IRIS FailSafe™ Installation and Maintenance Instructions

Document Number 108-0144-005

Contributors

Written by Carolyn Curtis
Illustrated by Cheri Brown, Dan Young, and Carolyn Curtis
Production by Michael Dixon
Engineering contributions by Paddy Sreenivasan, Michael Nishimoto, Chander Kant, Ajit Dandapani, and Ewan McKissock

FCC Warning

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

Attention

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

VDE 0871/6.78

This equipment has been tested to and is in compliance with the Level A limits per VDE 0871.

European Union Statement

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

CE

NRTL/C

TUV GS geprufte Sicherheit

International Special Committee on Radio Interference (CISPR)

This equipment has been tested to and is in compliance with the Class A limits per CISPR publication 22.

Canadian Department of Communications Statement

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

Attention

Le present appareil numerique n'emet pas de bruits radioelectriques depassant les limites applicables aux appareils numeriques de Classe A prescrites dans le Reglement sur le Brouillage Radioelectrique etabli par le Ministere des Communications du Canada.

Japanese Compliance Statement

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

IRIS FailSafe Installation and Maintenance Instructions
Document Number 108-0144-005

Silicon Graphics, Inc.
Mountain View, California

Contents

Figures

Tables

About This Guide

The Silicon Graphics IRIS FailSafe product provides a general facility for high- availability services. The IRIS FailSafe system is based on two Origin2000, Origin200, CHALLENGE, or POWER CHALLENGE deskside or rackmount servers* (or certain combinations), each offering services such as NFS and Netscape Enterprise Server. While running these services, the servers can also run database or other application software. Storage devices are physically attached to the two nodes in the system, but are owned and accessed by one node at a time. For this version of IRIS FailSafe (1.2), the servers must be running IRIX 6.2 or 6.4.

Note: For ease in reading, CHALLENGE is written as Challenge in the balance of this guide. The term "Challenge" also refers to POWER CHALLENGE servers except where noted.

Audience

The instructions in this guide are intended for Silicon Graphics System Support Engineers only. Because of the complexity of the IRIS FailSafe system, you must be familiar with the explanations of the system in the latest version of the *IRIS FailSafe Administrator's Guide* (007-3109-003 or later).

Structure of This Document

This guide contains the following chapters:

- Chapter 1, "IRIS FailSafe System Components," lists the hardware upgrades, software, cables, and options included in the IRIS FailSafe system.
- Chapter 2, "Setting Up and Cabling the IRIS FailSafe System With Origin2000 and Origin200 Servers," explains how to install and set up an IRIS FailSafe system with Origin2000 or Origin200 servers.

---

* Onyx and Onyx2 are also supported as FailSafe nodes in the same combinations as Challenge L and Origin2000 servers, respectively.

- Chapter 3, “Setting Up and Cabling the IRIS FailSafe System With Large Challenge Servers,” describes how to install and set up Challenge DM, L, and XL servers and shared storage systems and how to cable them for an IRIS FailSafe system.
- Chapter 4, “Setting Up and Cabling the IRIS FailSafe System With Challenge S Servers,” explains how to install the hardware for an IRIS FailSafe system in which one or both servers are Challenge S systems.
- Chapter 5, “Setting Up and Cabling Mixed Configurations (Challenge and Origin Servers),” explains cabling an IRIS FailSafe system in which one IRIS FailSafe host is an Origin2000 or Origin200 deskside or rackmount server and the other is a Challenge S, DM, L, or XL server.
- Chapter 6, “Configuring Challenge RAID Storage Systems for IRIS FailSafe,” explains distribution of disks on RAID buses, the RAID configuration file, LUNs, and SCSI device nodes with respect to IRIS FailSafe.
- Chapter 7, “Maintaining and Upgrading the High-Availability System,” explains troubleshooting system problems, isolating a node, replacing disks on the Challenge RAID or Fibre Channel storage system, and other maintenance tasks.

An index completes this guide.

Other Required Documentation

This section lists Silicon Graphics documents that are required for installing the IRIS FailSafe system. Be completely familiar with them, or bring the latest versions with you to the installation site.

Note: The final digits of the document numbers in the following list might not be the latest versions. Many current installation guides are available in PostScript form from *guest@comrade.wpd:/usr/people/guest/docdist*. On the Web, you can access the Technical Publications library:

inside the firewall:
http://techpubs.engr.sgi.com/library/

outside the firewall:
http://techpubs.sgi.com/library/

- *IRIS FailSafe Administrator’s Guide* (007-3109-003)

  Note: You must have the latest version of this manual for software installation and configuration instructions; these instructions are only in the *IRIS FailSafe Administrator’s Guide* and are not included in this manual.

- *Site Preparation for Origin Family, Onyx2, OCTANE, and O2* (007-3452-002)
- *Origin2000 and Onyx2 Deskside and Rackmount Installation Instructions* (108-0155-003)
- *Origin2000 Rackmount Owner’s Guide* (007-3456-003)
- *Origin2000 Deskside Server Owner’s Guide* (007-3453-002)
- *Origin200 and Origin Vault Installation Instructions* (108-0153-002)

- *O2 Workstation Hardware Reference Manual* (007-3275-002)
- *CHALLENGE/Onyx Site Preparation Guide* (108-7040-040)
- *CHALLENGE/Onyx XL Rackmount Installation Instructions* (108-7042-020; contains SCSI channel information in Appendix D)
- *CHALLENGE/Onyx Deskside Installation Instructions* (108-7039-020)
- *CHALLENGE Vault Rack and SCSIBox 2 Installation Instructions* (108-7044-040)
- *CHALLENGE Vault L Installation Instructions (*108-0124-001*)*
- *CHALLENGE S Server Owner's Guide* (007-2314-005)
- *Ultra SCSI XIO Board Installation Instructions* (108-0157-001)
- *IRIS 4-Port Fast Ethernet Adapter with Asynchronous Serial XIO Board Installation Instructions* (108-0151-001)
- *Fast Ethernet PCI Option Installation Instructions* (007-3535-001)
- *100Base-T VME Board Installation Instructions* (108-0148-001)
- *FDDIXPress User's Guide for the PCI Local Bus* (007-3447-001)
- *FDDIXpress Mezzanine Board for CHALLENGE and Onyx Installation Instructions* (108-0116-002)
- *FDDIXpress Administration Guide* (007-0813-050; included if FDDI components are ordered)
- *FDDIXpress Release Notes,* included in FDDI board shipment
- *CHALLENGE RAID Installation and Maintenance Instructions* (108-0128-006)
- *Origin FibreVault and Fibre Channel RAID Installation Instructions* (108-0154-002)
- *Remote System Control Port Installation Guide* (108-0140-001)

In addition to the manuals listed above, the following manuals are useful for maintaining or troubleshooting the IRIS FailSafe system; bring them to a site requiring maintenance:

- *CHALLENGE RAID Owner's Guide* (007-2532-006 or later; included with optional Challenge RAID)
- IRIX NetWorker documentation
- *ONC3/NFS Administrator's Guide* (007-0850-090)
- *IRIX Admin: Disks and Filesystems* volume, which covers XLV volumes for IRIX 6.2 (007-2825-002)

Customer manuals shipped with the IRIS FailSafe system are

- *IRIS FailSafe Administrator's Guide* (007-3109-003)
- *IRIS FailSafe Programmer's Guide* (007-3298-001)
- administrator's guides for software options that the customer has purchased

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
“”	(Double quotation marks) Onscreen menu items and references in text to document section titles
[]	(Brackets) Surrounding optional syntax statement arguments

Chapter 1

IRIS FailSafe System Components

The Silicon Graphics IRIS FailSafe system consists of the following hardware:

- two servers—deskside or rackmount Challenge or POWER Challenge S, DM, L, or XL, or deskside or rackmount Origin2000 or Origin200:
  - two Challenge or POWER Challenge S, DM, L, or XL in any combination
  - two Origin2000 or Origin200 in any combination
  - one Origin2000 and one Challenge or POWER Challenge DM, L, or XL
  - one Origin200 and one Challenge or POWER Challenge S, DM, L, or XL

  Note: The combination Origin2000 and Challenge S is not a supported IRIS FailSafe configuration.
- shared storage (with correct interface cards in the servers):
  - Challenge Vault XL, L, or DM
  - Challenge RAID deskside or rackmount storage system; each chassis assembly has two storage-control processors (SPs) and at least five disk modules, caching enabled
  - for Origin2000 or Origin200 only: Origin Vault differential deskside or rackmount
  - for Origin2000 or Origin200 only: Fibre Channel storage options, either RAID or FibreVault (JBOD: “just a bunch of disks”), either rackmount or tower

  Note: Many other restrictions apply besides those given above. Some configurations require specific patches, firmware versions, versions of storage options, or versions of interface boards. For the latest supported configuration information, see

  http://origin.engr.sgi.com/product/hiavailability/support.html
- required hardware upgrades and cables
- Ethernet or FDDI networking adapters and facilities

The software for this release (1.2) of the IRIS FailSafe system includes graphical user interface (GUI) software. You must also access and install IRIS FailSafe rollup patches and the high-end recommended patch set. In addition, you install the software included with any attached peripherals, and applicable patches for them. For information on recommended patches for each platform, see

http://bits.csd.sgi.com/digest/patches/recommended/

Besides the basic configuration, the IRIS FailSafe 1.2 product is available with an NFS, Web, Oracle, Informix, Sybase, WebFORCE, WebFORCE MediaBase, or Gauntlet server option; other options are scheduled for later release. For information about configurations, see

http://origin.engr.sgi.com/product/hiavailability/support.html

IRIS FailSafe enhances the Silicon Graphics Oracle Parallel Server (OPS) by providing IP failover in an OPS hardware configuration. However, the two products are not merged administratively, so different tools are required to maintain a combined system.

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

Table 1-1 through Table 1-8 list IRIS FailSafe system kit components.

Note: These tables do not include the customer documentation or the envelope containing the SSE manuals. For the 40-foot cable options, see Table 1-9.

Table 1-1 lists components for symmetric Origin2000 to Origin2000 configurations (marketing code FailSafe-2000C-1.2).

Table 1-1 Origin2000 to Origin2000 (FAILSAFE-2000C-1.2)

Component	Quantity	Part Number
Ethernet cable assembly RJ45-RJ45 (null modem), 20-foot	1	9290131
Differential SCSI cable, 20-foot	2	9290111
Cable assembly, 8-pin mini-DIN to PC9 standard pinout, 20-foot	2	018-0668-001
Cable assembly, PC9 to PC9 standard pinout, 20-foot	2	018-0691-001
CD PACK Failsafe 1.2	1	SC4-FSAFE-1.2
IRIS FailSafe GUI configuration CD 1.2	1	SC4-FSAFE-GUI-1.2
License: XFS Volume Plexing RTU 2.0	2	SR4-PLEX-2.0

Table 1-2 lists components for symmetric Origin200-to-Origin200 configurations (marketing code FailSafe-200C-1.2).

Table 1-2 Origin200 to Origin200 (FAILSAFE-200C-1.2)

Component	Quantity	Part Number
Ethernet cable assembly RJ45-RJ45 (null modem), 20-foot	1	9290131
Differential SCSI cable, 20-foot	2	9290111
Cable assembly, 8-pin mini-DIN to PC9 standard pinout, 20-foot	2	018-0668-001
CD PACK Failsafe 1.2	1	SC4-FSAFE-1.2
IRIS FailSafe GUI configuration CD 1.2	1	SC4-FSAFE-GUI-1.2
License: XFS Volume Plexing RTU 2.0	2	SR4-PLEX-2.0

Table 1-3 lists components for asymmetric configurations with an Origin2000 server and an Origin200 server or Challenge XL / L / DM server (marketing code FailSafe-2000M-1.2).

Table 1-3 Origin2000 to Origin200 or Challenge XL / L / DM (FAILSAFE-2000M-1.2)

Component	Quantity	Part Number
Ethernet cable assembly RJ45-RJ45 (null modem), 20-foot	1	9290131
Differential SCSI cable, 20-foot	2	9290111
Cable assembly, 8-pin mini-DIN to PC9 standard pinout, 20-foot	2	018-0668-001
Cable assembly, PC9 to PC9 standard pinout, 20-foot	1	018-0691-001
Cable assembly, PC9 to SGI 9-pin, 20-foot	2	018-0669-001
Cable assembly, 8-pin mini-DIN to SGI 9-pin, 20-foot	1	018-0690-001
CD PACK Failsafe 1.2	1	SC4-FSAFE-1.2
IRIS FailSafe GUI configuration CD 1.2	1	SC4-FSAFE-GUI-1.2
License: XFS Volume Plexing RTU 2.0	2	SR4-PLEX-2.0

Table 1-4 lists components for Origin200 to Challenge XL / L / DM configurations (marketing code FailSafe-200M-1.2).

Table 1-4 Origin200 to Challenge XL / L / DM (FAILSAFE-200M-1.2)

Component	Quantity	Part Number
Ethernet cable assembly RJ45-RJ45 (null modem), 20-foot	1	9290131
Differential SCSI cable, 20-foot	2	9290111
Cable assembly, PC9 to SGI 9-pin, 20-foot	1	018-0669-001
Cable assembly, 8-pin mini-DIN to SGI 9-pin, 20-foot	1	018-0690-001
CD PACK Failsafe 1.2	1	SC4-FSAFE-1.2
IRIS FailSafe GUI configuration CD 1.2	1	SC4-FSAFE-GUI-1.2
License: XFS Volume Plexing RTU 2.0	2	SR4-PLEX-2.0

Table 1-5 lists components for Origin200 to Challenge S configurations (marketing code FailSafe-200S-1.2).

Table 1-5 Origin200 to Challenge S (FAILSAFE-200S-1.2)

Component	Quantity	Part Number
Remote power control unit	1	9110120
Ethernet cable assembly RJ45-RJ45 (null modem), 20-foot	1	9290131
Differential SCSI cable, 20-foot	2	9290111
Cable assembly, PC9 to RJ45	1	018-0667-001

Table 1-5 (continued) Origin200 to Challenge S (FAILSAFE-200S-1.2)

Component	Quantity	Part Number
Cable assembly, 8-pin mini-DIN to 8-pin mini-DIN	1	018-8223-001
CD PACK Failsafe 1.2	1	SC4-FSAFE-1.2
IRIS FailSafe GUI configuration CD 1.2	1	SC4-FSAFE-GUI-1.2
License: XFS Volume Plexing RTU 2.0	2	SR4-PLEX-2.0

Table 1-6 lists components for Challenge S to Challenge S systems (marketing code FailSafe-SS-1.2).

Table 1-6 Challenge S to Challenge S Configurations (SS Kit)

Component	Quantity	Part Number
Remote power control unit	1	9110120
Ethernet cable assembly RJ45-RJ45 (null modem), 20 -foot	1	9290131
SCSI cable, 20-foot	2	9290111
CD PACK Failsafe 1.2	1	SC4-FSAFE-1.2
IRIS FailSafe GUI configuration CD 1.2	1	SC4-FSAFE-GUI-1.2
License: XFS Volume Plexing RTU 2.0	2	SR4-PLEX-2.0

Table 1-7 lists components for symmetric or asymmetric all Challenge systems (Challenge DM, L, or XL in any combination) that do not include Challenge S (marketing code FailSafe-LL-1.2).

Table 1-7 Symmetric or Asymmetric Challenge Configurations Not Including a Challenge S Server (LL Kit)

Component	Quantity	Part Number
Ethernet cable assembly RJ45-RJ45 (null modem), 20-foot	1	9290131
Cable assembly 9PDSUB-9PDSUB (null modem serial reset), 20-foot	1	018-0552-001
SCSI cable, 20-foot	2	9290111
CD PACK Failsafe 1.2	1	SC4-FSAFE-1.2
IRIS FailSafe GUI configuration CD 1.2	1	SC4-FSAFE-GUI-1.2
License: XFS Volume Plexing RTU 2.0	2	SR4-PLEX-2.0

Table 1-8 lists components for asymmetric IRIS FailSafe systems consisting of a Challenge S and any other Challenge server (marketing code FailSafe-LS-1.2).

Table 1-8 Asymmetric Configurations Including a Challenge S Server (LS Kit)

Component	Quantity	Part Number
Remote power control unit	1	9110120
Ethernet cable assembly RJ45-RJ45 (null modem)	1	9290131

Table 1-8 (continued) Asymmetric Configurations Including a Challenge S Server (LS

Component	Quantity	Part Number
SCSI cable, 20-foot	2	9290111
Cable assembly 9PDSUB-8PMINI (null modem serial reset)	1	018-0542-001
CD PACK Failsafe 1.2	1	SC4-FSAFE-1.2
IRIS FailSafe GUI configuration CD 1.2	1	SC4-FSAFE-GUI-1.2
License: XFS Volume Plexing RTU 2.0	2	SR4-PLEX-2.0

The customer might have ordered one or more of the options listed in Table 1-9.

Table 1-9 Options for the IRIS FailSafe System

Component	Quantity	Part Number/Marketing Code
IRIS FailSafe NFS 1.2		SC4-FSAFE-NFS-1.2
IRIS FailSafe GUI Configurator 1.2		SC4-FSAFE-GUI-1.2
IRIS FailSafe WebFORCE 1.2		SC4-FSAFE-WEB-1.2
IRIS FailSafe ORACLE 1.2		SC4-FSAFE-ORCL-1.2
IRIS FailSafe INFORMIX 1.2		SC4-FSAFE-IFMX-1.2
IRIS FailSafe Sybase 1.2		SC4-FSAFE-SYBS-1.2
IRIS FailSafe Gauntlet 1.2		SC4-FSAFE-GLET-1.2
IRIS FailSafe WebFORCE MediaBase 1.2		SC4-FSAFE-MBAS-1.2
40-foot cable option for two Origin2000 servers:		X-FSAFE-40FT-2000C
Cable assembly, 8-pin mini-DIN to PC9 standard pinout, 40-foot	2	018-0668-101
Cable assembly, PC9 to PC9 standard pinout, 40-foot	2	018-0691-101
Ethernet cable assembly RJ45-RJ45 (null modem), 40-foot	1	9290132
40-foot cable option for two Origin200 servers:		X-FSAFE-40FT-200C
Cable assembly, 8-pin mini-DIN to PC9 standard pinout, 40-foot	*2*	018-0668-101
Ethernet cable assembly RJ45-RJ45 (null modem), 40-foot	*1*	9290132
40-foot cable option for mixed configurations with Origin2000 server, version 1.2:		X-FSAFE-40FT-2000M
Cable assembly, 8-pin mini-DIN to PC9 standard pinout, 40-foot	*2*	018-0668-101
Cable assembly, PC9 to SGI 9-pin, 40-foot	*2*	018-0669-101
Cable assembly, 8-pin mini-DIN to SGI 9-pin, 40-foot	*1*	018-0690-101
Cable assembly, PC9 to PC9 standard pinout, 40-foot	*1*	018-0691-101
Ethernet cable assembly RJ45-RJ45 (null modem), 40-foot	*1*	9290132
40-foot cable option for mixed configurations with Origin200 server:		X-FSAFE-40FT-200M
Cable assembly, PC9 to SGI 9-pin, 40-foot	*1*	018-0669-101
Cable assembly, 8-pin mini-DIN to SGI 9-pin, 40-foot	*1*	018-0690-101
Ethernet cable assembly RJ45-RJ45 (null modem), 40-foot	*1*	9290132
40-foot cable option for Challenge only configurations (no Challenge S):		X-FSAFE-40FT-CHAL
40-foot SCSI cable 9PDSUB-8PMINI, 40-foot	2	018-0553-001
Ethernet cable assembly RJ45-RJ45 (null modem), 40-foot	1	9290132

Optional hardware includes servers and storage, as well as cables of different lengths from those supplied with the IRIS FailSafe option.

Chapter 2

Setting Up and Cabling the IRIS FailSafe System With Origin2000 and Origin200 Servers

This chapter explains how to set up and cable an IRIS FailSafe system in which at least one server (IRIS FailSafe host) is a deskside or rackmount Origin2000 server or Origin200 server.* It also explains cabling the shared storage—Origin Vault, Challenge RAID, Challenge vault, or Fibre Channel storage option—for use as the storage for the IRIS FailSafe system.

