk capacity in blocks
Private	Amount of physical disk reserved for private space
Bind Signature	Unique value assigned to each disk in a logical unit at bind time
Hard Read Errors	Number of hard errors encountered on reads for this disk
Hard Write Errors	Number of hard errors encountered on writes for this disk
Soft Read Errors	Number of soft errors encountered on reads for this disk
Soft Write Errors	Number of soft errors encountered on writes for this disk
Read Retries	Number of retries occurring during reads
Write Retries	Number of retries occurring during writes
Remapped Sectors	Number of sectors that have been remapped
Number of Reads	Number of reads this disk has seen
Number of Writes	Number of writes this disk has seen

B.11 getlog

The SP maintains a log of event messages in its memory. These events include hard errors, startups, and shutdowns involving disk modules, fans, SPs, power supplies, and the battery backup unit. Periodically, the SP writes this log to disk to maintain it when SP power is off. The log can hold over 2,000 event messages; it has a filter feature that lets you select events by device or error message.

The event messages are in chronological order, with the most recent messages at the end. The command *raidcli getlog* is used to display the entire log:

```
raidcli -d device getlog [+N | -N] [-v]
```

The **+N** argument displays the newest *n* entries in the log, starting with the oldest entry:

```
raidcli -d device getlog +N
```

The **-N** argument displays the oldest *n* entries in the log, starting with the oldest entry:

```
raidcli -d device getlog -N
```

The following is a possible output of the command `raidcli getlog +5`:

```
12/17/95 09:59;51 A3: (A07) Cru Powered Down [0x47]
12/17/95 09:59;51 A3: (608) Cru Ready [0x0]
12/17/95 09:59;51 A3: (603) Cru Rebuild Started [0x0]
12/17/95 09:59;51 A3: (604) Cru Rebuild Complete [0x0]
12/17/95 09:59;51 A3: (602) Cru Enabled [0x0]
```

These entries show that a field-replaceable unit (disk module, fan unit, battery backup unit, SP) has failed, been replaced, been rebuilt, and been enabled.

At the tail of each log entry is an error code in brackets that gives diagnostic information when it is available. Table B-6 summarizes the error codes.

Table B-6 *getlog* Error Codes

Error Code	Explanation
0	No error
1	Invalid command-line parameters
2	Too few parameters
3	Too many parameters
4	Invalid bind type
5	Invalid LUN number
6	Invalid rebuild time
8	Invalid number of disks in *bind* command
9	Valid disk names are of format a0, b1, ... e3, etc.
10	Invalid stripe size

Table B-6 (continued) *getlog* Error Codes

Error Code	Explanation
11	Invalid disk name
12	Invalid cache flags
13	Invalid default owner
14	Incorrect number of *chglun* parameters
15	Unable to determine name of target host machine
16	Enable-disable flag invalid
17	Invalid usable cache size
18	Invalid page size
19	Invalid watermark value
20	High watermark less than low watermark
21	No device name listed
22	Invalid idle threshold
23	Invalid idle delay
24	Invalid write aside size
25	Disks must be on separate buses for bind
26	The agent was unable to configure any devices
27	LUN does not exist
28	LUN already exists
29	Cannot get current working directory for *firmware* command
50	Agent encountered an error during SCSI execution
51	Agent encountered an error during serial port execution
52	Agent returned an operating system error
53	Agent returned an internal agent error code
54	Cannot communicate with agent
55	RPC error encountered

More detailed descriptions of the above errors can be seen when verbose mode is enabled (-v).

B.12 getlun

The command *raidcli getlun* is used to display information on a logical unit and the components in it. The syntax is

`raidcli -d` *device* `getlun` *lun-number*

The following example displays information about LUN 0:

```
raidcli -d sc4d2l0 getlun 0
```

Note: If numbers in the output look unreasonable, particularly for numerical statistics such as Queue Max, use *clearstats* to reset all statistics.

Below is a truncated output for a RAID-5 group of five disks.