This chapter consists of these sections:

- Section 2.1, “Installing the Software”
- Section 2.2, “Setting Up the Component Systems”
- Section 2.3, “Installing Interface Boards”
- Section 2.4, “Cabling the Private and Public Networks”
- Section 2.5, “Setting Up the Serial Connection”
- Section 2.6, “Testing the Serial Connection”
- Section 2.7, “Setting Automatic Power-On for Origin200 Servers”
- Section 2.8, “Cabling the Challenge or Origin Vaults”
- Section 2.9, “Cabling the Challenge RAID Storage System to the Origin Servers”
- Section 2.10, “Cabling the Fibre Channel Storage Options”
- Section 2.11, “Setting the IRIS FailSafe Host SCSI IDs”

Note: Before installing an IRIS FailSafe system, make sure that the installation site meets the operating limits and AC power requirements as explained in *Site Preparation for Origin Family and Onyx2* and, if applicable, the *CHALLENGE/Onyx Site Preparation Guide*, chapters 3 and 4.

The following equipment is required for installation:

- installation guides for the component systems (see “About This Guide” for part numbers)

---

* Onyx2 systems are also supported as FailSafe nodes in the same combinations as Challenge L and Origin2000 servers, respectively.

  - *Site Preparation for Origin Family and Onyx2*
  - *Origin2000 and Onyx2 Deskside and Rackmount Installation Instructions*
  - *Origin200 and Origin Vault Installation Instructions*
  - *Challenge RAID Maintenance and Installation Instructions*
  - *Origin FibreVault and Fibre Channel RAID Installation Instructions*
  - manuals for interface boards (FDDI, Ethernet, SCSI); see "About This Guide" for names and part numbers
- laptop, ASCII terminal, or IRISconsole
- Phillips and small flat-blade screwdrivers

2.1 Installing the Software

Follow instructions in the latest version of the *IRIS FailSafe Administrator's Guide* (007-3109-003 or later) to install the software needed to run IRIS FailSafe.

For this version of IRIS FailSafe (1.2), the Origin servers must be running IRIX 6.2 or 6.4.

Depending on the servers and storage in the configuration and the IRIX revision level, you might need to install various patches.

- For the latest information on server and storage compatibility, see

  http://origin.engr.sgi.com/product/hiavailability/support.html
- For information on recommended patches for each platform, see

  http://bits.csd.sgi.com/digest/patches/recommended/

For IRIS FailSafe operation, you must determine the MMSC and MSC passwords on the Origin system used as an IRIS FailSafe host, if the system administrator has changed them from the default. This password is used when one node resets the other. To specify this password so that the IRIS FailSafe software knows about it, use

ha_spng -w *password* **-d** *dst_sysctlr_type*

For more information on MSC passwords, see the *Origin2000 and Onyx2 Deskside and Rackmount Installation Instructions*, the *Origin200 and Origin Vault Installation Instructions*, or the latest version of the *IRIS FailSafe Administrator's Guide*. For information on ha_spng(1M), see its man page.

2.2 Setting Up the Component Systems

This section consists of

- Section 2.2.1, "Setting Up the Hardware"
- Section 2.2.2, "Checking the Grounding in Configurations Using Fibre Channel Storage"

- Section 2.2.3, “Planning the Connections Between IRIS FailSafe Hosts”

2.2.1 Setting Up the Hardware

Read through *Site Preparation for Origin Family and Onyx2* before unpacking equipment at the site. Follow instructions in that manual, specifically:

1. Make sure that the installation site meets the operating limits and AC power requirements for the hardware.
2. Make sure required tools and personnel are on hand for unloading and opening crates and for moving large systems.
3. Prepare the physical location to allow for space, air flow, and floor-loading requirements for all component systems.

   Use only the cables included in the shipment. Plan to situate the servers and vaults fairly close together. Differential Fast-20 SCSI (installations with Origin family only), including cabling inside the chassis, should be no longer than 25 meters (82 feet). Differential SCSI-2 cables, including cabling inside the chassis, should be no longer than 18.3 meters (60 feet).
4. Make sure the site meets safety and operating considerations.
5. Prepare the site for the systems’ electrical requirements.
6. Connect your laptop or ASCII terminal and keyboard to the first IRIS FailSafe host, as explained in the server installation instructions.
7. Connect the power cord for the first IRIS FailSafe host. Turn on the main power switch on the back of the unit; turn on the laptop or ASCII terminal. If necessary, see the server installation instructions for details.
8. Turn on the IRIS FailSafe host, following instructions in its manual. When power-on diagnostics are completed, a login prompt appears on the console.
9. Enter **`hinv`** on each IRIS FailSafe host.

2.2.2 Checking the Grounding in Configurations Using Fibre Channel Storage

Grounding is very important in Origin family and Onyx2 systems. Each chassis must be well grounded through its power connector. All chassis connected by XIO copper cables must share the same transformer, be grounded through the same earthing rod, and be on the same branch circuit. If you have any doubts about the quality of the ground connection, consult a qualified electrician.

Use of an optical cable between the fibre enclosure(s) and the host XIO connection eliminates any problems related to common grounding.

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

The branch circuit wiring must have an insulated grounding conductor that is identical in size, insulation material, and thickness to the earthed and unearthed branch-circuit supply conductors. The grounding conductor should be green, with or without one or more yellow stripes. This grounding or earthing conductor should be connected to earth at the service equipment or, if supplied by a separately derived system, at the supply transformer or motor-generator set.The power receptacles in the vicinity of the systems should all be of an earthing type, and the grounding or earthing conductors serving these receptacles should be connected to earth at the service equipment.

2.2.3 Planning the Connections Between IRIS FailSafe Hosts

For the IRIS FailSafe system, you are setting up these networks on each server:

- public network interface(s)

  This network connects the IRIS FailSafe server to clients and the outside world. The I/O panel Ethernet port or one from an option board can be used.

  Consult with the customer on the number and type of public network interface(s) for each server. If a server has two public network interfaces, they must be on different networks.

- serial connection between nodes: the MSC (Origin2000 deskside system controller), **AUX** (Origin200 deskside system) or MMSC **ALTERNATE CONSOLE** port (Origin2000 rackmount module) is cabled to the tty_2 port on the other node

  If one node fails, this connection enables the surviving node to power cycle the other.

- private (heartbeat) network (Ethernet)

  This network, going only between the two nodes, is used exclusively for the keep-alive heartbeat that the two nodes use for monitoring each other's status and IRIS FailSafe control messages. The I/O panel Ethernet port or one from an option board can be used.

- shared fibre channel or differential SCSI connection to storage

Figure 2-1 diagrams an example IRIS FailSafe system with two deskside Origin2000 servers as hosts. Note the four types of connections.

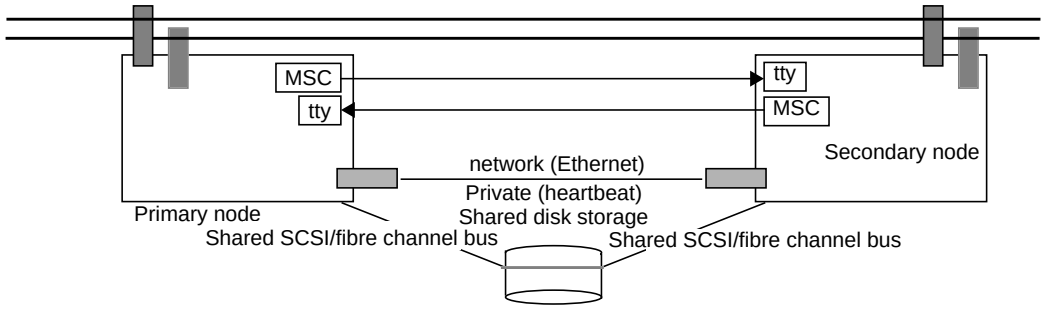

Figure 2-1 Two Deskside Origin2000 Servers

Figure 2-2 diagrams an example IRIS FailSafe system with one deskside Origin2000 server and one rackmount Origin2000 server as hosts.

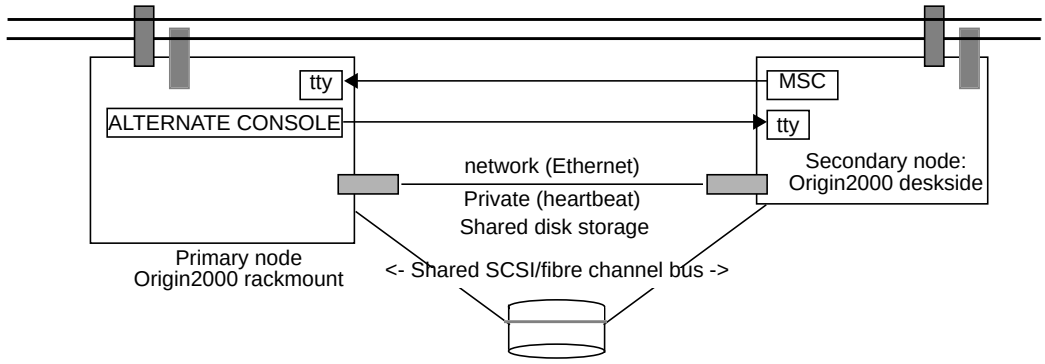

Figure 2-2 Origin2000 Rackmount and Origin2000 Deskside Servers

Figure 2-3 diagrams an example IRIS FailSafe system with two rackmount Origin2000 servers. In this configuration, the serial connection between the **ALTERNATE CONSOLE** ports on the Origin Rack MMSCs and the tty ports enables one server to power-cycle the other server in case of failure. For configurations with multiple Origin Racks, this connection is between the first (master) rack in the series. The servers in multiple-rack configurations must be set up as one system.

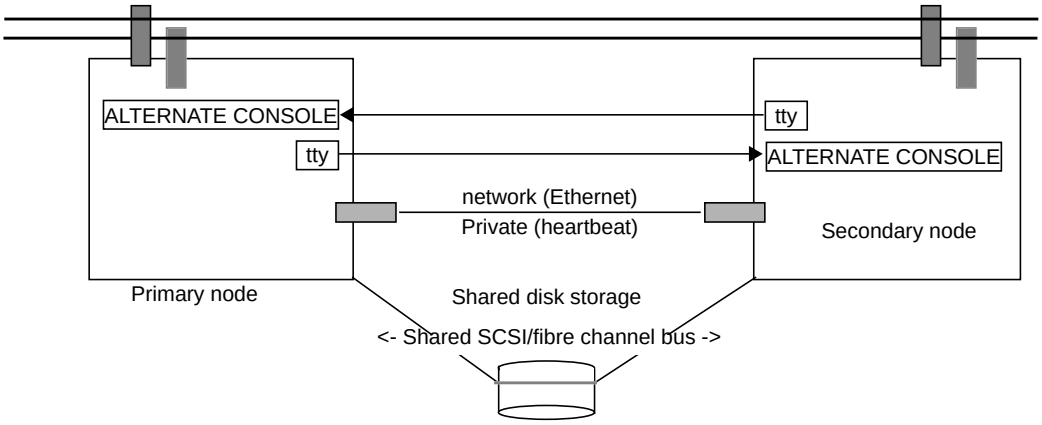

Figure 2-3 Two Origin2000 Rackmount Servers

Figure 2-4 diagrams an example IRIS FailSafe system with two Origin200 servers as hosts.

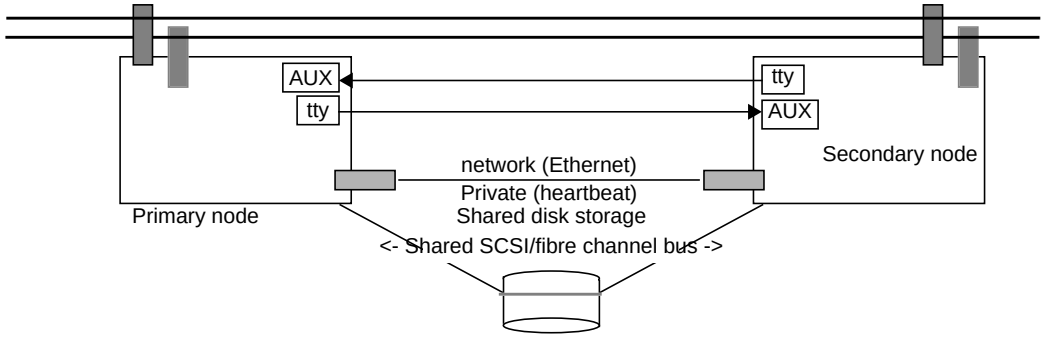

Figure 2-4 Two Origin200 Servers

Figure 2-5 diagrams an example IRIS FailSafe system with one Origin200 server and one rackmount Origin2000 server as hosts.

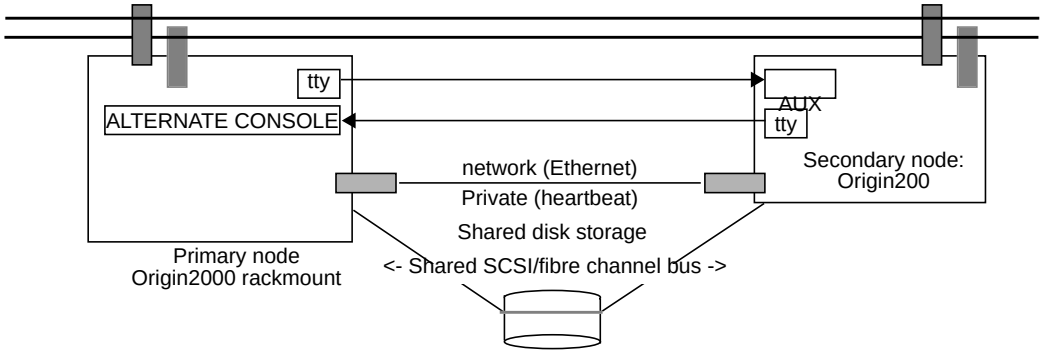

Figure 2-5 Origin2000 Rackmount and Origin200 Server

2.3 Installing Interface Boards

Before you cable the networks, install any required interface boards:

- Ethernet:

  The Origin2000 IO6 panel has one 10-Base-T / 100-Base-T Ethernet port. The optional ENET XIO board holds four additional Fast Ethernet ports For information on installing the ENET board, see the *IRIS 4-Port Fast Ethernet with Serial XIO Board Installation Instructions*.

  The Origin200 has one 10-Base-T / 100-Base-T Ethernet port. The optional ENET PCI board holds an additional Fast Ethernet port. For information on this board, see the *Fast Ethernet PCI Option Installation Instructions*.

- FDDI: If the customer is using the FDDI PCI board for the private network, see the *Origin200 and Origin Vault Installation Instructions* or the Origin2000 and Onyx2 installation manuals for instructions on installing PCI boards in the PCI module or slots.

  Note: As of this writing, no FDDI XIO board is available from Silicon Graphics, although one is planned for calendar 1998.

- SCSI: If the customer has ordered an Ultra SCSI XIO or PCI board, see
  - XIO board: *Ultra SCSI XIO Board Installation Instructions*
  - PCI board: *Origin200 and Origin Vault Installation Instructions*
- Fibre Channel: If the customer has ordered a Fibre Channel XIO or PCI board, see
  - XIO board: *Origin FibreVault and Fibre Channel RAID Installation Instructions*
  - PCI board: *Origin200 and Origin Vault Installation Instructions*

2.4 Cabling the Private and Public Networks

This section explains network cabling:

- Section 2.4.1, “Cabling the Private Network”
- Section 2.4.2, “Setting Up a Public Network Connection”

Note: IRIS FailSafe can fail over IP addresses from one public network interface to another interface of the same type.

2.4.1 Cabling the Private Network

The private network between the servers supplies the heartbeat of each server to the other. This section explains cabling the private network for various combinations of servers. Figure 2-6 shows the Ethernet connector on the Origin2000 IO6 panel. The server might also have one or more ENET boards.

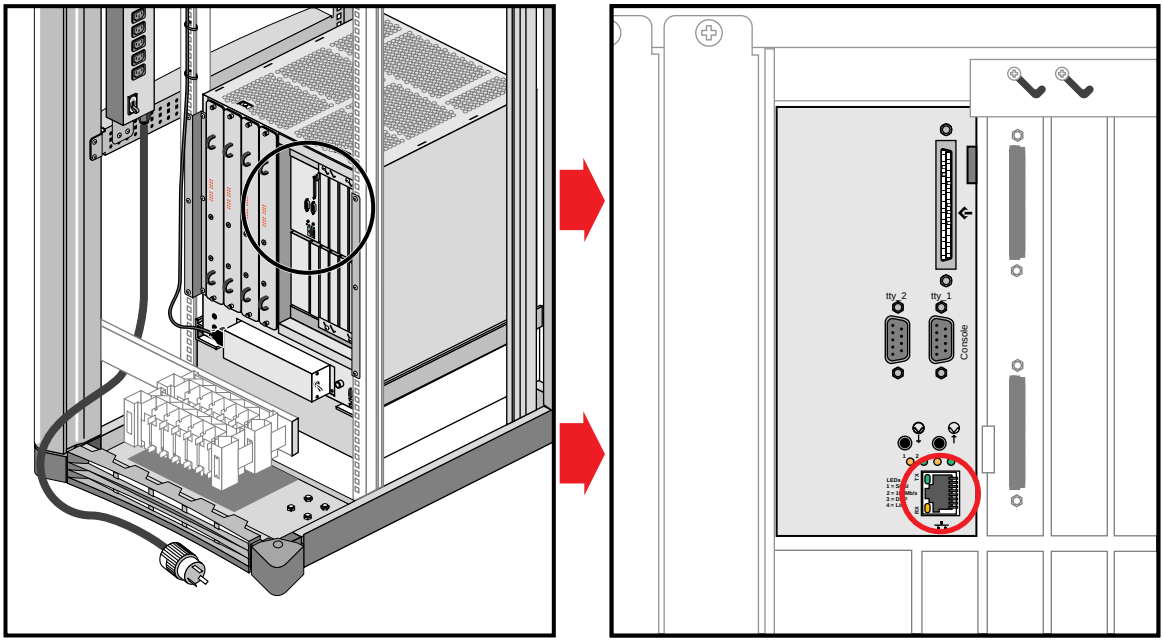

Figure 2-6 Origin2000 Rackmount Server Ethernet Connector

The Origin200 panel has one 10-Base-T / 100-Base-T Ethernet port. The optional ENET PCI board holds an additional Fast Ethernet port. Figure 2-7 shows Ethernet ports on the rear of the Origin200 server.

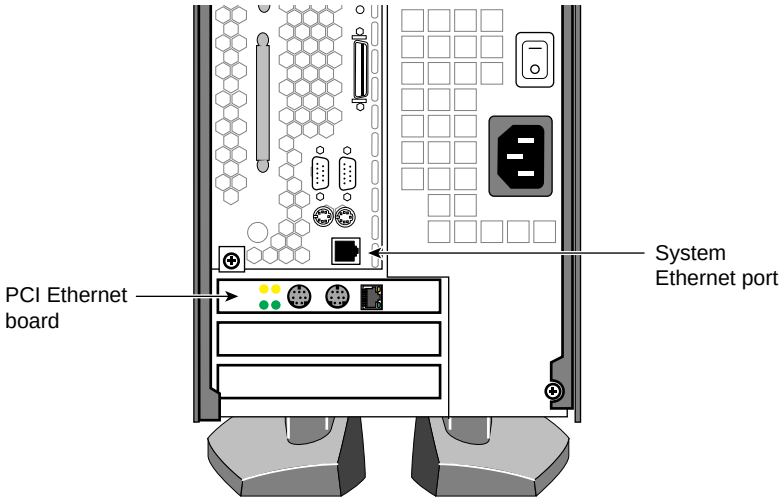

Figure 2-7 Origin200 Deskside Server Ethernet Ports

To cable the private network, attach an end of the null modem Ethernet cable supplied with the IRIS FailSafe system (9290131, 20-foot, or 9290132, 40-foot) to an Ethernet port on each host module.

2.4.2 Setting Up a Public Network Connection

On each server, connect the public network drop cable to an Ethernet or FDDI port.

2.5 Setting Up the Serial Connection

If one server fails, a serial cable enables the other server to power-cycle the failed server. This serial cable connects one server's system controller port to a serial (tty) port on the other server's I/O board panel. In each IRIS FailSafe system, two such serial cables are required, so that each server can power-cycle the other.

This section explains this cabling for various combinations of servers:

- Section 2.5.1, "Ports for the Serial Connection"
- Section 2.5.2, "Cabling the Serial Connection"

2.5.1 Ports for the Serial Connection

Table 2-1 lists the ports to use on each type of server.

Table 2-1 Ports for Serial Connection

Server	Controller (Serial) Port	Serial Port (9-Pin Sub-D)
Origin2000 deskside	MSC serial port on rear (9-pin sub-D)	tty_2 on IO6 board
Origin2000 rackmount	MMSC **ALTERNATE CONSOLE** port; 8-pin mini-DIN	tty_2 on IO6 board
Origin200	**AUX** port on rear (8-pin mini-DIN)	tty_2 on rear

Figure 2-8 shows the serial port on the rear of the Origin2000 module.

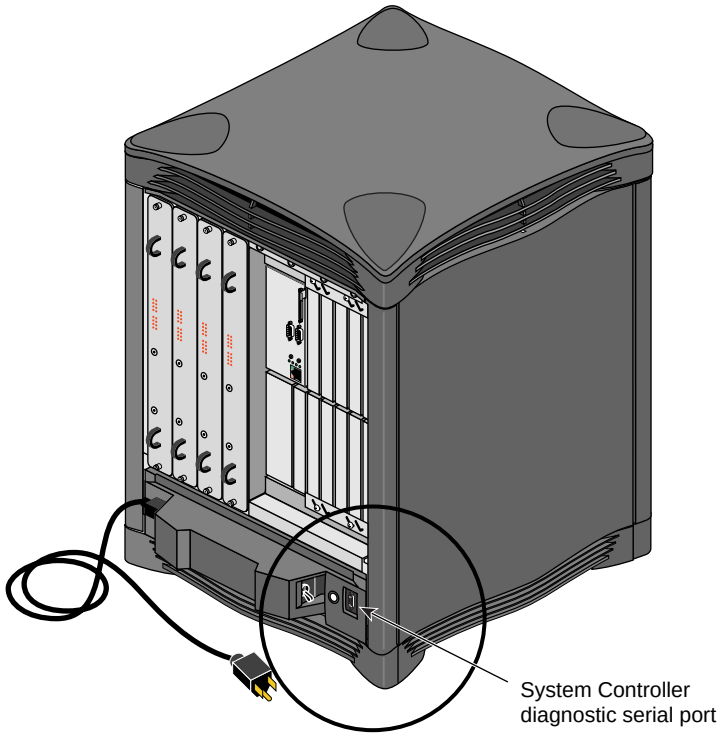

Figure 2-8 Origin2000 Deskside Server Serial Port (Rear)

Caution: The 9-pin sub-D serial port on the rear of the system is the same electrically as the 8-pin mini-DIN serial port on the front of the Origin2000 server. If anything is attached to either port before you begin installing the IRIS FailSafe system, it must be removed.

Figure 2-9 shows the **ALTERNATE CONSOLE** port on the Origin Rack MMSC.

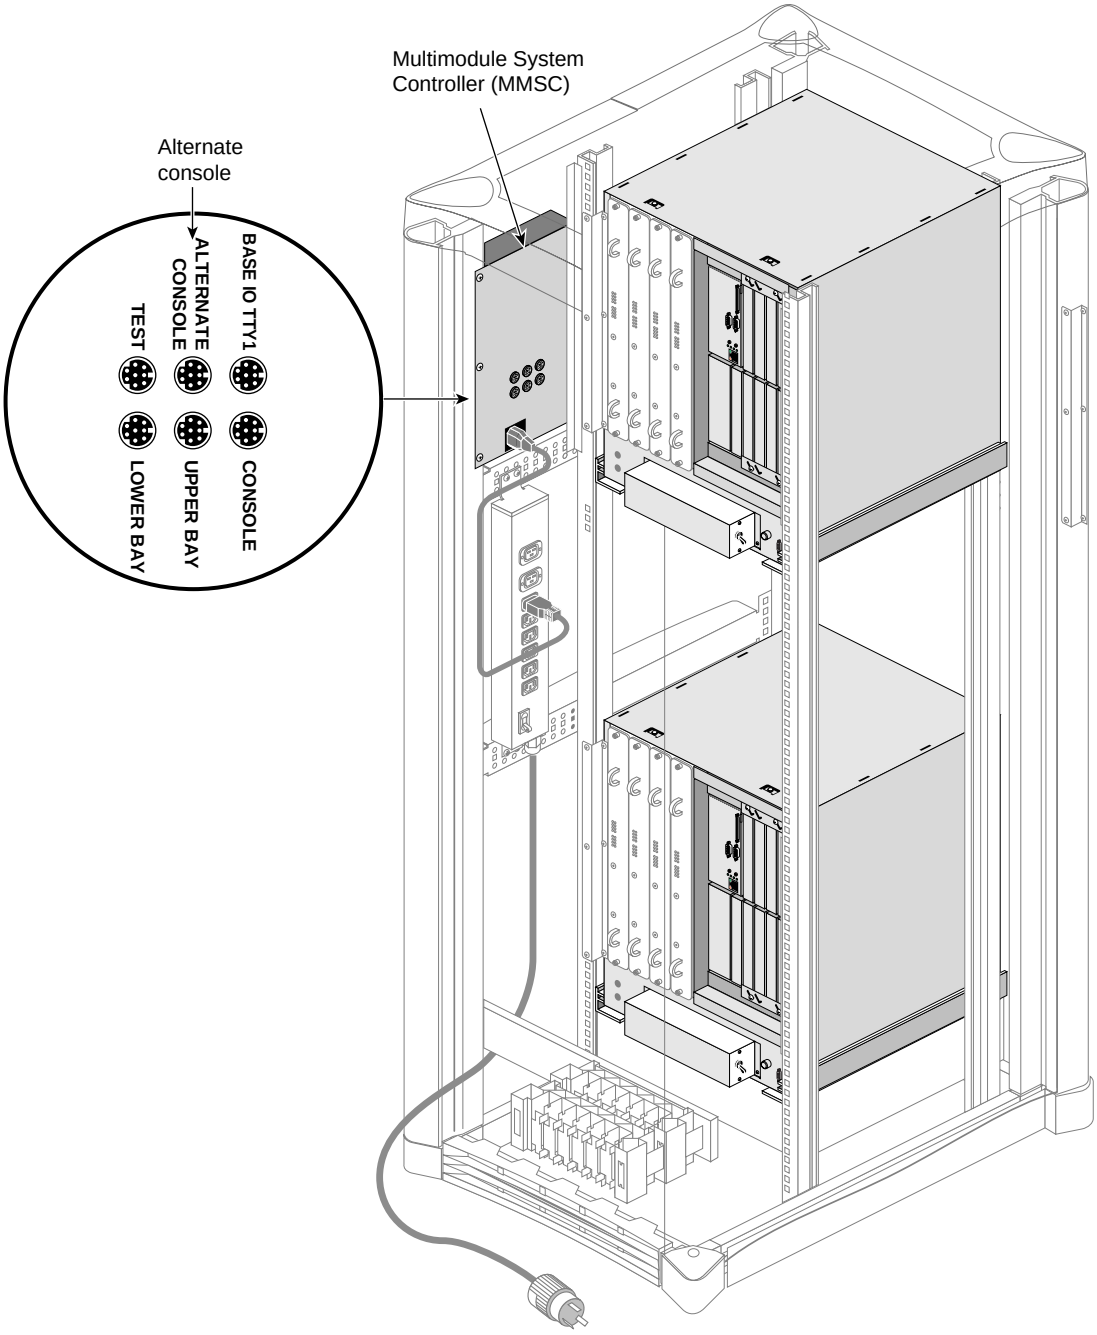

Figure 2-9 Origin Rack MMSC ALTERNATE CONSOLE Port

Figure 2-10 shows the tty ports on the Origin2000 server IO6 panel. The right-hand one is the console port, to which the system console is connected. The IRIS FailSafe serial connection uses the left-hand tty port.

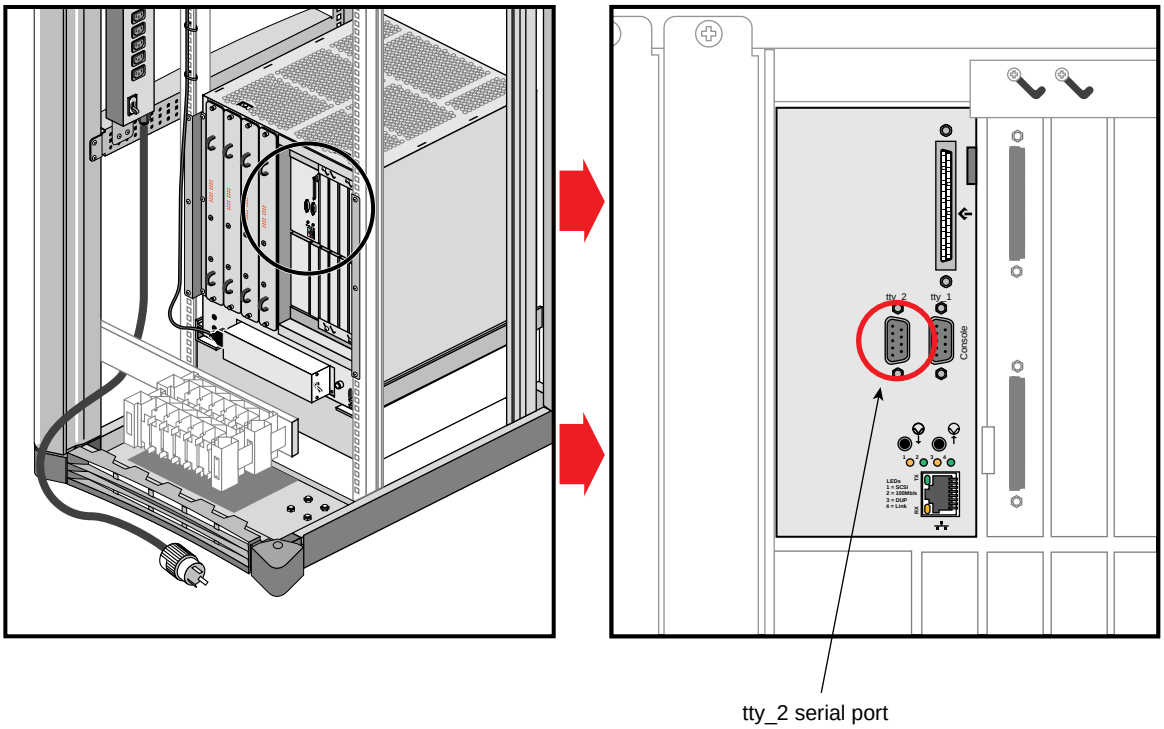

Figure 2-10 Origin2000 Rackmount Server tty Ports

Figure 2-11 shows the serial ports on the Origin200 server.

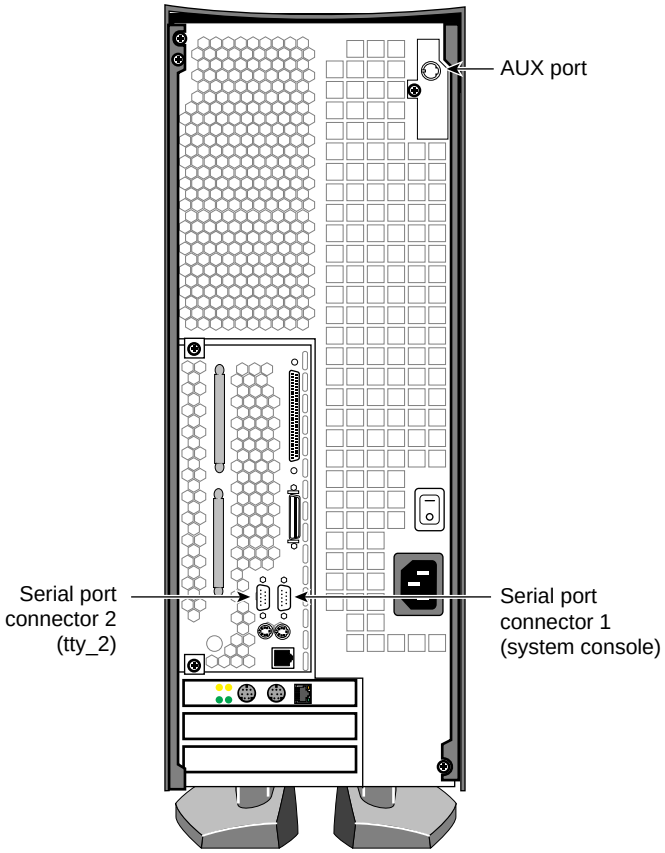

Figure 2-11 Origin200 Deskside Serial Ports

2.5.2 Cabling the Serial Connection

Follow these steps:

1. For an Origin2000 deskside server, make sure that nothing is connected to the serial port on the front or the back.
2. In the Origin2000 Rack, make sure that the cabling to the MMSC is correct. For IRIS FailSafe, Origin family rackmount servers in a rack must be set up as one system (machine); all Origin family servers in a multirack system must be set up as one system for IRIS FailSafe.

3. Cable the serial connection; Table 2-2 summarizes the connection for various combinations of servers.

Table 2-2 Serial Cables, Origin2000 and Origin200 Servers

Servers	Cable	Connections
Two Origin2000 deskside	018-0691-001 x 2	MSC serial port to tty_2 on the IO6 board
Origin2000 deskside and Origin2000 rackmount	018-0691-001 018-0668-001	Deskside: MSC serial port on rear to tty_2 on rack module IO6 board Rackmount: MMSC **ALTERNATE CONSOLE** to tty on deskside IO6
Two Origin2000 rackmount	018-0668-001 x 2	MMSC **ALTERNATE CONSOLE** port to tty_2 on IO6 board
Two Origin200	018-0668-001 x 2	**AUX** port to tty_2
Origin2000 deskside and Origin200	018-0691-001 018-0668-001	Origin2000: MSC to tty_2 on Origin200 IO6 board Origin200: **AUX** port to tty_2 on Origin2000 IO6 board
Origin2000 rackmount and Origin200	018-0668-001 018-0668-001	Origin2000: MMSC **ALTERNATE CONSOLE** to tty_2 on Origin200 IO6 Origin200: **AUX** port to tty_2 on Origin2000 IO6 board

Note: If the customer has ordered the 40-foot cable option, the cable part numbers are the same except for the last three digits, which are 101 instead of 001. See Table 1-9 in Chapter 1 for part numbers included in each 40-foot option kit.

Figure 2-12 shows the serial connection between two Origin2000 deskside servers.

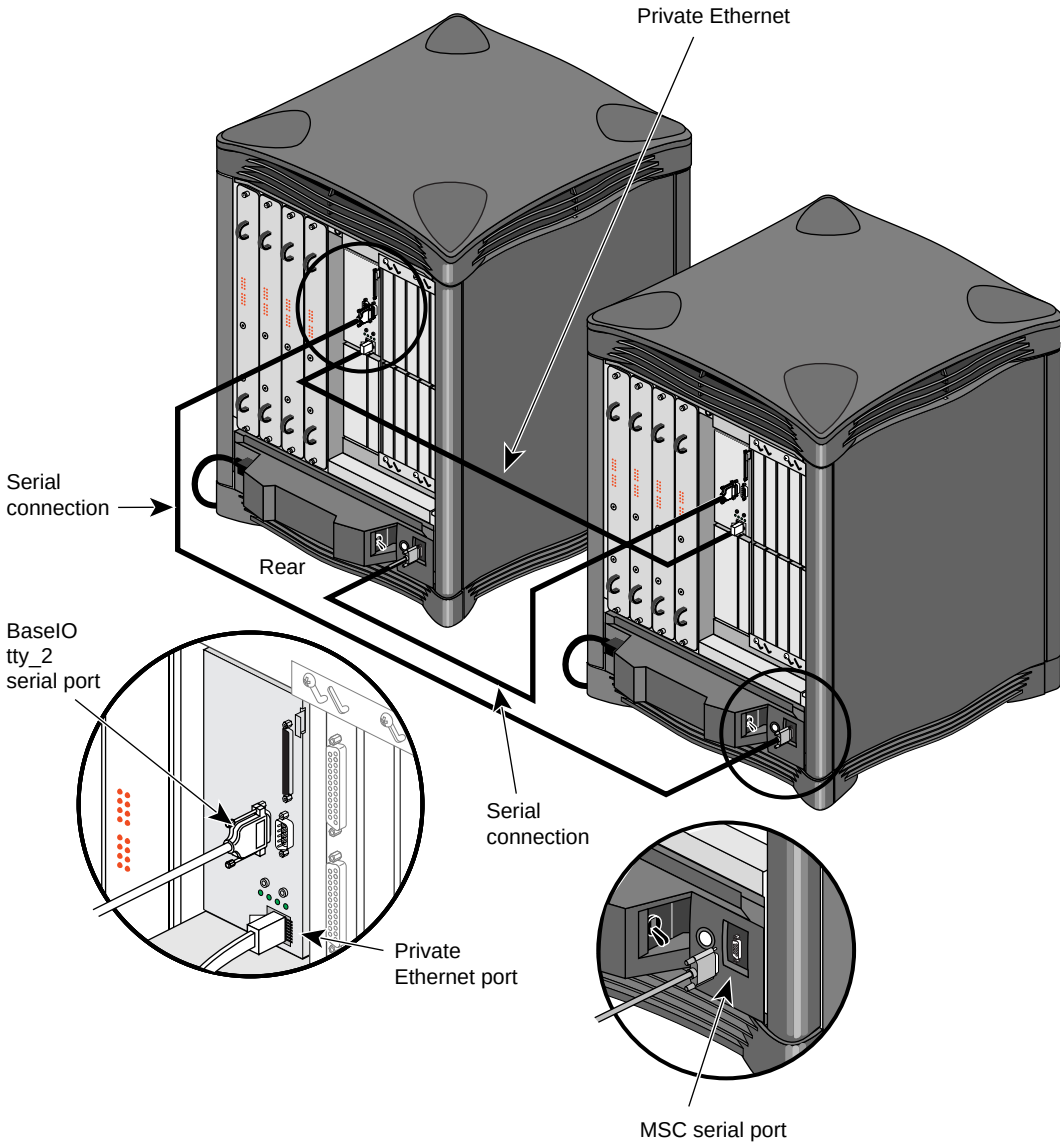

Figure 2-12 Serial Connection and Private Ethernet Connection, Origin2000 Deskside Servers

Figure 2-13 shows the serial connection between two Origin2000 rackmount servers.

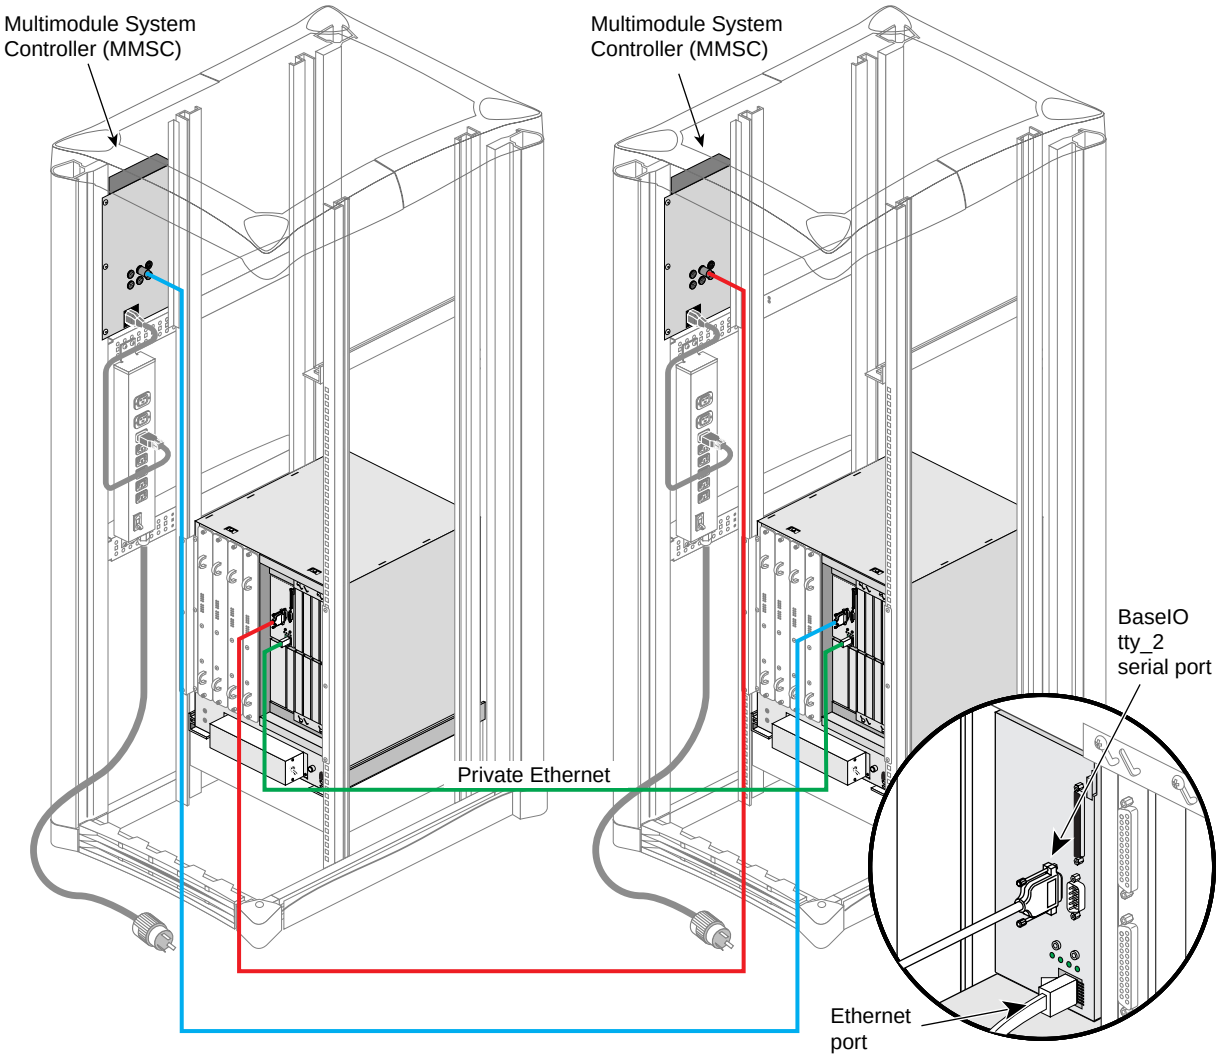

Figure 2-13 Serial Connection and Private Ethernet Connection, Origin2000 Rackmount Servers

Figure 2-14 shows the serial connection between two Origin200 deskside servers.