Note: Information on individual disks is not displayed unless statistics logging is enabled with *raidcli setstats*. See Section B.8, "getcontrol," earlier in this appendix.

```
Prefetch size multiplier = 4
Segment size/multiplier = 4
Maximum prefetch = 512
Prefetch Disable = 129
Prefetch idle content = 40

Variable length prefetching

Prefetch data retained

Read cache configured according to specified parameters
Read Histogram[1] 194560Write Histogram[1] 1945 60
Write Histogram[6053104] 1945 60
Write Histogram[10186] 1945 60
Write Histogram[6053103] 1945 60
Write Histogram[12109245] 1945 60
Write Histogram[0] 1945 60
Write Histogram[0] 1945 60
Write Histogram[0] 1945 60
Write Histogram[0] 1945 60
Write Histogram[0] 1945 60
Read Histogram[5] 1945 60
Read Histogram[0] 1945 60
Read Histogram[2] 1945 60
Read Histogram[2] 1945 60
Read Histogram[0] 1945 60
Read Histogram[0] 1945 60
Read Histogram[0] 1945 60
Read Histogram[0] 1945 60
Read Histogram[0] 1945 60

Read histogram overflows 0
Write Histogram[1] 194560
Write Histogram[6053104] 194560
Write Histogram[10186] 194560
Write Histogram[6053103] 194560
Write Histogram[12109245] 194560
Write Histogram[0] 194560
Write Histogram[0] 194560
```

```
Write Histogram[0] 194560
Write Histogram[0] 194560
Write Histogram[0] 194560

Write histogram overflows 0
Read requests 10
Write Requests 24225638
Blocks read 61
Blocks written
Read cache hits 4
Read cache misses 6
Prefetched blocks 51
Unused prefetched blocks 0
Write cache hits 24225043
Forced flushes 0

Type: RAID5
Stripe size: 128
Capacity: 0x10000
Current owner: YES
Auto-trespass: Disabled
Auto-assign; Enabled
Write cache: Disabled
Read cache: Disabled
Idle Threshold: 0
Idle Delay Time: 20
Write Aside Size: 2048
Default Owner: YES
Rebuild Time: 0
Read Hit Ratio: 0
Write Hit Ratio: 0
Prct Reads Forced Flushed: 0
Prct Writes Forced Flushed: 0
Prct Rebuilt: 100
Prct Bound: 100

A0 Enabled
A0 Reads: 62667
A0 Writes: 29248
A0 Blocks Read: 3212517
A0 Blocks Written: 471642
A0 Queue Max: 26
A0 Queue Avg: 1
A0 Avg Service Time: 14
A0 Prct Idle: 100
A0 Prct Busy: 0
A0 Remapped Sectors: 0
A0 Read Retries: 50
A0 Write Retries: 0

B0 Enabled
B0 Reads: 66946
[etc.]

C0 Enabled
C0 Reads: 69342
[etc.]
```

```
D0 Enabled
D0 Reads: 68558
[etc.]

E0 Enabled
Eo Reads: 69721
[etc.]
```

Table B-7 summarizes entries in the *raidcli getlun* output.

Table B-7 Output of *raidcli getlun*

Entry	Meaning
Prefetch size / multiplier	For prefetching (read-ahead caching), prefetch size and multiplier determine the amount of data prefetched for one host read request. For constant-length prefetching, the prefetch size is the number of blocks to prefetch. The range for the prefetch size is 1-2048 blocks. The range for the prefetch multiplier is 1-32; the default is 4.
Segment size / multiplier	Segment size and multiplier determine the size of the segments that make up a prefetch operation. An SP reads one segment at a time from the LUN because smaller prefetch requests interfere less with other host requests. The range for the segment size is 0-254 blocks. The range for the prefetch multiplier is 1-32; the default is 4.
Maximum prefetch	Maximum number of blocks to prefetch for variable-length prefetching. The range is 1-2048; the default is 512.
Prefetch Disable	Number read request sectors that disables prefetch; this number is so large that prefetching data would not be beneficial; for example, if the amount of requested data is equal to or greater than the size of the read cache. The range is 1-65279; the default 129.
Prefetch idle content	Leave this parameter set to 40 (the default).
Variable / Constant length prefetching	Prefetch type: determines whether to prefetch data of a variable or constant length or to disable prefetching. The default is variable-length prefetching.
Prefetch data retained	Retention determines whether or not prefetched data has equal or favored priority over host requested data when the read cache becomes full. The default is Favor prefetch, which is suitable for most applications.

Table B-7 (continued) Output of *raidcli getlun*

Entry	Meaning
Type	RAID0: nonredundant individual access array RAID1: mirrored pair RAID10: mirrored RAID-0 group (RAID-1_0) RAID3: parallel access array RAID5: individual access array DISK: individual disk unit SPARE: hot spare UNKNOWN: disk was bound in a firmware version that is not compatible with the current firmware version
Stripe Size	Sectors per disk per stripe with which the unit was bound
Capacity	Number of sectors total for use by user
Current owner	YES if this SP owns the unit; NO if it does not Set with *chglun*
Auto-trespass	Always Disabled
Auto-assign	Enabled or Disabled; Disabled is recommended
Write Cache	Enabled means this LUN is write caching; otherwise, Disabled
Read Cache	Enabled means this LUN is read caching; otherwise, Disabled
Idle Threshold	Maximum number of I/Os outstanding; used to determine cache flush start time; set with *chglun*
Idle Delay Time	Amount of time in 100-ms intervals that unit is below idle threshold; set with *chglun*
Write Aside Size	Smallest write-request size in blocks that can bypass the cache and go directly to the disk; set with *chglun*
Default Owner	YES if this SP is the default owner (not necessarily current owner) of this LUN, otherwise, NO
Rebuild Time	Amount of time in hours in which a rebuild should be performed. 0 means rebuild as fast as possible, but means a degradation in host I/O performance
Read Hit Ratio	Percentage of read requests to the controller that can be satisfied from the cache without requiring disk access
Write Hit Ratio	Percentage of write requests to the cache that can be satisfied with the cache without requiring a disk access
Prct Reads Forced Flushed	Percentage of read requests that flushed the cache