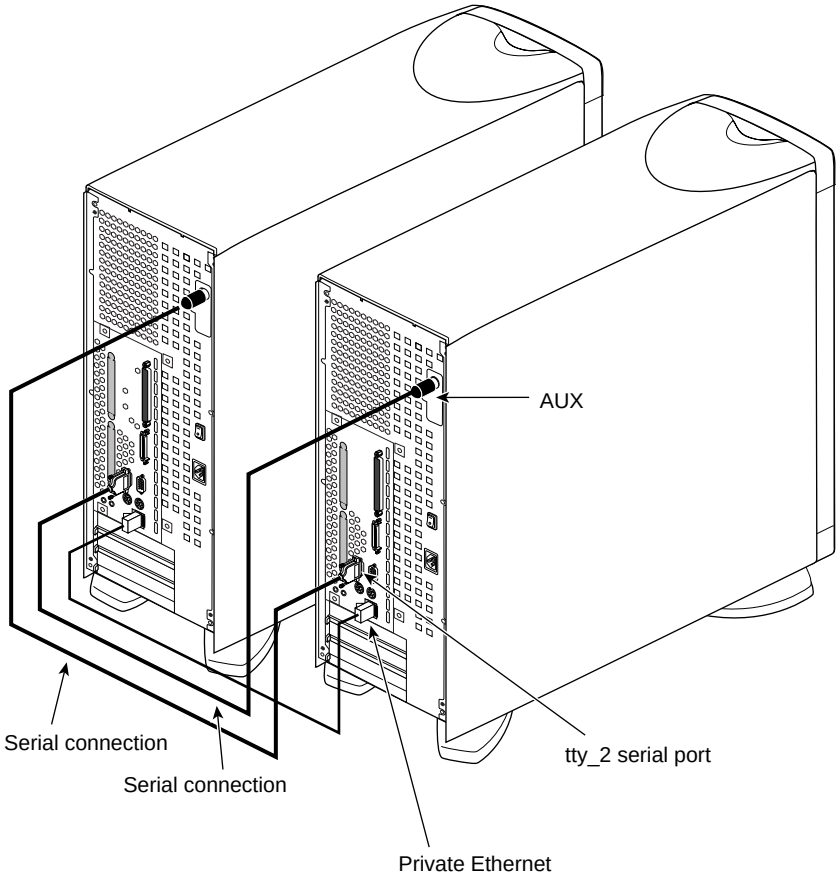

Figure 2-14 Serial Connection and Private Ethernet Connection, Origin200 Deskside Servers

Figure 2-15 shows the serial connection between an Origin2000 deskside server and an Origin200 deskside server.

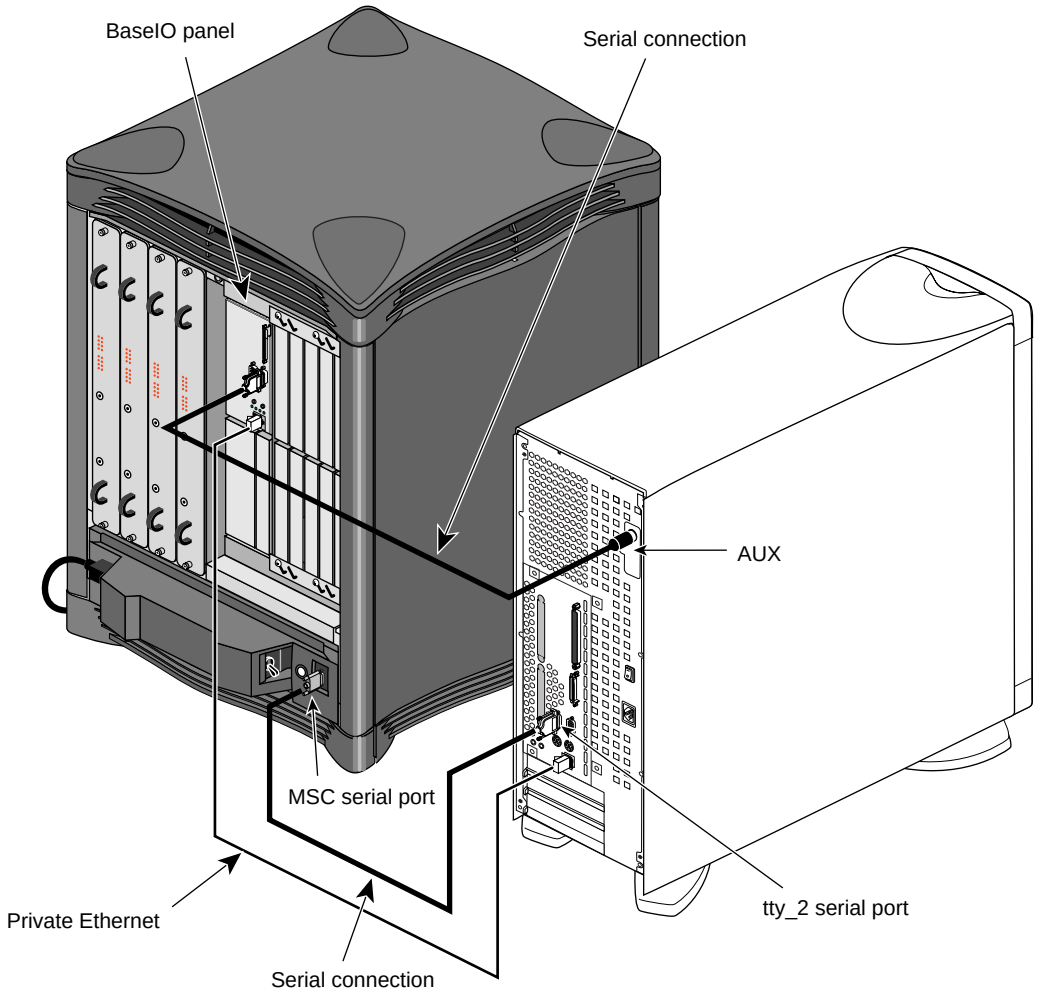

Figure 2-15 Serial Connection and Private Ethernet Connection, Origin2000 and Origin200 Deskside Servers

2.6 Testing the Serial Connection

To test the serial connection between the IRIS FailSafe servers, follow these steps:

1. Make sure the IRIS FailSafe servers are powered on.
2. Stop IRIS FailSafe on both nodes:

   ```
   /etc/init.d/failsafe stop
   ```

3. Test the connection. For example, if the serial cable was connected to tty_2 and the other host is a rackmount Origin2000 server, enter

   ```
   /usr/etc/ha_spng -i 10 -f /dev/ttyd2 -d MMSC
   ```

If the serial cable was connected to tty_2 and the other host is a deskside Origin2000 or Origin200 server, enter

```
/usr/etc/ha_spng -i 10 -f /dev/ttyd2 -d MSC
```

In this command, `/dev/ttyd2` is the tty of the node on which you are entering this command, and `MMSC` or `MSC` is the system controller of the other node.

No output appears; check the return value of the command. If the return value is 0, the connection is good.

If the return value is 1, perform these checks:

- Verify that the IRIS FailSafe server is powered on.
- Verify the cable connections from one server's serial port or remote power control unit and the other server's System Console port.

4. Repeat step 3 on the second node.

Note: If the system administrator has changed the MMSC or MSC password from the default, the IRIS FailSafe software must be notified; see Section 2.1, "Installing the Software," or the ha_spng(1M) man page.

If *ha_spng* fails, make sure all the cables are seated properly and rerun the command with a higher verbosity level using the **-v** option. This level shows all the commands and responses to and from the MMSC / MSC. To do more debugging, or if the messages from *ha_spng* are not clear, run the MMSC / MSC commands by hand directly after connecting to the MMSC / MSC through **cu -l <tty_dev>**.

Note: Precede all MMSC and MSC commands with Ctrl+T.

The commands *ha_spng* and *ha_killd* send the following commands to the MSC for pings and resets:

- The following command determines the version.

```
MSC> ver
        ok VER 3.0
```

- The following commands replace the password (if set) with none, and reset the MSC.

```
MSC> pas none
        ok

MSC> rst
        ok
```

For MMSC pings:

- The **ALTERNATE CONSOLE** port (COM5) is in RAT mode by default, so the MMSC prompt might not be visible. To get to the MMSC prompt, *ha_spng* and *ha_killd* send Ctrl+T to the tty port, followed by

```
R . MMSC

 MMSC> ^U
       CANCEL
MMSC> ver
       R1:MMSC 3.1
```

- The following commands reset the MMSC:

```
MMSC> ^U
      CANCEL
MMSC> authority service <password_if_present>
      R1:OK
MMSC> ^U
      CANCEL
MMSC> R all pwr c 5
      R1U:ok
      R1L:ok
```

2.7 Setting Automatic Power-On for Origin200 Servers

You must configure each Origin200 server so that it can power on automatically after it has been powered off. To do so, you set the *aut* level in the System Controller to 1:

Note: Precede all MMSC and MSC commands with Ctrl+T.

1. At the MSC prompt, check the *aut* level in the system controller:

```
MSC> aut
ok 0
```

2. Set the *aut* level to 1:

```
MSC> aut 1
ok
```

```
MSC> aut
ok 1
```

2.8 Cabling the Challenge or Origin Vaults

This section explains how to cable each SCSI bus on each Challenge vault or Origin Vault to each IRIS FailSafe host. Note the following:

- Use the Origin Vault option in configurations with Origin family servers only.
- Use the Challenge vault in mixed Origin and Challenge configurations or in configurations with Challenge servers only.
- The SCSI buses in the Challenge vault are named A and B. The recommended procedure is to connect even-numbered SCSI buses to A and odd-numbered buses to B.
- The Origin Vault contains only one SCSI bus.
- For connection to differential Fast-20 devices, the Origin2000 module must have one or more Ultra SCSI XIO boards, each with four SCSI ports. Figure 2-16 shows the location of example Ultra SCSI XIO boards.

  The ports are implemented as double connectors, as shown in Figure 2-16. Ports 1, 2, and 3 are DF ports; port 0 is an autosensing DF/SE port.

Caution: Do not use port 0 on the Ultra SCSI XIO board for dual-hosting purposes.

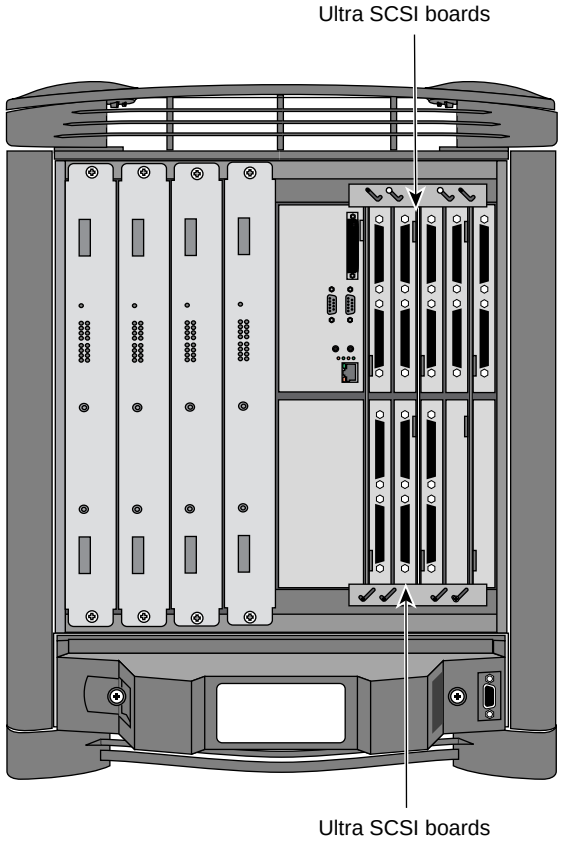

Figure 2-16 Ultra SCSI Ports

- The Origin200 (and the Origin2000) server can use a differential PCI SCSI board to connect to differential Fast-20 devices. PCI SCSI cards are available in single-ended or differential versions, denoted by icons near the connectors; see Figure 2-17.

Figure 2-17 PCI SCSI Option Board Icons

Note: IRIS FailSafe requires the differential version of the Origin Vault, which must be cabled to the differential version of the PCI SCSI board only.

- The Y cables shipped with the Ultra SCSI board do not connect directly to the Challenge vault or Origin Vault, nor do the SCSI cables shipped with the IRIS FailSafe option connect directly to the Ultra SCSI board; you must connect the storage option's SCSI cable to the Y cable to make the connection. For more information on the Ultra SCSI board, see the *Ultra SCSI XIO Board Owner's Guide* or *Ultra SCSI XIO Board Installation Instructions.*

In an IRIS FailSafe configuration, each SCSI bus runs from a controller on one server through a vault (or RAID storage option) to the second IRIS FailSafe host. Cabling the SCSI buses on the vaults to each IRIS FailSafe host is explained in

- Section 2.8.1, “Cabling the Challenge Vault”
- Section 2.8.2, “Cabling the Origin Vault”

2.8.1 Cabling the Challenge Vault

To cable the SCSI buses on the Challenge vaults to each IRIS FailSafe host, follow these steps:

1. Power off both servers and the vaults.
2. If necessary, for the Ultra SCSI XIO board on each host, insert an Ultra SCSI Y cable into the connector for ports 2 and 3 on each host. Because the recommended procedure is to connect even-numbered SCSI buses to A, and you cannot use port 0 on the Ultra SCSI board for dual-hosting purposes, you must use port 2.

   The connectors or cables should be labeled, for example, **MOD 1 IO6 SCSI2** and **MOD 1 IO6 SCSI3**.
3. On the back of the Challenge vault, insert one of the cables (p / n 9290111) into the top SCSI channel A socket, as shown in Figure 2-18.
4. Connect the other end of this cable to the Ultra SCSI cable coming from port 2 (such as **MOD 1 IO6 SCSI2**) on the first IRIS FailSafe host, as shown in Figure 2-18. Label the connector or cable end.

   Note: Because the recommended procedure is to connect even-numbered SCSI buses to A, and you cannot use port 0 on the Ultra SCSI board for dual-hosting purposes, you must use port 2.
5. Plug a cable into the lower SCSI channel A port on the back of the vault. Label the cable end.
6. Connect the other end of this cable to the cable coming from Ultra SCSI port 2 on the second IRIS FailSafe host (such as **MOD 2 IO6 SCSI2**).
7. Repeat steps 3 through 6 for the second vault, using Ultra SCSI port 3 (for example, **MOD 1 IO6 SCSI3** and **MOD 2 IO6 SCSI3**). Figure 2-18 diagrams Challenge vault cabling.

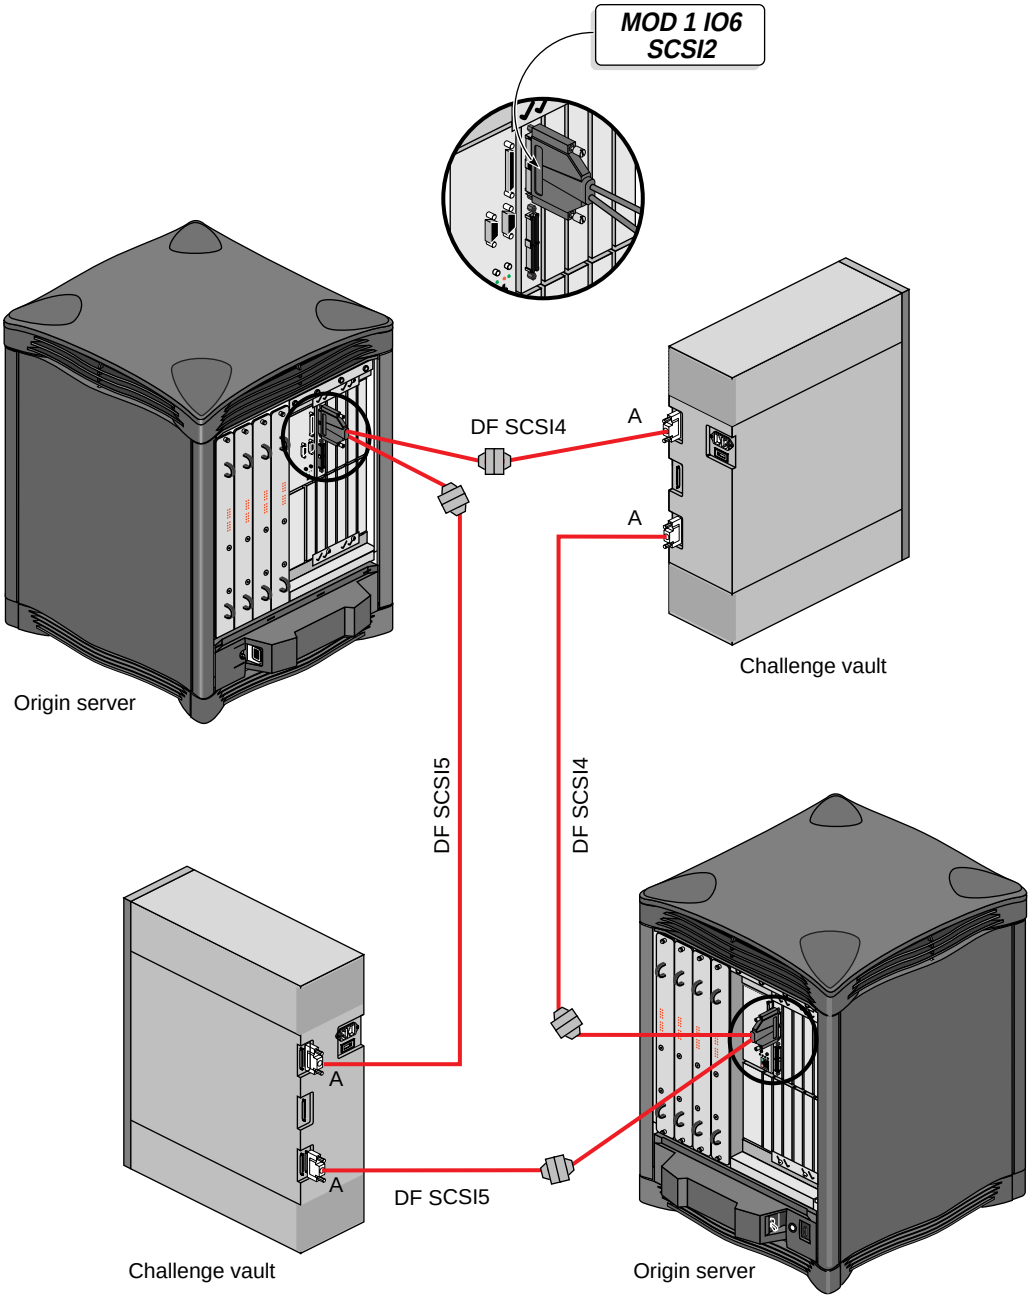

Figure 2-18 Cabling Challenge Vaults and Origin2000 Servers

8. If necessary, reset the SCSI bus selector settings on the front of each Challenge vault if they have been unset accidentally.
9. Power on the vaults and servers.
10. After the disks are set up, do a dummy I/O to test each link. For example, from the first host, read data on disk 8 on SCSI bus 5 (*dks5d8s0*) to see if the disk light for disk 8 on the vault lights up. Repeat this test from the second IRIS FailSafe host and check if the same LED is activated.