Table B-7 (continued) Output of *raidcli getlun*

Entry	Meaning
Prct Writes Forced Flushed	Percentage of write requests that flushed the cache
Prct Rebuilt	Percentage complete during a rebuild
Prct Bound	Percentage complete during a bind
Diskname State	Enabled, Binding, etc. (same as for *getdisk*)
Diskname Reads	Total number of reads this disk has done
Diskname Writes	Total number of writes this disk has done
Diskname Blocks Read	Total number of blocks this disk has read
Diskname Blocks Written	Total number of blocks this disk has written
Diskname: Queue Max	Maximum number of I/Os queued up to this drive
Diskname: Queue Avg	Average number of I/Os queued up to this drive
Diskname: Avg Service Time	Average service time in milliseconds
Diskname: Prct Idle	Percentage of time disk is not servicing request
Diskname: Prct Busy	Percentage of time disk is servicing request
Diskname: Remapped Sectors	Number of remaps that have occurred on this disk
Diskname Read Retries	Number of read retries that have occurred on this disk
Diskname Write Retries	Number of write retries that have occurred on this disk

B.13 getsp

The *getsp* parameter returns the firmware (FLARE) revision number and the SP model number for the specified device:

`raidcli -d` *device* `getsp`

In the output, Model 7624 is the AMD-based SP and Model 7305 is the PowerPC-based SP. Example output:

```
Revision 8.20.0, Model: 7624
```

For systems containing many Challenge RAID chassis assemblies, this parameter is especially useful as an alternative to *raidcli getagent* when the firmware revision number and model number of a specific SP are desired.

This parameter is available for RAID agent 1.55 and later.

B.14 setcache

The cache parameters you specify for the entire storage system are the cache size of 8 or 64 MB, depending on the amount of memory the SP has, and the cache page size, as 2, 4, 8, or 16 KB.

The command *raidcli setcache* is used to set up caching:

`raidcli -d` *device* `setcache` *enable* | *disable* `[-u` *usable*`] [-p` *page*`]`
`[-l` *low_watermark]* `[-h` *high_watermark*`]`
`[-sma` *system-memory-SPA*`] [-smb` *system-memory-SPB*`]`
`[-sta` *read-cache-state-SPA*`] [-stb` *read-cache-state-SPB*`]`
`[-rca` *read-cache-size-SPA*`] [-rcb` *read-cache-size-SPB*`]`

In this syntax, variables have the following meanings:

enable | *disable*

Enter 1 to enable write caching or 0 to disable write caching.

Note: When enable / disable is the only option used with *setcache*, 1 enables system read and write cache and 0 disables system read and write cache.

Using *setcache* 1 with options takes 45 seconds while cache states are disabled; cache dumping might take some time.

-u *usable* Size in megabytes to use for caching, not greater than the SP memory size as displayed by *raidcli getagent* (see Section B.6, "getagent," earlier in this appendix). Valid values are 0 through 64 in increments of 1 MB.

The command-line interface does not let you specify more memory than you have. If you specify less than you have, the remaining memory is unused.

Note: For mirrored caching, both SPs must have the same amount of cache memory in order for caching to be preserved in the event of shutdown or other power loss.

-p *page* Size in KB of pages into which to partition the cache. Valid sizes are 2, 4, 8, and 16. The default is 2, regardless of whether caching is enabled or disabled.

-l *low* Percentage of cache full that discontinues flushing. Valid values are 0 through 100; the default is 50, regardless of whether caching is enabled or disabled.

-h *high* Percentage of cache full that initiates flushing. Valid values are 0 through 100; the setting must be greater than the low watermark. The default is 75, regardless of whether caching is enabled or disabled.

-sma *system-memory-SPA*, -smb *system-memory-SPB*

Set the system memory on SP A or B, respectively.

Note: This option reboots each SP in the system. Rebooting takes several minutes, during which *raidcli getagent* is the only command accepted. When the reboot completes, stop and restart the agent before any *raidcli* commands are accepted.

-sta *read-cache-state-SPA*, stb *read-cache-state-SPB*

Set the read cache state for SP A or B, respectively: 0 to disable, 1 to enable.

-rca *read-cache-size-SPA*, -rcb *read-cache-size-SPB*

Set the read cache size in MB for SP A and B, respectively.

-r3 *raid3-memory-size*

Sets the fast RAID-3 memory size for this SP and its peer SP (if present). When this option is used, all options other than *enable* / *disable* are ignored. If the peer SP is present and the configuration requires failover, you must allocate twice as much memory; for example, if 6 MB is required for a RAID-3 LUN, you must allocate 12 MB with this option.

Note: This option reboots each SP in the chassis. The reboot process takes several minutes, during which no commands are accepted except *raidcli getagent*. When the reboot completes, you must stop and restart the agent before any commands are accepted.

The sum of an SP's write cache size, system memory size, and read cache size must be less than or equal to the SP's physical memory size.