2.8.2 Cabling the Origin Vault

The Origin Vault expansion option is available as a rackmountable or deskside unit. Either format is available with a single-ended controller for the Origin Vault six 3.5-inch SCSI disks, or a differential controller board. Only the DF version can be used for IRIS FailSafe. Besides the 3.5-inch SCSI disks, the option can contain 5.25-inch media (always single-ended), but these are not relevant for IRIS FailSafe operation.

Figure 2-19 shows SCSI ports on the rackmount version of the DF Origin Vault option. (The SE version has only two SCSI connectors, one for the 5.25-inch drives and one for the 3.5-inch drives.)

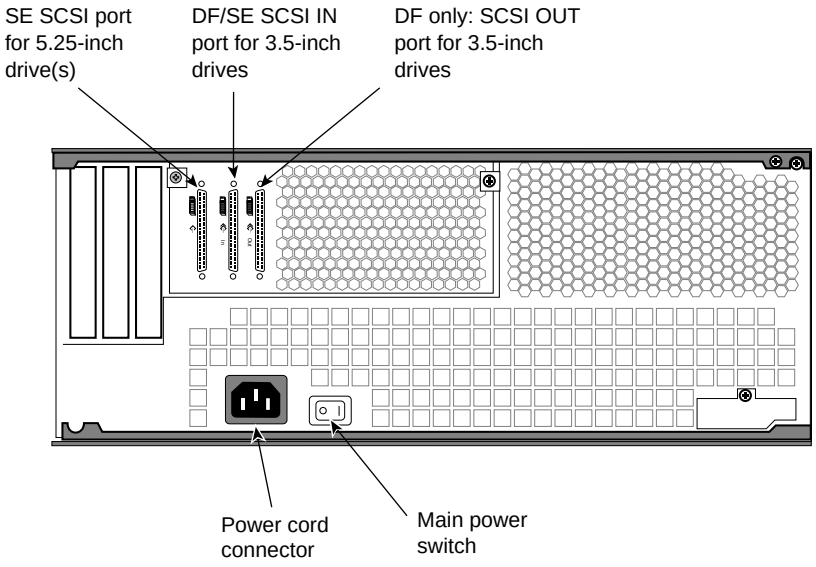

Figure 2-19 Origin Vault Rackmountable Enclosure Rear View (Differential)

Although the ports for the 3.5-inch drives are labeled for in and out connections, they can be used interchangeably.

The Origin Vault has both a main (master) power switch at the rear and a power button in the front, as shown in Figure 2-20.

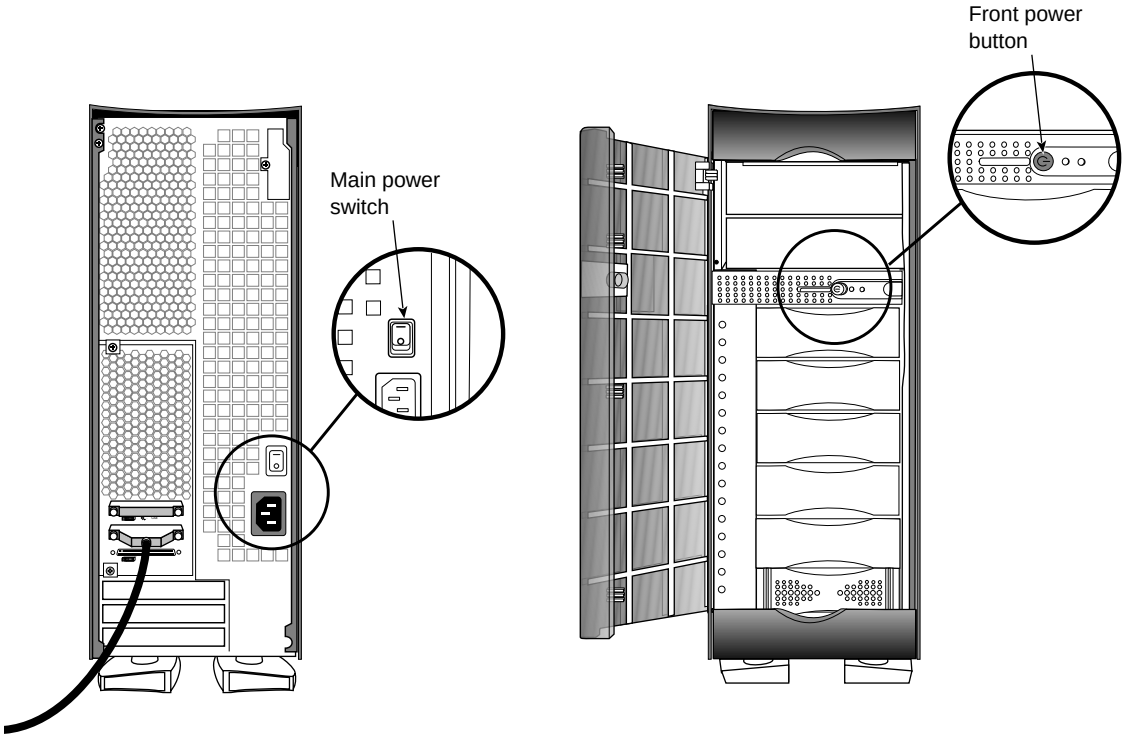

Figure 2-20 Main Power Switch (Circuit Breaker) and Power Button

Pressing the power button at the front causes a soft power-off. Residual power is still supplied to the unit as long as the main power switch at the rear is on. The main power switch is a circuit breaker that cuts off all power to the Origin Vault system. If you turn off the main switch and do not press the power button at the front, the Origin Vault system immediately powers back on when you turn the main switch back on. The system controller remembers the state of the front power switch independently of the main power switch so that when power is restored after a power failure, the system immediately powers back on.

The Origin Vault DF version includes one 1-meter and one 10-meter SCSI cable. For dual-hosting, the customer must order an additional 10-meter SCSI cable. The 1-meter cable can be used to connect the 5.25-inch drive controller to a host IO6 SCSI connector or to daisy-chain two Origin Vault options; to set up both connections requires ordering an additional 1-meter SCSI cable.

To cable the SCSI buses on the Origin Vault(s) to each IRIS FailSafe host, follow these steps:

1. Power off both servers and the vaults.
2. If necessary, for the Ultra SCSI XIO board on each host, insert an Ultra SCSI Y cable into the connector for ports 2 and 3 on each host. Remember that you cannot use port 0 on the Ultra SCSI board for dual-hosting purposes.

   The connector or cable should be labeled, for example, **MOD 1 IO6 SCSI2**.

3. Connect a 10-meter SCSI cable (supplied with the Origin Vault) to a DF SCSI connector in the first host, such as a PCI DF SCSI connector or the Y cable end for the Ultra SCSI port you are using, such as port 2.
4. Insert the other end of the 10-meter SCSI cable into the IN (middle) SCSI connector on the back of the Origin Vault expansion option, as shown in Figure 2-21.

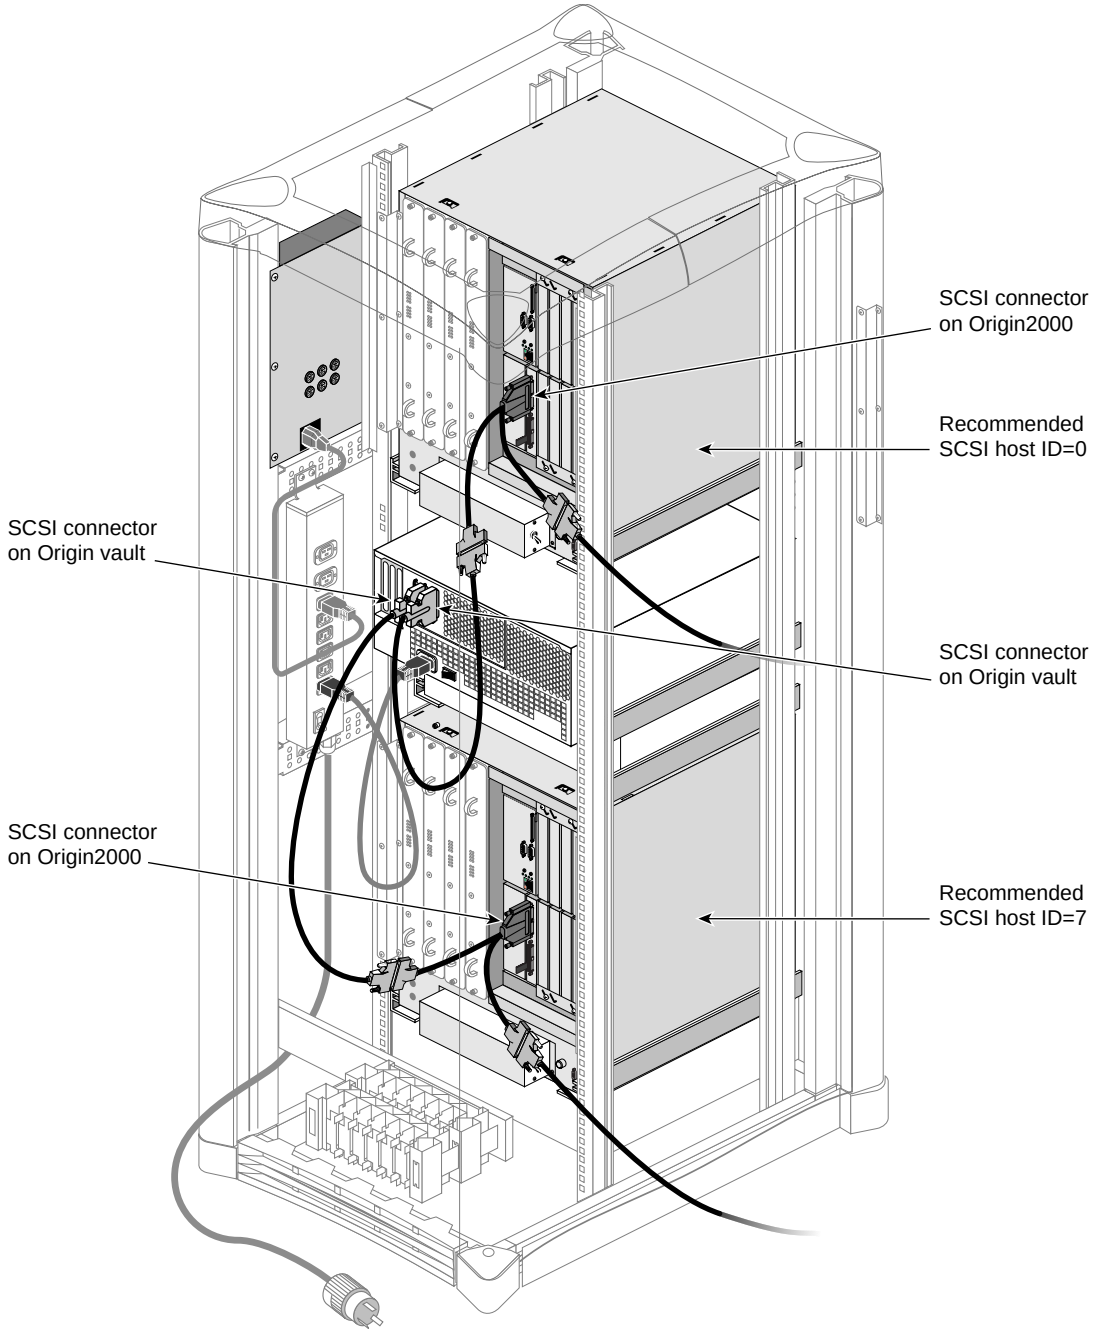

Figure 2-21 Connecting the Rackmount SCSI Cable to an Origin2000 Host

5. Remove the terminator from the Origin Vault's OUT SCSI connector.

6. If you are daisy-chaining two Origin Vault options, use the 1-meter cable to connector the first Origin Vault OUT SCSI connector to the second Origin Vault IN SCSI connector. Figure 2-22 shows two daisy-chained Origin Vault towers.

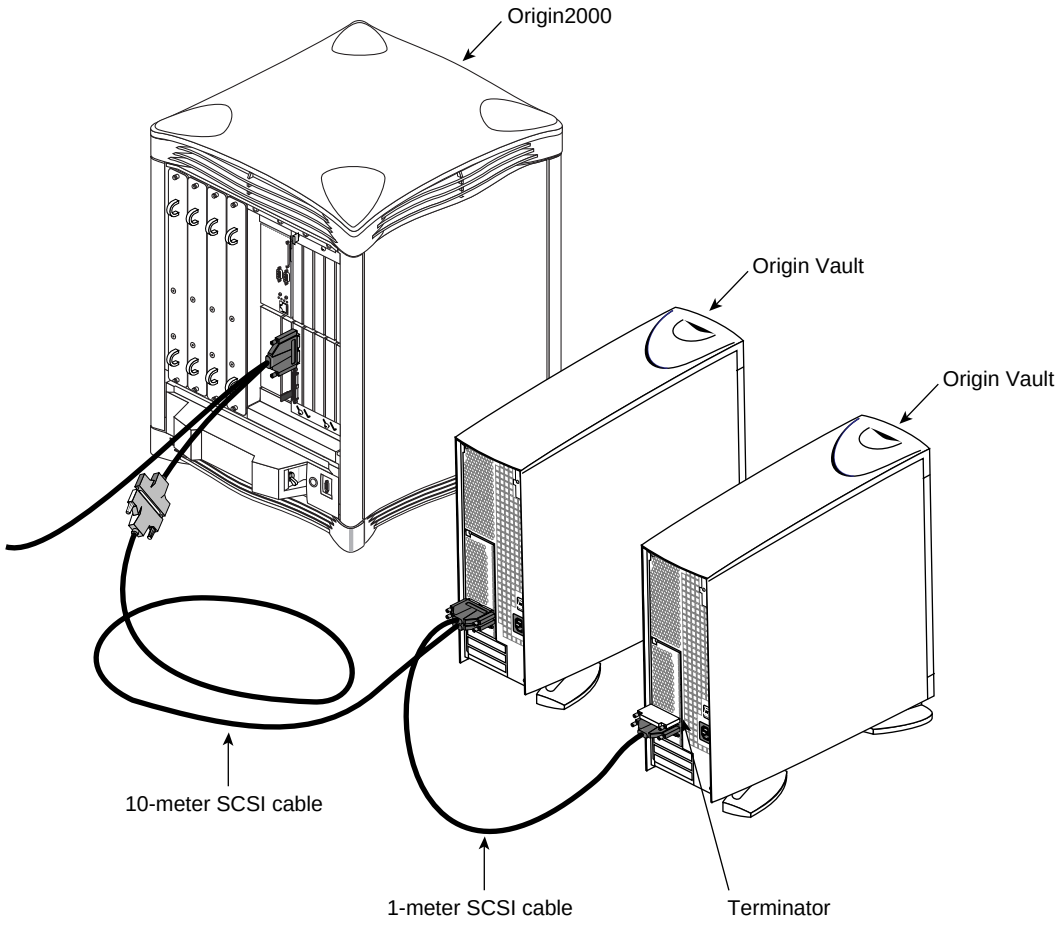

Figure 2-22 Daisy-Chained Origin Vault Towers

7. To cable the second host to the Origin Vault(s), plug one end of a 10-meter SCSI cable into the Origin Vault option's 3.5-inch OUT controller port. Attach the other end of this cable to a differential SCSI connector in the second host, such as a PCI DF SCSI connector, or the Y cable to channel 1, 2, or 3 on an Ultra SCSI XIO board.

   Caution: Do not use port 0 on the Ultra SCSI board for dual-hosting purposes.

Figure 2-23 shows an example with a single Origin Vault option; Figure 2-24 shows an example with daisy-chained Origin Vault options.