This command has no output.

You can change the cache size, the cache page size values, or the type of caching for any physical disk unit without affecting the information stored on it. Follow these steps:

1. Disable the cache:

   raidcli -d device setcache 0

2. Wait for the cache memory to be written to disk, which may take several minutes.

3. Reassign cache size and re-enable caching:

   raidcli -d device setcache 1 -u 64 [parameters]

Note: Before changing a cache parameter, you must always disable the cache.

The following example enables the system cache with an 8 MB cache partitioned into 8 KB pages, with a 50% low watermark value and a 75% high watermark value:

```
setcache -d sc4d2l0 1 -u 8 -p 8 -l 50 -h 75
```

B.15 setstats

The *setstats* parameter sets statistics logging:

```
raidcli setstats [-on | -off]
```

If no optional switch is present, this command returns the current state of statistics logging.

You must be root to use this parameter.

The command *raidcli setstats -on* enables statistics logging; *raidcli setstats -off* disables it. With the optional switch, this command has no output.

If statistics logging is off, certain options in other commands are disabled. The option descriptions for these other commands specify whether statistics logging must be on for their options to be valid.

B.16 trespass

Note: This parameter is available only with RAID agent revision level 1.55. For systems with earlier revisions of the agent, use */sbin/trespass*.

To restore the original status of an SP in a dual-SP RAID system when the SP is turned off and then turned back on again, use the *trespass* parameter:

`raidcli trespass -d` *device action*

where *action* is one of the following:

- mine: tells the specified device to take ownership for all LUNs for which it is the default owner
- *LUN-number*: tells the specified device to take ownership of the specified LUN
- all: tells the specified device to take ownership of all LUNs

For example, if SP A was turned off and is now operational again, and SP A is connected to SCSI bus 6 with target (SCSI ID) 6, you would enter

```
raidcli trespass sc6d6l0 mine
```

The following example assigns LUN 3 to this device.

```
raidcli trespass sc6d6l0 3
```

This command has no output; errors are printed to *stderr*.

B.17 unbind

The *unbind* parameter deconfigures physical disks from their current logical configuration into LUNs.

Caution: This parameter destroys all data on the LUN (disk group).

You must be root to use this parameter. The command prompts you for verification before issuing the unbind, unless you override it with the -o flag.

The command *raidcli unbind* is used to unbind disks:

`raidcli -d` *device* `unbind` *lun-number* `[-o]`

In this syntax, variables and options mean:

lun-number	Number of the logical unit (LUN) to deconfigure.
-o	When *raidcli unbind* is entered, a prompt appears asking the user for verification before the unbind is issued. This flag disables this prompt.

This command has no output.

The following example destroys LUN 3 and frees its disks to be reconfigured, with no prompting to the user:

```
unbind -d sc4d2l0 3 -o
```

Appendix C

Troubleshooting

This chapter explains how to solve common system problems:

- determining firmware failure
- restarting the agent
- reconfiguring the system as system type 9
- reenabling command-tagged queuing
- determining caching problems
- reassigning LUN ownership

C.1 Determining Firmware Failure

When you start the RAID system via serial connection, the software outputs the alphabet to the serial device (such as a laptop or ASCII terminal). If the output remains at letter W, the meaning is that the system cannot find valid firmware. That is, drive A0, B0, C0, and A3 are all missing or are all defective. The firmware in these disks is required for SP operation. Replace these disks with functional ones.

C.2 Restarting the Agent

Turn on the Challenge RAID storage system. To see if the agent is active, type

```
/usr/raid5/getagent
```

Note these error messages and their solutions:

- `Can't open device`

  The agent is unable to start. The Challenge RAID storage system might have been configured as system type 1, instead of system type 9. To remedy this situation, follow instructions in the next section.

- `Cannot connect to agent`

  The agent is not running and must be restarted. See "Starting the Agent" in Chapter 4 for instructions.

- in *syslog*: `Is the device configured correctly?`

  Check *raid5-agent.config*, particularly the device specification; make corrections as needed.

Note: Error messages from the agent are not reported to the console, but are stored in the *syslog* file.

C.3 Reconfiguring the Challenge RAID Storage System as System Type 9

You can access the Challenge RAID storage system two ways: using the command-line interface in an IRIX shell or using *dassmgr*, Data General®'s software interface to the Challenge RAID storage system. The *dassmgr* program varies, depending on whether the storage system is running Sauna and Phoenix FLARE code.

The *dassmgr* program runs on an ASCII terminal attached to the serial port on the system. The program is also on diskette for use on your laptop; the diskette is available only in training classes for Challenge RAID.

Caution: This section contains information about *dassmgr*. The command line interface explained in Appendix B in this guide contains most of the functionality of *dassmgr*, omitting features not supported by Silicon Graphics. Use *dassmgr* only in circumstances when the command line interface is not available, that is, when you can access the Challenge RAID storage system only through the serial port. Ignore features of *dassmgr* that Silicon Graphics does not support, such as the non-mirrored cache configuration.