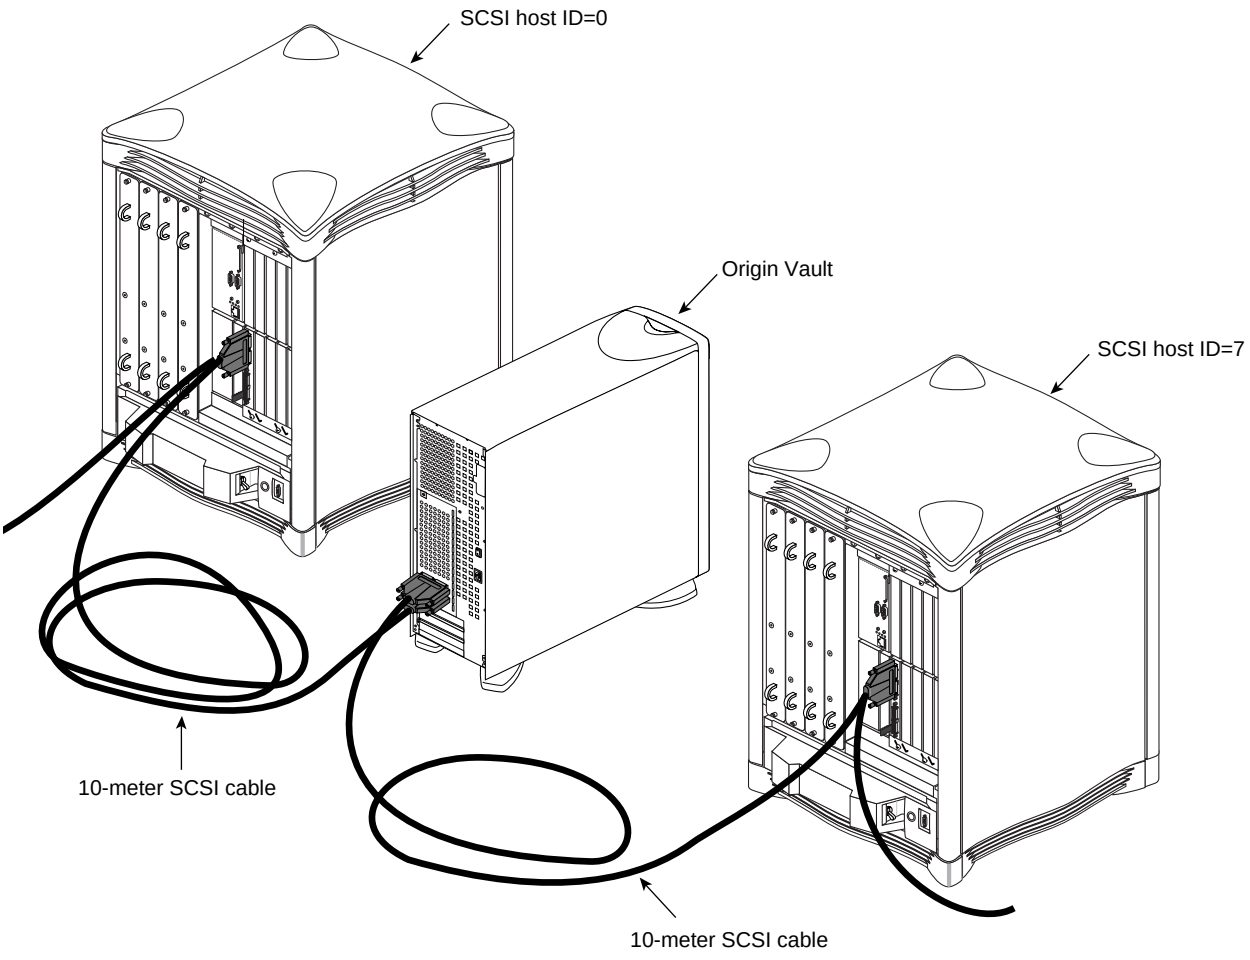

Figure 2-23 Dual-Hosted Configuration Sharing One Origin Vault Option

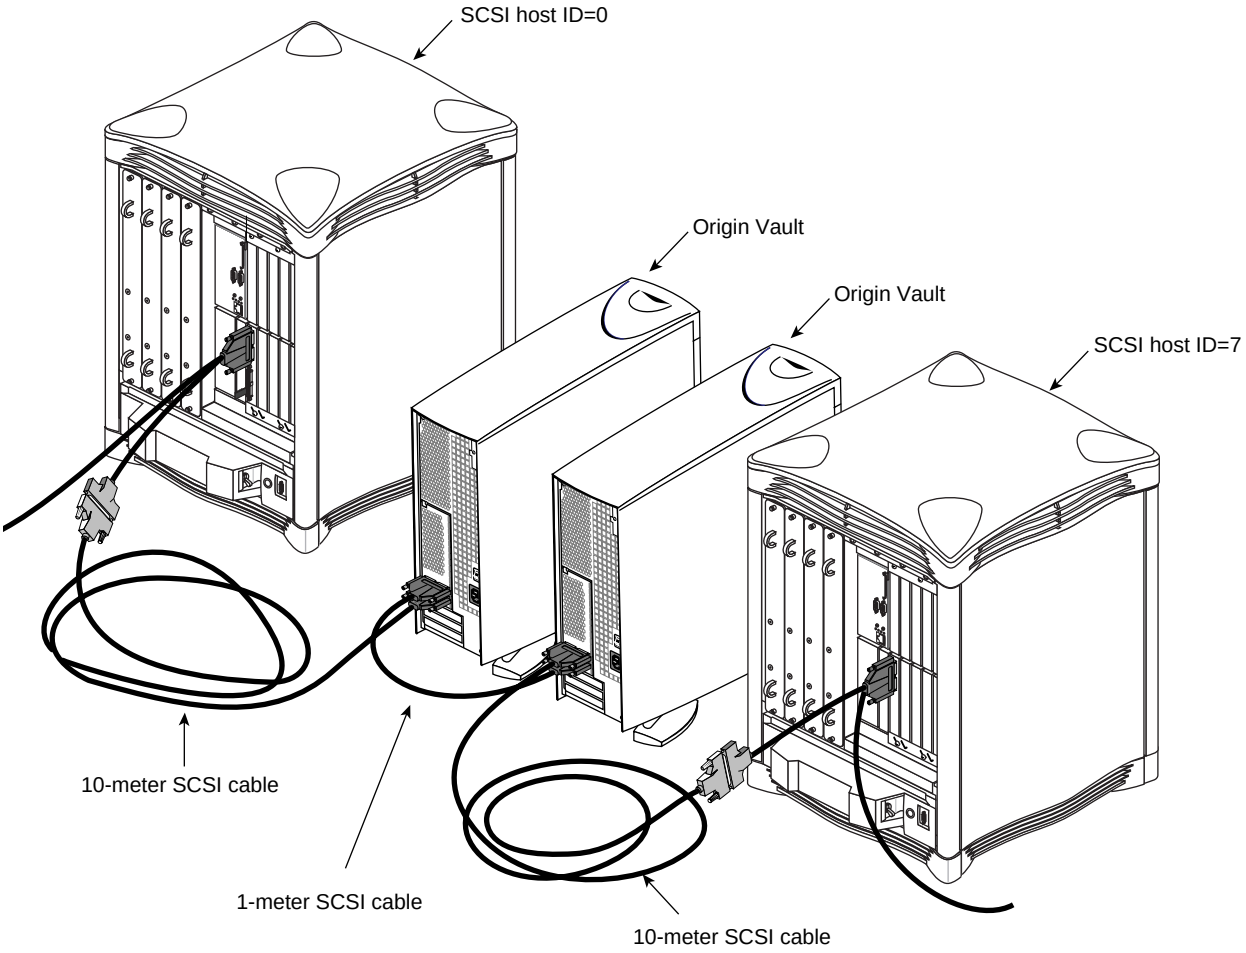

Figure 2-24 Dual-Hosted Configuration Sharing Two Daisy-Chained Origin Vault Options

8. Repeat steps 2 through 7 for the second Origin Vault option, using Ultra SCSI port 3 (for example, **MOD 1 IO6 SCSI3** and **MOD 2 IO6 SCSI3**). Figure 2-25 diagrams IRIS FailSafe Origin Vault cabling.

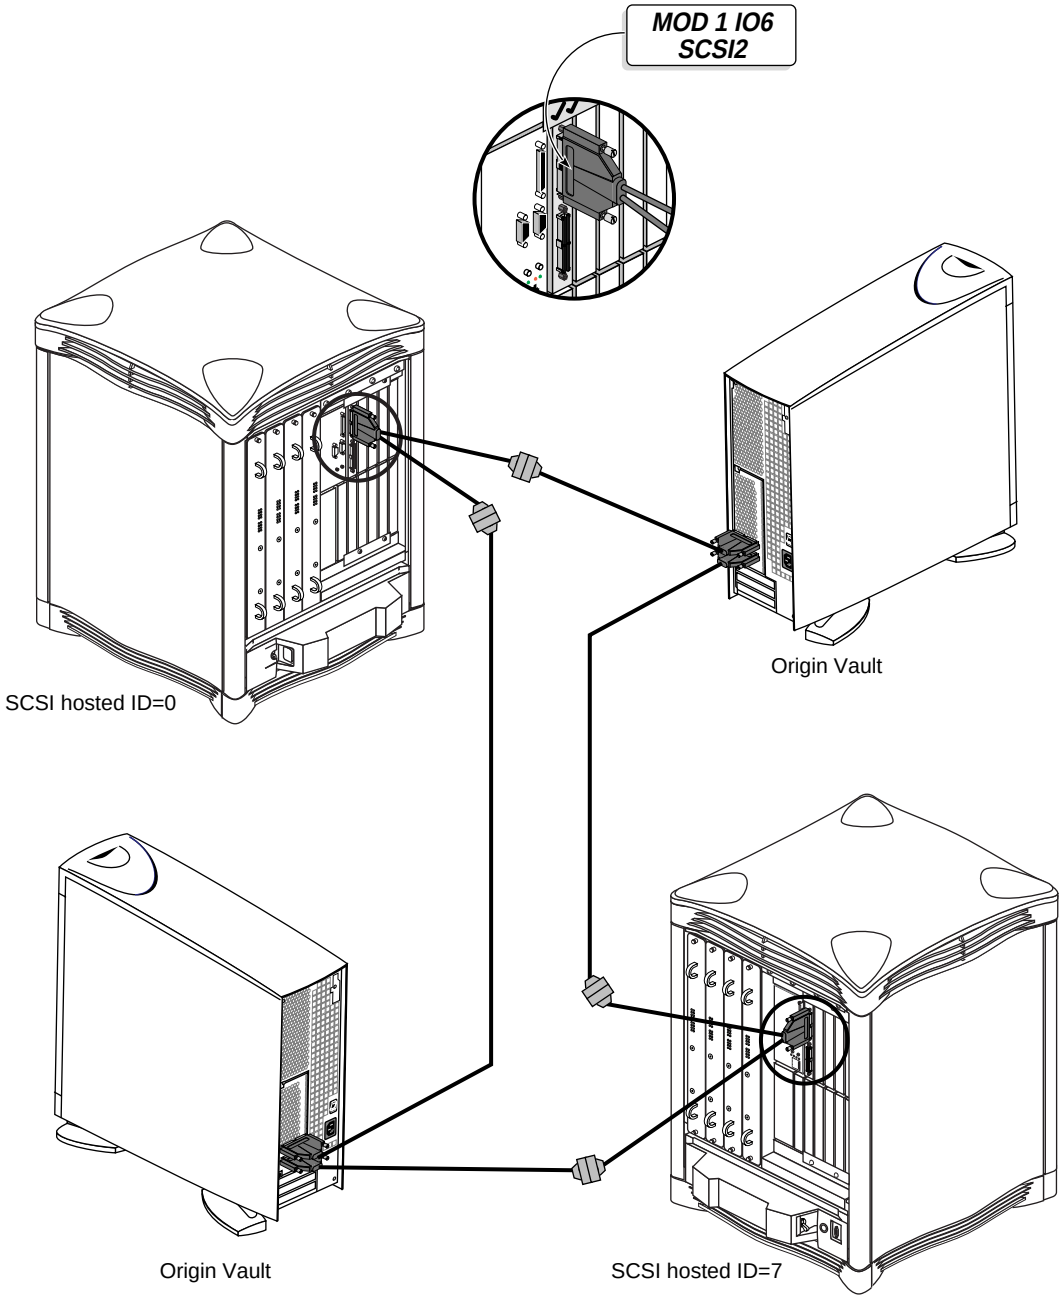

Figure 2-25 Cabling Origin Vaults and Servers

9. Power on the Origin Vault options and the servers.

10. After the disks are set up, do a dummy I/O to test each link. For example, from the first host, read data on disk 8 on SCSI bus 5 (*dks5d8s0*) to see if the disk light for disk 8 on the vault lights up. Repeat this test from the second IRIS FailSafe host and check if the same LED is activated.

2.9 Cabling the Challenge RAID Storage System to the Origin Servers

This section explains how to cable each SCSI bus on a Challenge RAID storage system to each IRIS FailSafe host. Note the following:

- Use the Challenge RAID storage system in these configurations:
  - Origin2000 and Challenge L
  - Origin200 and Challenge L
  - Origin200 and Challenge S
  - Challenge family servers only
- The SCSI buses in the Challenge RAID are named A and B. The recommended procedure is to connect even-numbered SCSI buses to A, and odd-numbered buses to B.
- The Origin2000 server IO6 board's single-ended SCSI-3 port cannot be used for IRIS FailSafe. For connection to differential Fast-20 devices, the Origin2000 module must have one or more Ultra SCSI XIO boards, each with four SCSI ports. The ports are implemented as double connectors, as shown in Figure 2-16. Ports 1, 2, and 3 are DF ports; port 0 is an autosensing DF/SE port.

  Caution: Do not use port 0 on the Ultra SCSI board for dual-hosting purposes.
- The Y cables shipped with the Ultra SCSI board do not connect directly to the Challenge RAID, nor do the SCSI cables shipped with the IRIS FailSafe option connect directly to the Ultra SCSI board; you must connect the SCSI cable to the Y cable to make the connection. For more information on the Ultra SCSI board, see the *Ultra SCSI XIO Board Owner's Guide* or *Ultra SCSI XIO Board Installation Instructions.* Figure 2-16 shows the location of example Ultra SCSI XIO boards.

Table 2-3 charts SCSI ID switch settings for the Challenge RAID storage system.

Note: Challenge RAID SCSI ID switch settings do not conform to frequently used numbering schemes.

Table 2-3 Challenge RAID SCSI ID Switch Settings

SCSI ID Number	Switch Number				Use With Origin2000	Use With Origin200
	ID 0	ID 1	ID 2	ID 3		
1	Off	On	On	On	Yes	Yes
2: Do not use	On	Off	On	On	No	No
3	Off	Off	On	On	Yes	Yes
4	On	On	Off	On	Yes	Yes

Table 2-3 (continued) Challenge RAID SCSI ID Switch Settings

SCSI ID Number	Switch Number				Use With Origin2000	Use With Origin200
	ID 0	ID 1	ID 2	ID 3		
5	Off	On	Off	On	Yes	Yes
6	On	Off	Off	On	Yes	Yes
7	Off	Off	Off	On	Yes	Yes
8	On	On	On	Off	Yes	No
9	Off	On	On	Off	Yes	No
10	On	Off	On	Off	Yes	No
11	Off	Off	On	Off	Yes	No
12	On	On	Off	Off	Yes	No
13	Off	On	Off	Off	Yes	No
14	On	Off	Off	Off	Yes	No
15	Off	Off	Off	Off	Yes	No

Note: The Origin200 server does not use all switch settings.

To cable the Challenge RAID storage system to the Origin servers, follow these steps:

1. Have ready the two SCSI cables shipped with Challenge RAID and the two 20-foot SCSI cables included in the Challenge RAID shipment (p / n 9290111). Figure 2-26 shows the Challenge RAID SCSI ports.

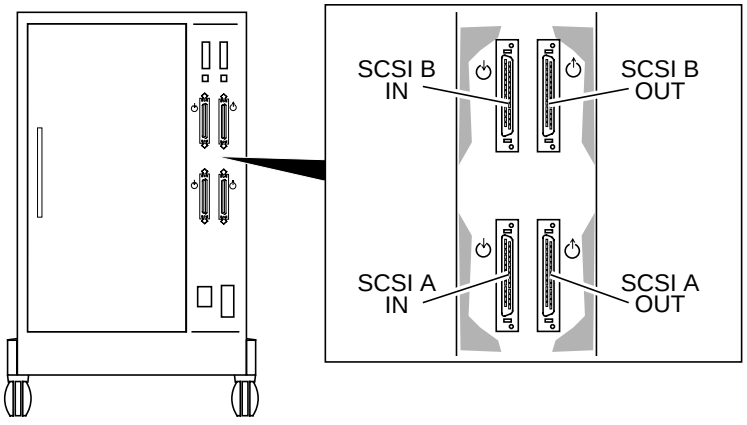

Figure 2-26 SCSI-2 Bus Connectors on Back of Challenge RAID Chassis

2. Power off both servers and the Challenge RAID storage system.
3. If necessary, insert an Ultra SCSI Y cable into the connector for ports 2 and 3 on both hosts. Since the recommended procedure is to connect even-numbered SCSI buses to A, and you cannot use port 0 on the Ultra SCSI board for dual-hosting purposes, you must use port 2.

The connectors or cables should be labeled, for example, **MOD 1 IO6 SCSI2** and **MOD 1 IO6 SCSI3**.

4. On the back of the Challenge RAID storage system, attach a SCSI cable to the Challenge RAID **SCSI A** in port, as shown in Figure 2-27, making sure the connector is inserted securely.

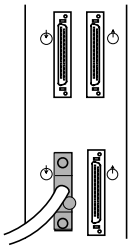

Figure 2-27 Connecting a SCSI Bus Cable to a Challenge RAID SCSI Port

5. Connect the other end of this cable to the cable attached to Ultra SCSI port 2 on the first host (for example, **MOD 1IO6 SCSI2**), as shown in Figure 2-28.

   Caution: Do not use port 0 on the Ultra SCSI board for dual-hosting purposes. Since the recommended procedure is to connect even-numbered SCSI buses to A, and you cannot use port 0 on the Ultra SCSI board for dual-hosting purposes, you must use port 2 on the Ultra SCSI board.

6. On the back of the Challenge RAID storage system, attach another SCSI cable to the Challenge RAID **SCSI A** out port.

7. Connect the other end of this cable to the cable attached to Ultra SCSI port 2 on the second host (for example, **MOD 2 IO6 SCSI2**).

8. Repeat steps 4 through 7 for the SCSI B port on the Challenge RAID storage system, using Ultra SCSI port 3 on each host (for example, **MOD 1 IO6 SCSI3** and **MOD 2 IO6 SCSI3**).

Figure 2-28 shows connections for Challenge RAID and Origin2000 servers in an IRIS FailSafe configuration.

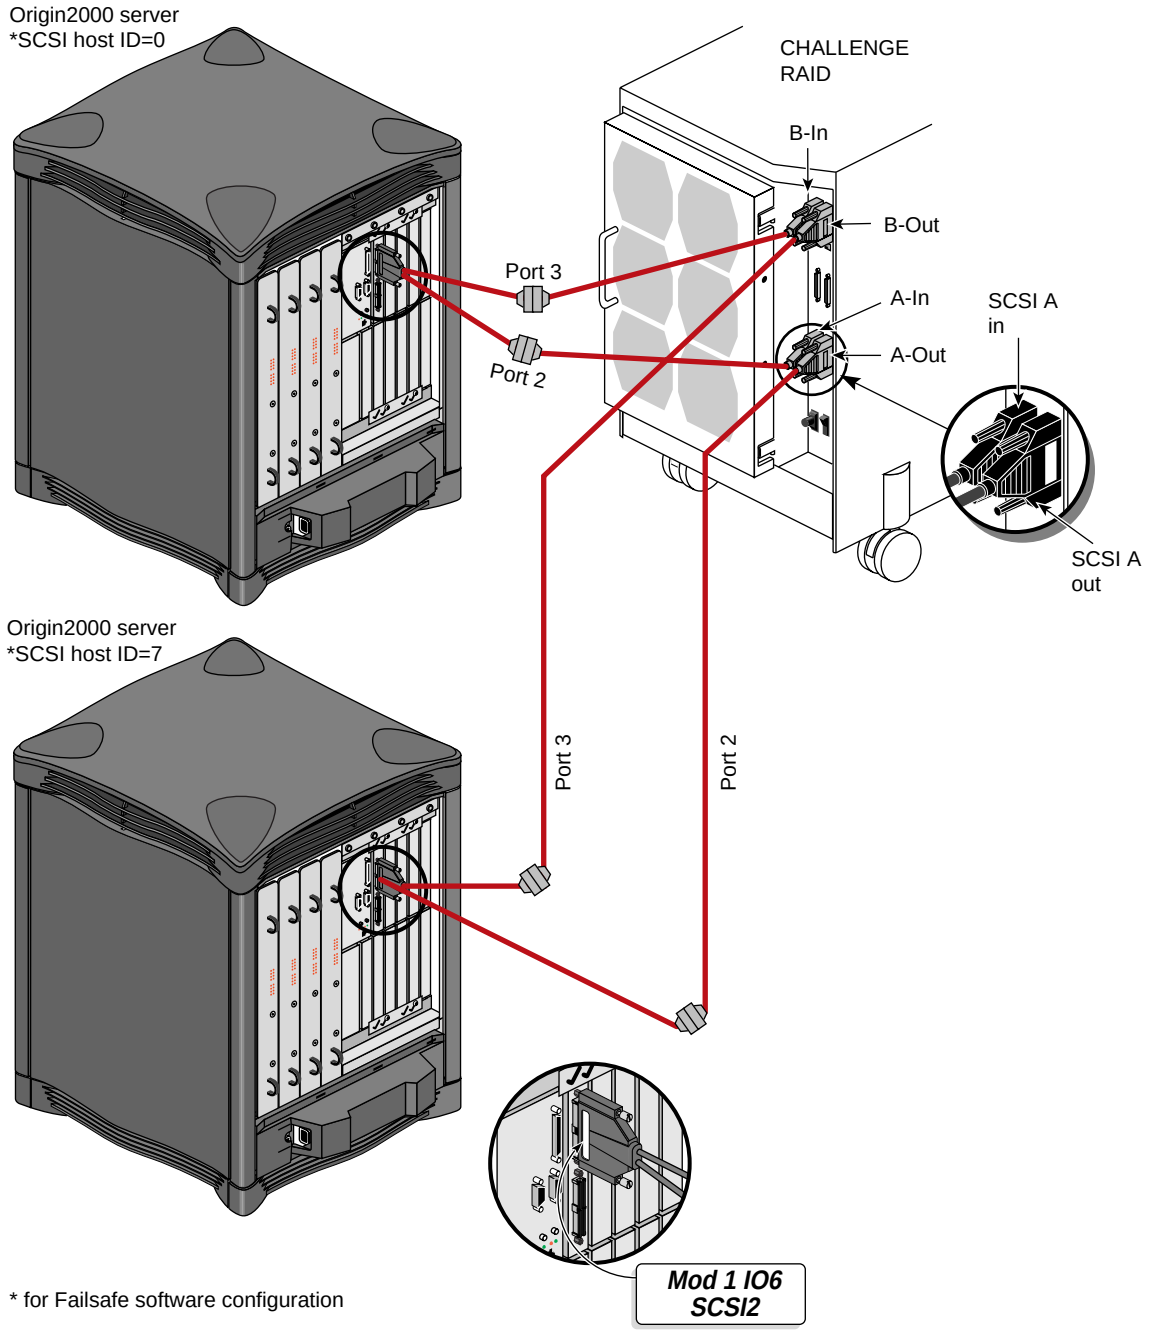

Figure 2-28 Example Cabling for Challenge RAID and Origin2000 Servers

2.10 Cabling the Fibre Channel Storage Options

Two Silicon Graphics Fibre Channel option boards for the Origin family—an XIO board with two FC ports and a PCI board with one FC port—connect to various Silicon Graphics Fibre Channel storage options. Cabling various fibre channel storage options for IRIS FailSafe is explained in these sections:

- Section 2.10.1, “Fibre Channel Enclosures and Disks”
- Section 2.10.2, “Fibre Channel Connectors”
- Section 2.10.3, “Cabling Fibre Rack Power Supplies for IRIS FailSafe”
- Section 2.10.4, “Cabling Fibre Channel RAID (DPEs) for IRIS FailSafe”
- Section 2.10.5, “Cabling FibreVault (JBOD) Enclosures for IRIS FailSafe”
- Section 2.10.6, “Using Optical FC Cables”

Note: For instructions on installing Fibre Channel storage options and boards, see the *Origin FibreVault and Fibre Channel RAID Installation Instructions*.