To reconfigure the system as system type 9, use an ASCII terminal. Follow these steps:

1. Run *raidcli getcache* to determine if caching is enabled. If it is enabled, make sure both storage-control processors are powered on; check the status lights or run *raidcli getcrus* (see Appendix B for information).
2. If caching is enabled, disable it by entering

   `raidcli -d` *device* **`setcache 0`**
3. Start *dassmgr*:
   - On a laptop: load *dassmgr* from the diskette.
   - On an ASCII terminal, *dassmgr* (actually the GridMgr screen of *dassmgr*) starts automatically when the ASCII terminal is powered on after being connected to the serial port on the Challenge RAID storage system.
4. If necessary, select the Grid Management option from the Disk Array Storage System Manager Main Menu.

   When you select the Grid Management option the first time after powering up the storage system, the Presentation Utility screen appears. Otherwise, the screen that appears is the one displayed the last time someone exited *dassmgr*. If the Main Menu screen is not replaced by another screen, press the caret (^) key several times.
5. Display the GridMgr Main Menu. If necessary, type a caret (**^**) and press **<Enter>**, repeating this process until it the GridMgr screen appears, as shown in Figure C-1.

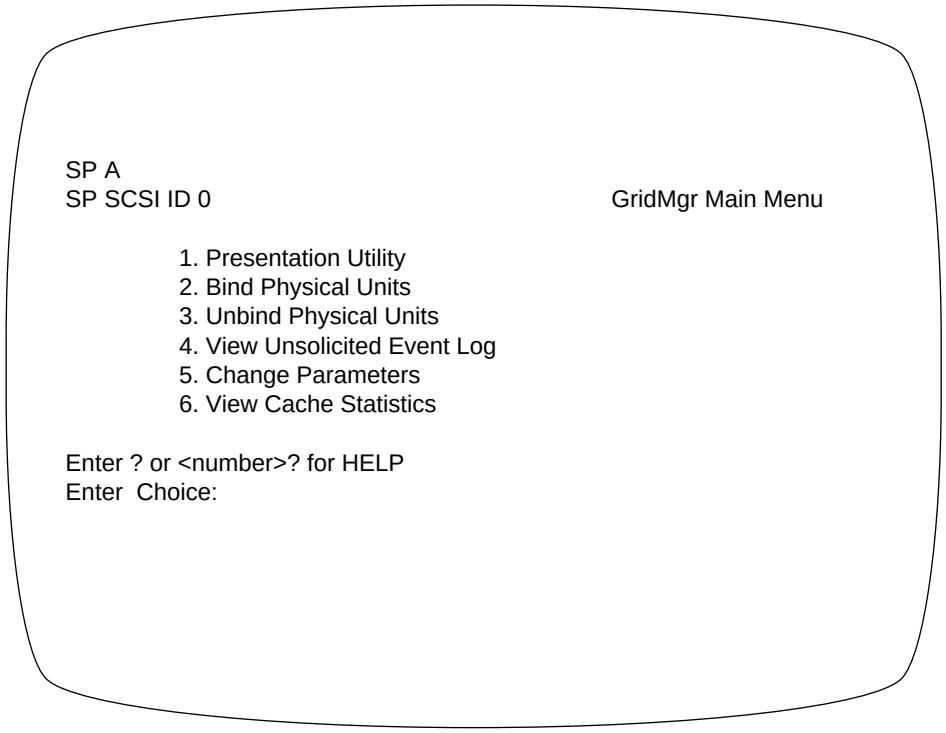

Figure C-1 GridMgr Main Menu

6. To get into engineering mode:
   - Press **<Ctrl-p>**. Be sure to use lowercase.
   - At the `Enter choice:` prompt, type **ENG** and press **<Enter>**. Be sure to use uppercase.

     This input does not appear on-screen.
   - Press **<Ctrl-j>**. Be sure to use lowercase.

   The letters `ENG` appear at the upper right corner of the GridMgr screen.
7. In the GridMgr screen, select 5, “Change Parameters.”
8. In the Change Parameters Menu, select 2, “Change Storage System Parameters.”
9. In the Change Storage System Parameters Menu, select “Select Storage System Options Package”; this choice is item 4 in the FLARE 7.xx (Sauna) version and item 5 in the FLARE 8.xx (both Sauna and Phoenix) version.
10. At the prompt

    ```
    Enter the system type (1-9) or ^ to return to previous screen

    Enter system type:
    ```

    type **9** and press **<Enter>**.
11. At the prompt

    ```
    Set system type and reboot controller(s) in system (Y/N)?
    ```

    type **y** and press **<Enter>**. After the reboot process is finished, the Presentation Utility appears.

12. In the Presentation Utility, check that the code for disk modules in the system are shown as ENA (enabled). Figure C-2 shows an example with two RAID-5 groups.