2.10.1 Fibre Channel Enclosures and Disks

Table 2-4 summarizes Silicon Graphics Fibre Channel storage options.

Table 2-4 Fibre Channel Storage Options

Marketing Code	Part Number	Model	Disk Type
P-F-BOX-R	9470123	Rackmount FibreVault disk array enclosure, or DAE; also known as “just a bunch of disks” (JBOD)	Non-RAID (9.1 GB)
P-F-VAULT-DS	9470126	Deskside (tower) FibreVault DAE (JBOD)	Non-RAID
P-F-RAID-B5X9-R	9470134 (chassis)	Rackmount Fibre Channel RAID enclosure (disk processor enclosure, or DPE) with ten FC RAID disks and one or two storage processors (SPs); two SPs are required for IRIS FailSafe	FC RAID (8.8 GB)
P-F-RAID-E5X9-R	9470136 (chassis)	Rackmount Fibre Channel RAID Vault expansion option; no SPs of its own, but connects to a DPE that contains at least nine FC RAID disks	RAID
P-F-RACK	9470157	Fibre Channel Rack (holds 11 DAEs or 5 DPEs); two power distribution units are required for IRIS FailSafe	Either

A Fiber Channel RAID enclosure (DPE) must contain only RAID disks; it has one or two storage processors (SPs) that control the RAID functionality. IRIS FailSafe requires two SPs (usually denoted SP A and SP B).

A FibreVault (DAE) has no SPs and can be either RAID or non-RAID, depending on the disks it holds:

- If it contains RAID disks, it is a Fibre Channel RAID expansion enclosure and must be cabled to a Fibre Channel RAID enclosure.

- If it contains non-RAID disks, it is a JBOD enclosure. For IRIS FailSafe, a JBOD enclosure can have a maximum of 110 disks per loop. It must have two link controller cards (LCCs).

The outside front and rear of both types of DAE—Fibre Channel RAID expansion enclosure and JBOD enclosure—look identical; the only difference is the drives inside them. Each drive module in a DAE (or any other Fibre Channel enclosure) has an identifying sticker with its part number:

- p/n 9470138: 8.8-GB RAID drives (520-byte sectors)
- p/n 9470140: 9.1-GB non-RAID drives (512-byte sectors)

Either type of FC storage must have two power supplies for IRIS FailSafe; the rack must have two power distribution units (PDUs).

Note the following:

- The FC RAID disks are not interchangeable with SCSI RAID disks.
- The FC non-RAID disks are not interchangeable with regular SCSI disks.
- The FC RAID and FC non-RAID disks are not interchangeable with each other and cannot be converted to the other type in the field.
- All disks in an FC enclosure must be either RAID or non-RAID. Disks in a DPE must be RAID.

2.10.2 Fibre Channel Connectors

Figure 2-29 shows the fibre channel connectors on the XIO and PCI FC boards.

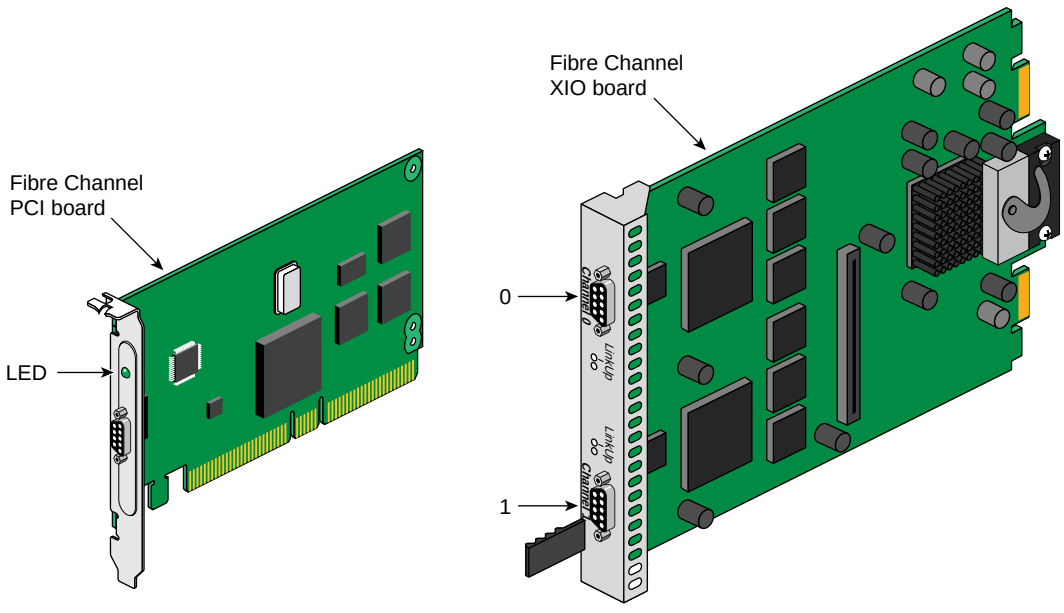

Figure 2-29 Fibre Channel PCI Board Connector

Figure 2-30 shows connectors on the Fibre Channel DAE (Fibre Vault).

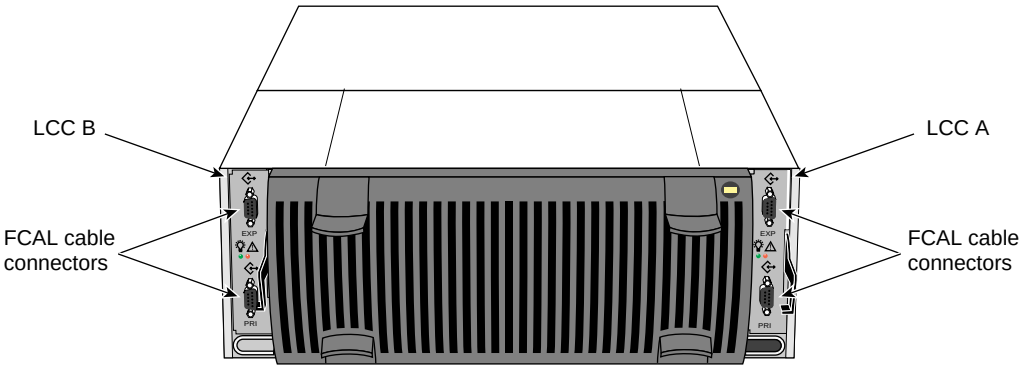

Figure 2-30 Fibre Channel DAE (FibreVault) Connectors and Indicators

Figure 2-31 shows connectors and indicators on the Fibre Channel RAID controller DPE.

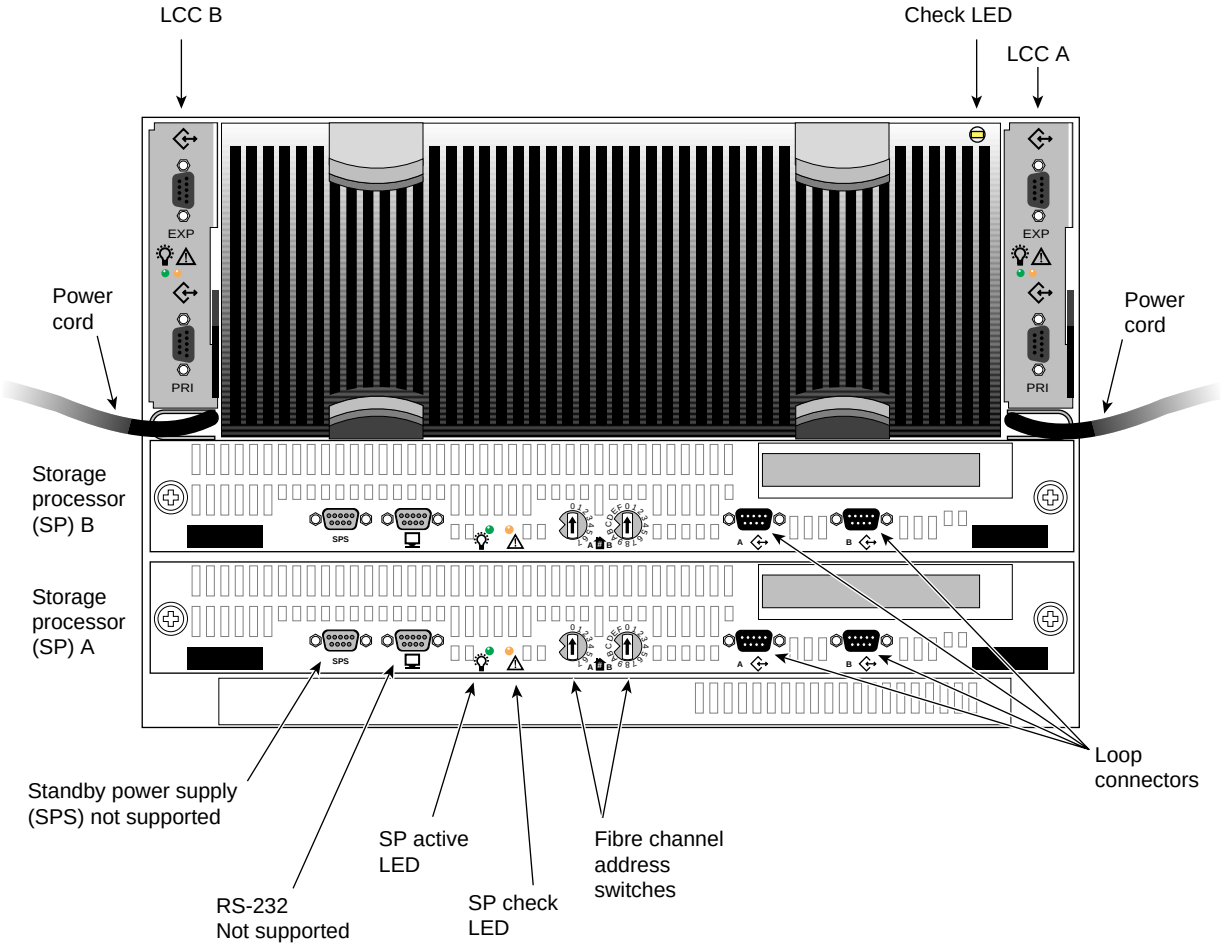

Figure 2-31 Fibre Channel RAID DPE Connectors and Indicators

The storage enclosures use either copper or optical cable of various lengths. Only 10-meter copper cabling is included with a Fibre Channel option board, one per FC port (marketing code X-F-COP-10M; part number 018-0570-001). The customer must order others separately. Optical cables require adapters.

Caution: If a host system and a fibre channel enclosure are on separate building grounds, grounding problems may arise when this cabling is connected with a copper cable utilizing a DC ground shield on both ends. In this case, the customer must use the optional fibre optical cabling to link host and enclosure.

2.10.3 Cabling Fibre Rack Power Supplies for IRIS FailSafe

If the customer's FC enclosures are in a Fibre Channel Rack, you must make sure of the following:

- The rack must have two power distribution units (PDUs).
- Grounding must be correct for the equipment, as explained in Section 2.2.2, "Checking the Grounding." If a host system and a fibre channel enclosure are on separate building grounds, grounding problems may arise when this cabling is connected with a copper cable utilizing a DC ground shield on both ends. In this case, the customer must use the optional fibre optical cabling to link host and enclosure.
- Each power distribution unit must be cabled to a different power strip.
- All SP As for all FC storage enclosures in the rack must be cabled to one PDU and all SP Bs must be cabled to the other PDU.

2.10.4 Cabling Fibre Channel RAID (DPEs) for IRIS FailSafe

Cabling Fibre Channel RAID enclosures for IRIS FailSafe consists of processes explained in

- Section 2.10.4.1, "Cabling DPE Power for IRIS FailSafe"
- Section 2.10.4.2, "Cabling DPE SPs for IRIS FailSafe"
- Section 2.10.4.3, "Setting the Enclosure Address on a Fibre Channel RAID Enclosure"

2.10.4.1 Cabling DPE Power for IRIS FailSafe

Each DPE has two power cords, one for each power supply. For IRIS FailSafe, make sure the power for each DPE power supply is connected to the PDU on that side of the rack, as shown in Figure 2-32.

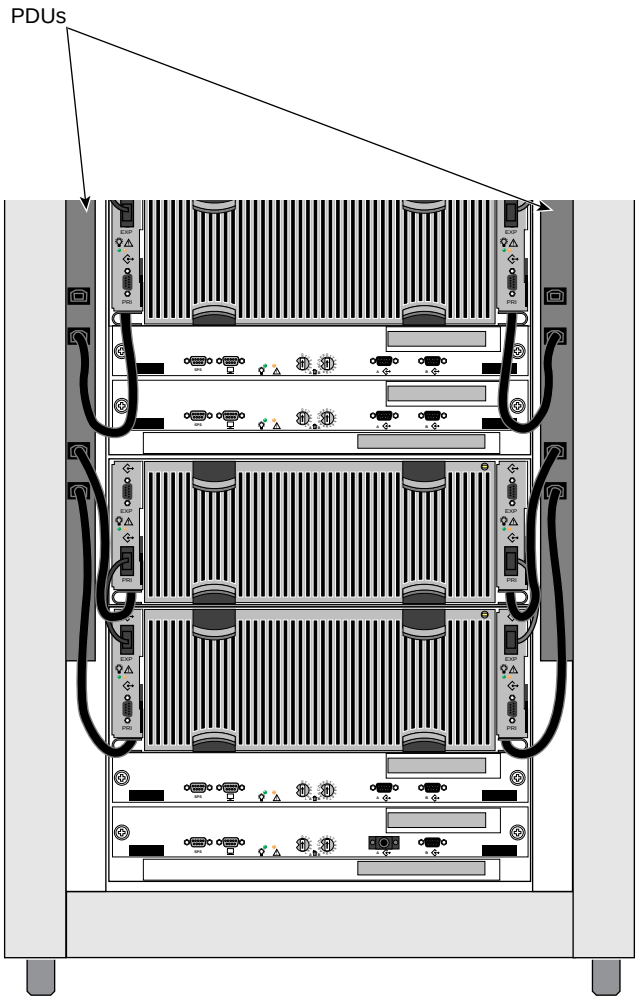

Figure 2-32 FC RAID Power Cabling for IRIS FailSafe

If necessary, consult the *Origin FibreVault and Fibre Channel RAID Installation Instructions* for instructions on removing the fan module and routing the power cords.

2.10.4.2 Cabling DPE SPs for IRIS FailSafe

Each IRIS FailSafe host must be connected to Fibre Channel RAID storage through each controller as diagrammed in Figure 2-33. Each host is connected to each of the two SPs.

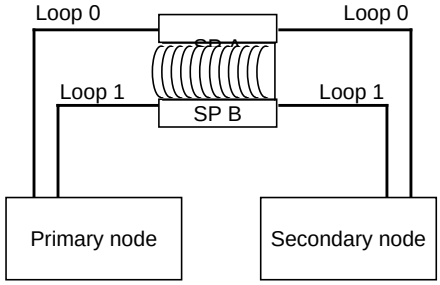

Figure 2-33 Fibre Channel RAID Storage for IRIS FailSafe

Cabling FC RAID DAEs to a DPE is explained in the *Origin FibreVault and Fibre Channel RAID Installation Instructions*.

Several DPEs can be daisy-chained to each other and connected to IRIS FailSafe hosts, as diagrammed in Figure 2-34.

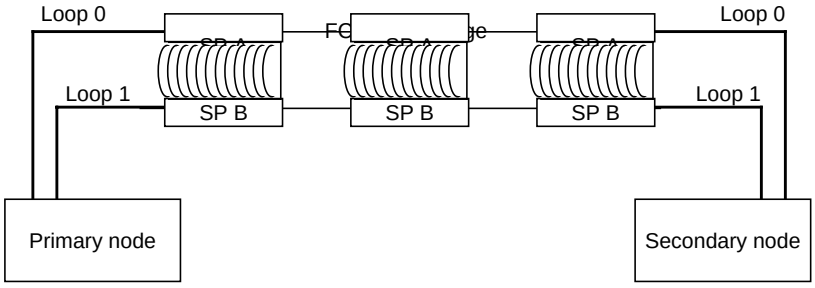

Figure 2-34 Fibre Channel RAID Storage for IRIS FailSafe With Daisy-Chained DPEs

To connect the Fibre Channel enclosure to the FC XIO or PCI board for IRIS FailSafe with copper cabling, follow these steps:

1. Make sure grounding is correct for the equipment, as explained in Section 2.2.2, "Checking the Grounding."

   Caution: If a host system and a fibre channel enclosure are on separate building grounds, grounding problems may arise when this cabling is connected with a copper cable utilizing a DC ground shield on both ends. In this case, the customer must use the optional fibre optical cabling to link host and enclosure.

2. Have ready both copper cables; one 10-meter cable for each port is included with each Fibre Channel option board for each host.

   Note: For instructions on using optical cables, see Section 2.10.6, "Using Optical FC Cables."

3. Attach labels from the board label kit to the cable connector and to the I/O panel for the Fibre Channel option board in the chassis. These labels are similar to those included with the Ultra SCSI board.
4. Attach one end of a copper cable to an FC option board port in the first host.
5. Attach the other end to fibre channel port A on SP A (SCSI bus 0) in the Fibre Channel RAID storage system. Figure 2-35 shows cabling for an Origin2000 server.

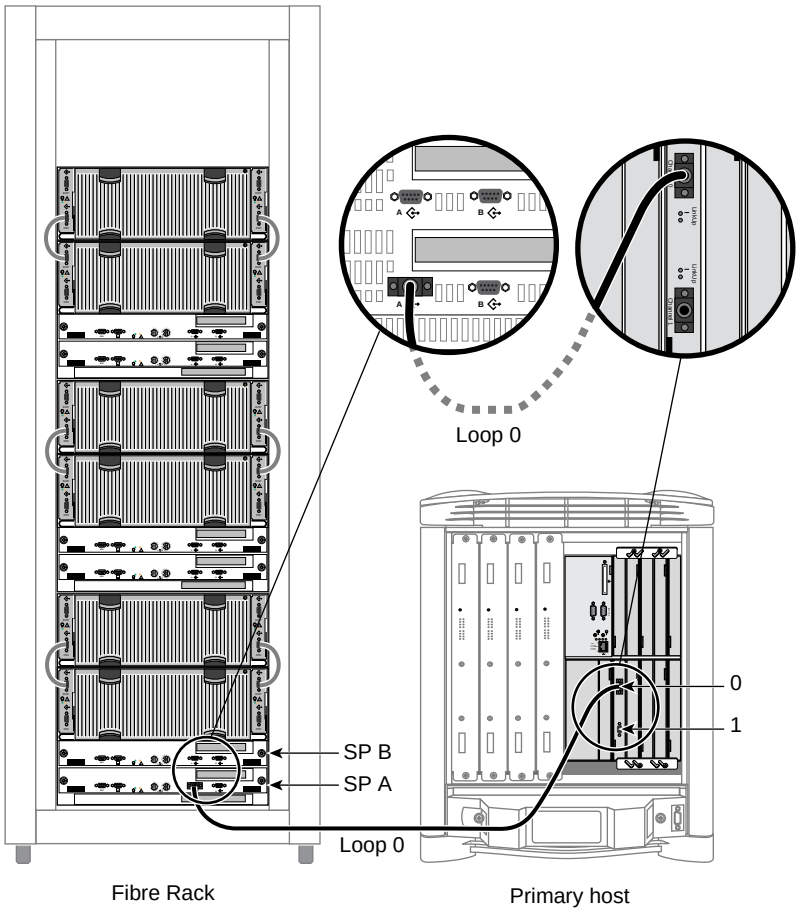

Figure 2-35 Cabling a Fibre Channel XIO Board to a Fibre Channel RAID Enclosure (DPE)

6. Attach one end of the second copper cable to the option board in the second host. Attach the other end to fibre channel port B on SP B (SCSI bus 1) in the Fibre Channel RAID storage system.