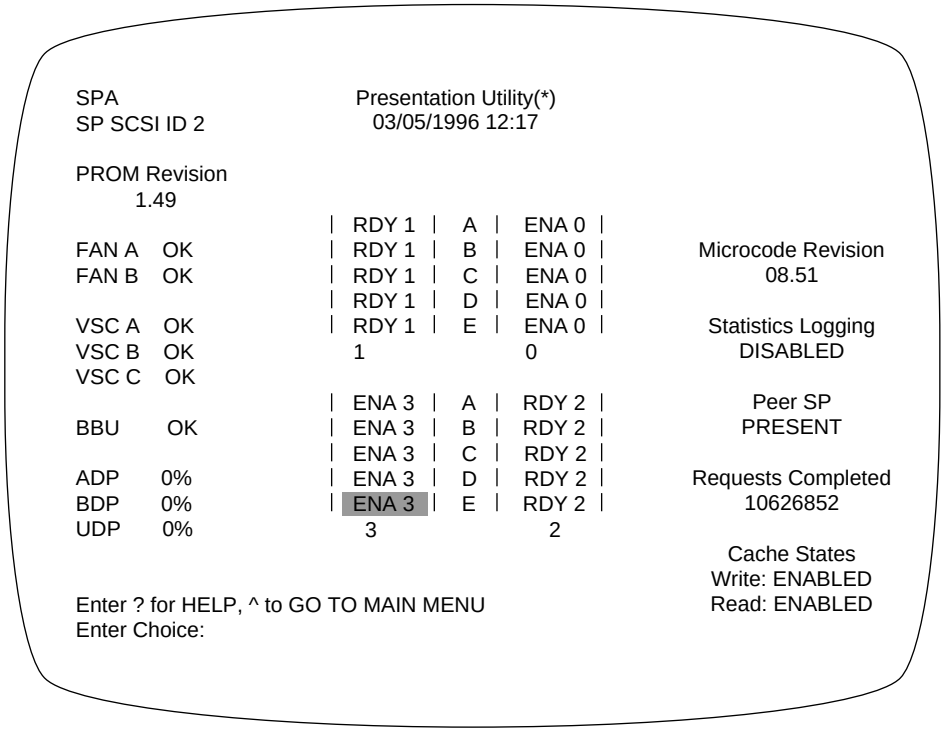

Figure C-2 Example Presentation Utility Screen

The FLARE code checks all drives. The fault indicator lights on any drive modules that are not Silicon Graphics drives. The agent will not start if any LUN is incorrect, that is, contains a drive that is not a Silicon Graphics drive.

Because the fault indicator light also indicates a drive fault for genuine Silicon Graphics drives, use *raid5 getdisk* to verify the problem with the disk drive.

13. Exit *dassmgr*:

    - On a laptop: press **<Esc>** to return to the Disk Array Storage System Manager Main Menu; select “Exit dassmgr.”
    - On an ASCII terminal: power off the terminal and disconnect it from the serial port on the Challenge RAID storage system.

C.4 Reenabling Command-Tagged Queuing

In the case of extremely sluggish performance, the volume header on one or more disk modules might have been inadvertently destroyed. To recreate the volume header, follow the steps in Section 4.10, “Enabling Command-Tagged Queuing,” in Chapter 4.

C.5 Determining Caching Problems

Disk modules in locations A0, B0, C0, D0, and E0 are vault drives. If a vault drive fails, caching is disabled until the failed module is replaced.

If an SP fails, caching is disabled until the failed SP is replaced.

C.6 Reassigning LUN Ownership

When an SP or SCSI-2 interface fails in a dual-SP configuration, such as the dual-interface / dual-processor configuration, the LUN ownership of the failing or unavailable SP is assigned to the surviving SP.

However, when the original SP is replaced or turned back on (at the end of the replacement procedure), all LUNs continue to be owned by the other SP. Although this situation is viable, it is not optimal. Performance in a dual-SP system is significantly improved if accesses to the LUNs are balanced across the two SPs, instead of all being funneled through one SP.

To restore the original status, use the *raidcli trespass* command (RAID agent 1.55; for earlier versions, use */sbin/trespass*) as follows. For example, if SP A was turned off and is now operational again, and SP A is connected to SCSI bus 6 with target (SCSI ID) 6, you would type:

```
/usr/raid5/raidcli trespass sc6d610 mine
```

This command does not require the RAID agent to be stopped. However, for SPs with RAID agent earlier than 1.55, use */sbin/trespass* as follows:

```
/etc/init.d/raid5 stop
/sbin/trespass sc6d610 mine
/etc/init.d/raid5 start
```

For */sbin/trespass*, the first command is required because the */sbin/trespass* command cannot work while the agent is running; the last command restarts the RAID agent. The */sbin/trespass* or *raidcli trespass* command tells SP A to take over the LUNs of which it is a default owner.

Follow this procedure whenever an SP (storage-control processor) of a dual-SP Challenge RAID storage system is turned off and then back on again.

Index

A

agent, 4-2 through 4-3
  error messages, C-2
  restarting, 4-11, 6-20, 6-22, C-2, C-6
  See also getagent
  starting, 4-2
  stopping, 4-11, 6-20, C-6
auto-reassign, 6-5 through 6-9

B

basic configuration, 2-11
  cabling, 3-9 through 3-11
battery backup unit, 2-10
  adding, 7-8 through 7-9
  location, 2-4, 2-5, 6-3, 6-4, 7-2, 7-3
  part number, 1-3
  replacing, 7-8 through 7-9
  seating, 3-6
  status lights, 2-10
BBU. *See* battery backup unit
bind, 4-11 through 4-14, B-2 through B-5
  legal configurations, 4-12, B-3
  planning, 4-5 through 4-6

C

cable
  attaching, 3-11
  part number, 1-1
cabling
  basic configuration, 3-9 through 3-11
  dual-bus / dual-initiator configuration, 3-15 through 3-17
  dual-interface / dual-processor, 3-13
  second CHALLENGE RAID, 3-20 through 3-23
  split-bus configuration, 3-14 through 3-15
cache, 4-17 through 4-20
  add-on, part number, 1-3
  disabling, 4-18, B-24
  enabling, 4-18, B-24
  flag in *bind*, 4-13, B-4
  flag in *chglun*, B-6
  hardware requirements, 2-3, 2-11
  page size, 4-18, B-25
  size, 4-18, B-25
  storage-control processor memory required, 2-6
Can't open device message, C-2
Cannot connect to agent message, C-2
CD-ROM, part number, 1-1
CHALLENGE server system, xiii
CHALLENGE S hookup, 3-12
chassis
  components, 2-3 through 2-10
  front view
    deskside, 2-1
    rack, 2-2
chassis assembly
  attaching to rack, 7-14
  installing additional, 7-13 through 7-14
chglun, B-5 through B-8
clearlog, B-8
clearstats, B-8
command-line interface, B-1 through B-27
command-tagged queuing, 4-14 through 4-17
components, 1-1, 1-2, 2-3 through 2-10
  replacement, 1-1, 1-3
configuration
  basic, 2-11
    cabling, 3-9 through 3-11
  comparison, 3-9
  disks, 4-1 through 4-20

dual-bus / dual-initiator, 2-13 through 2-14
cabling, 3-15 through 3-17
dual-interface / dual-processor, 2-12, C-6
cabling, 3-13
dual-SP, C-6
file, 4-2
split-bus, 2-13, 3-14
cabling, 3-14 through 3-15
system, 2-11 through 2-13
type 9, C-2 through C-5
conventions, xiv
CTQ. *See* command-tagged queuing

D

dassmgr, C-2 through C-5
device, 4-3, B-2
name, 4-4
node, creating, 5-15
node,creating, 4-1
disk module, 2-7 through 2-8
A0, B0, C0, A3
removing, 3-8
A0, B0, C0, D0, E0
failed, 5-5
and caching, 4-17, 5-4, C-6
removing, 3-8, 5-1, 5-12
add-on array, 5-12 through 5-15
installing, 5-12 through 5-15
order, 5-13
ordering, 5-12
configuring, 4-1 through 4-20
identifying failed, 5-4 through 5-5
installing replacement, 5-9 through 5-11
labeling, 3-7 through 3-8
location, 2-3, 3-8, 5-2, 5-14
part numbers, 1-1
removing failed, 5-6 through 5-9
seating, 3-8
status lights, 2-7, 5-4
unbinding, 4-4 through 4-5, 5-5
dual-bus / dual-initiator configuration, 2-13 through 2-14
cabling, 3-15 through 3-17
dual-interface / dual-processor configuration, 2-12, 3-13
cabling, 3-13
dual-SP configuration, C-6

E

electrostatic discharge damage (ESD), avoiding, 5-3
environment variable, 4-3, B-2
Equipment View, 6-7
error codes, B-18 through B-19
event log, B-18

F

failover, 6-6
RAID-3 memory requirement, 4-6, 4-8, 4-10, 4-11
See also high availability, IRIS FailSafe configuration
fan module, 2-9
closing, 6-2, 7-4
location, 2-4, 2-5, 6-3, 6-4, 7-2, 7-3
locking, 3-6, 3-27
opening, 6-2, 6-8, 7-4
part number, 1-3
replacing, 7-6 through 7-8
system operation when open, 6-8
unlocking, 3-6, 3-27
filesystem capacity and RAID, 4-5
filler board for storage-control processor, removing, 6-9 through 6-11
firmware, B-10
for fast RAID-3, 4-7
reboot, 4-10
See also FLARE code
updating after disk replacement, 5-11
firmware, 5-11, B-9
FLARE code
See also firmware
upgrading Sauna 7.xx to Sauna 8.xx, 6-18
upgrading Sauna to Phoenix, 6-19 through 6-20
front bezel, attaching, 7-14
FRU
location, 6-3, 6-4, 7-2, 7-3
deskside, 2-4
rack, 2-5
part numbers, 1-3
fx, 4-15 through 4-17

G

getagent, 4-4, B-10

getcache, B-11 through B-13
getcontrol, B-14
getcrus, B-15
getdisk, B-15 through B-17
getlog, B-18 through B-19
getlun, B-20 through B-24
getsp, 4-10

H

high availability, 2-13 through 2-14, 4-12, B-3
  See also failover
hot spare, binding, B-3

I

idle
  delay time, B-6
  threshold, B-6
inst, 4-2
IRIS FailSafe configuration, 2-13, 3-15, 6-5, 6-6
IRIX version required, xiii

K

kit contents, 1-1 through 1-3

L

LUN, 4-11 through 4-14
  ownership, 6-5
  restoring ownership, C-6
  See also disk module, *bind*, *unbind*

M

MAKE_VLUNS, 5-15
MAKEDEV, 4-1
memory module
  configurations, 6-13
  installing, 6-15 through 6-17
  part number, 1-3
  removing, 6-18
midplane, part number, 1-3
model number
  SP, B-24

O

Oracle Parallel Server configuration, 2-13, 3-15
owner's guide, part number, 1-1

P

Page Stream = chapter
  upgrading to LIC 9.X in storage system, 6-23
part numbers, 1-1 through 1-3
power cord, part number, 1-2
power distribution unit
  part number, 1-3
  replacing, 7-10 through 7-12
powering
  off, 3-6
  on, 3-25 through 3-27
power receptacle, part number, 1-3
power supply, 2-8 through 2-9
  adding, 7-5 through 7-6
  location, 2-4, 2-5, 6-3, 6-4, 7-2, 7-3
  part number, 1-3
  replacing, 7-5 through 7-6
  seating, 3-6
  status lights, 2-8
power switch location, 6-3, 7-2

R

rack
  adding chassis assembly, 7-13 through 7-14
  opening rear door, 3-5
  replacing power distribution unit, 7-10 through 7-12
  setting up, 3-2 through 3-4
RAID
  and filesystem capacity, 4-5
  level 1_0
    disk module order, 4-12, B-3
  level 3
    allocating memory, 4-10 through 4-11
    planning, 4-7 through 4-9
    setting memory, 4-19, B-25
    striping, 4-13
  levels
    planning, 4-5 through 4-6

raid5-agent.config, 4-2, C-2
RaidAgentDevice environment variable, 4-3, B-2
raidcli, B-1 through B-27
  parameters summarized, B-1
  parsing without calling API, 4-3
  syntax, 4-3, B-2
  verbose, 4-3
RAIDGUI, starting, 6-6
RAID levels, binding, B-3
rebuild time, 5-10 through 5-11, B-4
  bind, 4-13
  chglun, B-6
restarting
  agent after changing RAID-3 memory, 4-11
rip controller, part number, 1-3

S

Sauna
  upgrading 7.xx to 8.xx, 6-18
  upgrading to Phoenix, 6-19 through 6-23
SCSI-2
  adapter, 2-6
  bus length limit, 3-20, 3-23
  cable
    attaching, 3-11
    attaching. *See also* cabling
    internal I/O, part number, 1-3
    part number, 1-1
  connectors on CHALLENGE RAID chassis, 3-10
  ID
    SPs
      first CHALLENGE RAID, 3-18 through 3-19
      second CHALLENGE RAID, 3-25
    switches
      first CHALLENGE RAID, 3-18 through 3-19
      second CHALLENGE RAID, 3-25
  interface, 2-6
  internal I/O cable, part number, 1-3
  terminator plug, attaching, 3-11
setcache, B-24 through B-26
setstats, B-26
SIMM
  configurations, 6-13
  installing, 6-15 through 6-17
  part number, 1-3
  removing, 6-18
software, installing, 4-1 through 4-2
split-bus configuration, 2-13
  cabling, 3-14 through 3-15
SP. *See* storage-control processor
status light
  battery backup unit, 2-10
  disk module, 2-7
  fan module, 2-9
  power supply, 2-8
  storage-control processor
    ready, 2-6
    service, 2-6
storage-control processor, 2-6 through 2-7
  failed, and caching, 4-17, C-6
  filler board, removing, 6-9 through 6-11
  installing, 6-11 through 6-12
  location, 2-4, 2-5, 6-3, 6-4, 7-2, 7-3
  memory for caching, 2-6
  memory module
    installing, 6-15 through 6-17
    removing, 6-18
  minimum system memory, 2-6
  part number, 1-3
  powering off, 6-9
  powering on, 3-27, 6-12
  removing, 6-9 through 6-12
  replacing, 6-5 through 6-12
  seating, 3-6
  setting SCSI ID
    first CHALLENGE RAID, 3-18 through 3-19
    second CHALLENGE RAID, 3-25
  status lights, 2-6
stripe size, 4-13, B-4
striping
  count, 4-13, B-4
  RAID-3, 4-13
system type 9, reconfiguring as, C-2 through C-5

T

technical specifications, A-1 through A-3
terminator plug, part number, 1-2
trespass, B-26, C-6
troubleshooting, C-1 through C-6

U

unbind, 4-4 through 4-5, 5-5, B-27
user documentation, part number, 1-1

V

VSC. *See* power supply