Figure 2-36 shows both hosts cabled to the RAID storage system.

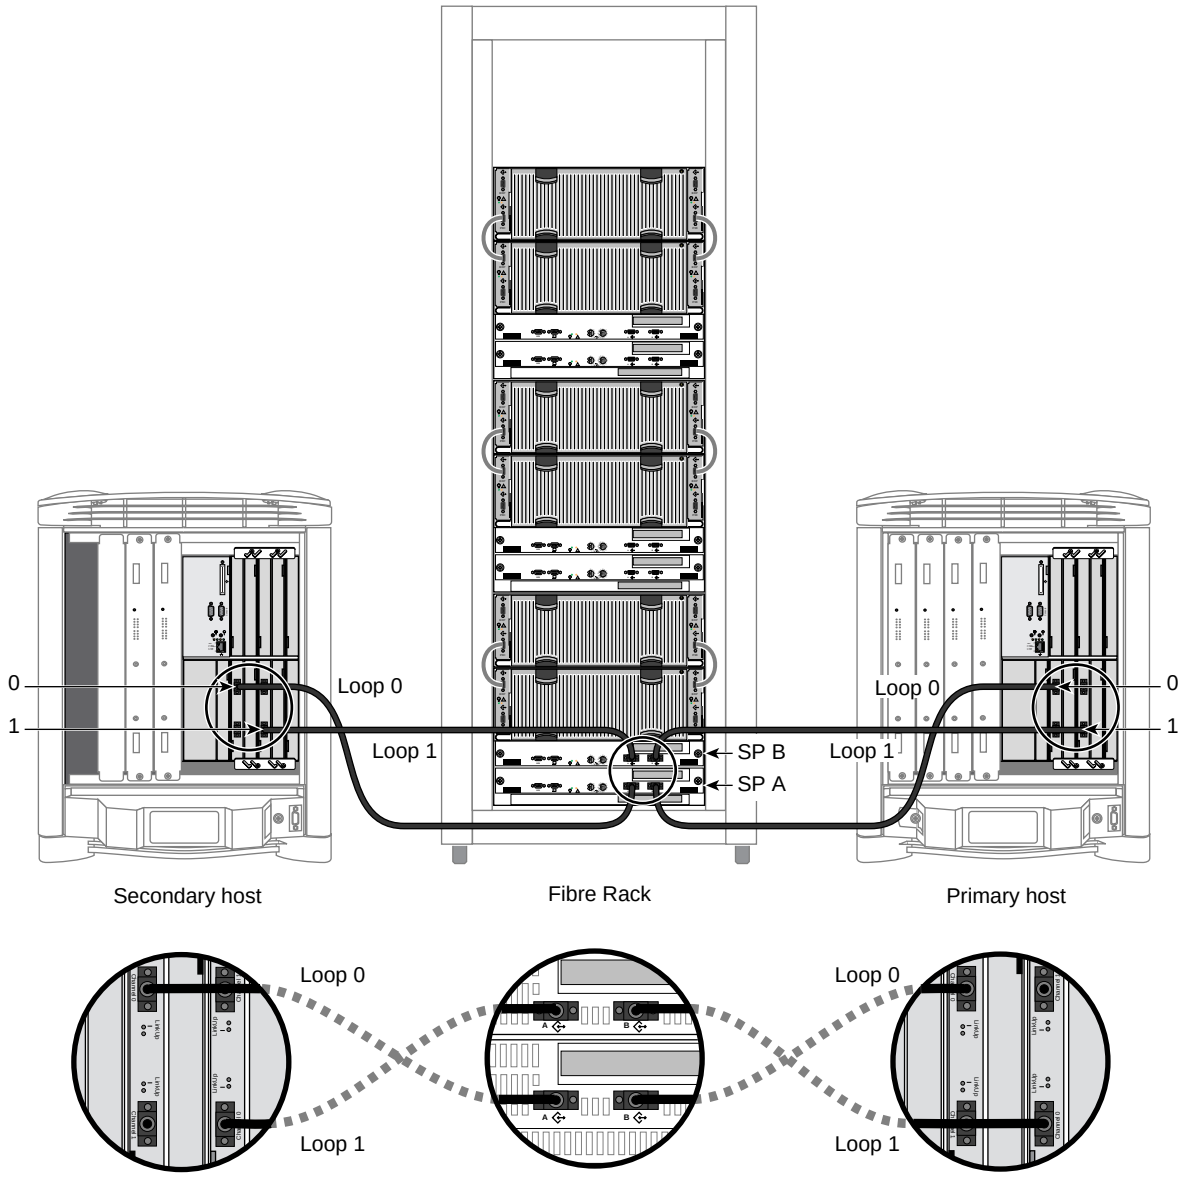

Figure 2-36 Example Fibre Channel RAID Cabling: Origin2000 Servers

Figure 2-37 shows cabling for Origin200 servers.

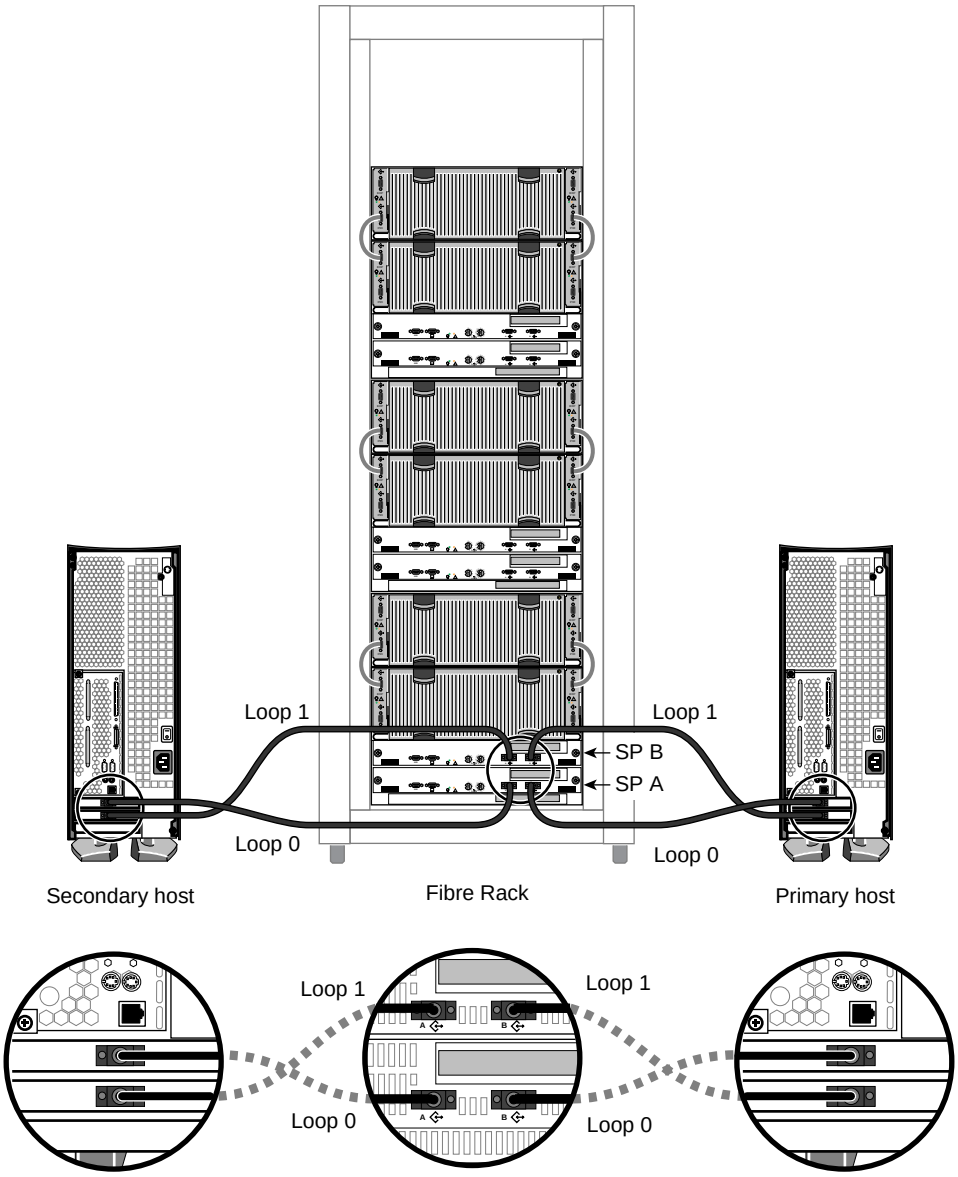

Figure 2-37 Example Fibre Channel RAID Cabling: Origin200 Servers

2.10.4.3 Setting the Enclosure Address on a Fibre Channel RAID Enclosure

Each SP must have a unique fibre channel front-end address. The Fibre Channel RAID enclosure address is set on the system processor (SP) board, as shown in Figure 2-38.

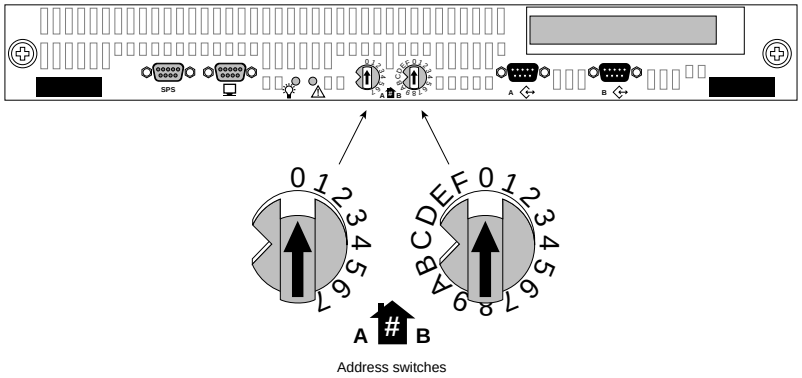

Figure 2-38 Fibre Channel RAID Enclosure Channel Address Switches

Valid address ranges are 0 through 109 (decimal), or 0 through 6D hexadecimal: 0 on both switches through 6 on the left switch and D on the right switch.

Give each SP an address that is unique in the system. To see available addresses, refer to the scsiha(1M) man page.

The disk slots in a FibreVault or Fibre Channel RAID enclosure are numbered 0-9. Disk modules are numbered based on the ID of the Fibre Channel RAID enclosure. Table 2-5 summarizes disk numbering.

Table 2-5 Fibre Channel RAID Enclosure and Disk Slot Numbering

Enclosure ID	Disk Slot Number
0	0-9
1	10-19
2	20-29
3	30-39
...	...
9	90-99
10	100-109

When you are finished setting the address, push the enclosure's front door closed until it latches into place. If desired, lock the enclosure's front door.

2.10.5 Cabling FibreVault (JBOD) Enclosures for IRIS FailSafe

The FibreVault JBOD enclosures have

- no processors (SPs)
- 9.1 GB non-RAID disks: part number 9470140 (512-byte sectors)

  These drives are factory-formatted; you cannot change the format between RAID and non-RAID in the field.
- up to 110 disks per FC loop

This section consists of

- Section 2.10.5.1, “Cabling FibreVault Power for IRIS FailSafe”
- Section 2.10.5.2, “Cabling FibreVault LCCs for IRIS FailSafe”
- Section 2.10.5.3, “Setting Enclosure Addresses on the FibreVault”
- Section 2.10.5.4, “Setting Up Mirroring for FC JBOD Storage”

2.10.5.1 Cabling FibreVault Power for IRIS FailSafe

Each FibreVault DAE (JBOD) enclosure has two power cords, one for each LCC. Make sure the power for each DAE is cabled to a different PDU; each PDU should be cabled to a different power strip.

Figure 2-39 shows DAE power cabling for IRIS FailSafe.

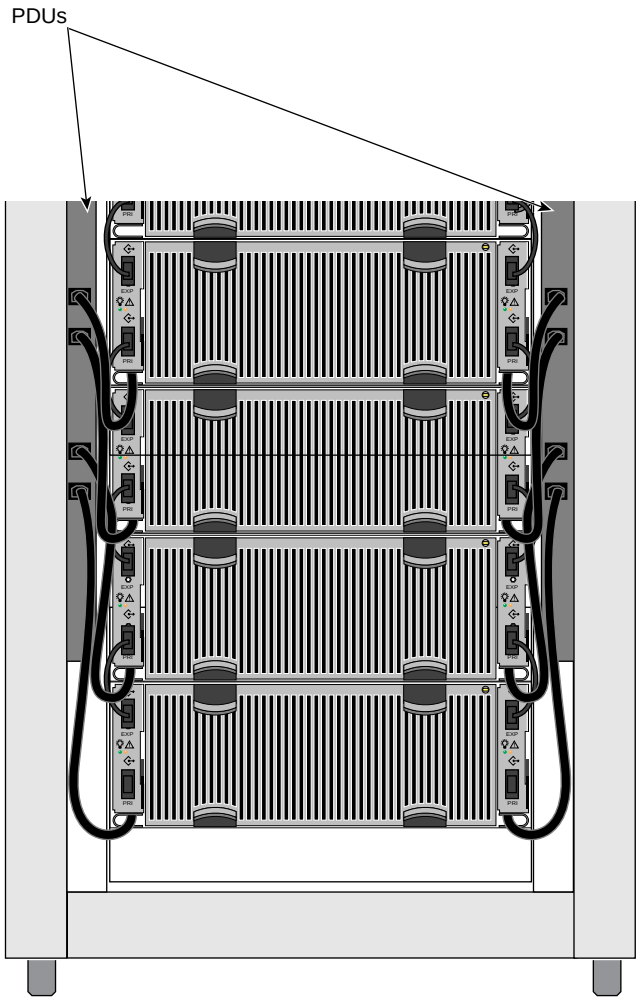

Figure 2-39 FibreVault Power Cabling for IRIS FailSafe

If necessary, consult the *Origin FibreVault and Fibre Channel RAID Installation Instructions* for instructions on removing the fan module and routing the power cords.

2.10.5.2 Cabling FibreVault LCCs for IRIS FailSafe

For JBOD, each IRIS FailSafe host is connected to an LCC on the enclosure; loops 1 and 2 in Figure 2-40 diagram these connections.

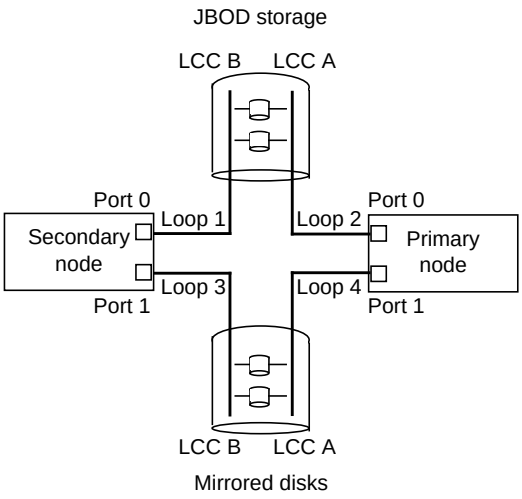

Figure 2-40 FibreVault JBOD Storage for IRIS FailSafe

Note that a configuration using JBOD storage must be mirrored on another JBOD enclosure of equal capacity; loops 3 and 4 in Figure 2-41 diagram these connections. Follow these steps:

1. Make sure grounding is correct for the equipment, as explained in Section 2.2.2, "Checking the Grounding in Configurations Using Fibre Channel Storage."

   Caution: If a host system and a fibre channel enclosure are on separate building grounds, grounding problems may arise when this cabling is connected with a copper cable utilizing a DC ground shield on both ends. In this case, the customer must use the optional fibre optical cabling to link host and enclosure.

2. Have ready four copper cables; one 10-meter cable for each port is included with each Fibre Channel option board for each host.

   Note: For instructions on using optical cables, see Section 2.10.6, "Using Optical FC Cables."

3. Attach labels from the board label kit to the cable connector and to the I/O panel for the Fibre Channel option board in the chassis. These labels are similar to those included with the Ultra SCSI board.

4. Cable the JBOD storage:

   - Connect FC option board port 0 in the primary host to LCC A (SCSI bus 0) in the FibreVault storage system.
   - Connect FC option board port 0 in the secondary host to LCC B in the FibreVault storage system.

Figure 2-41 shows an example.

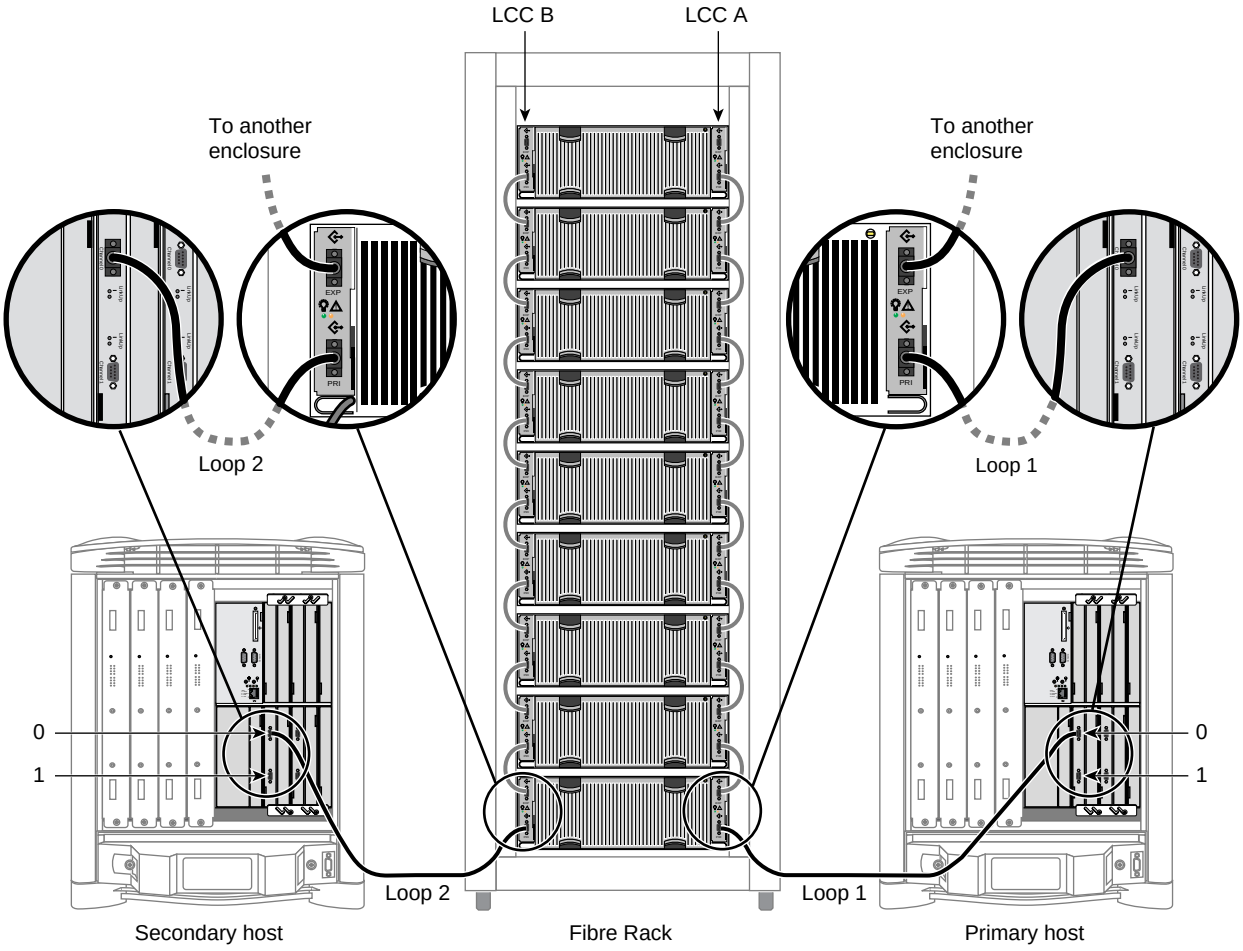

Figure 2-41 Cabling FibreVault Storage for IRIS FailSafe

5. Cable the mirrored JBOD storage:
   - Connect port 1 in the primary host to LCC A in the mirrored JBOD storage.
   - Connect port 1 in the secondary host to LCC B in the mirrored JBOD storage.

   Figure 2-42 shows an example.

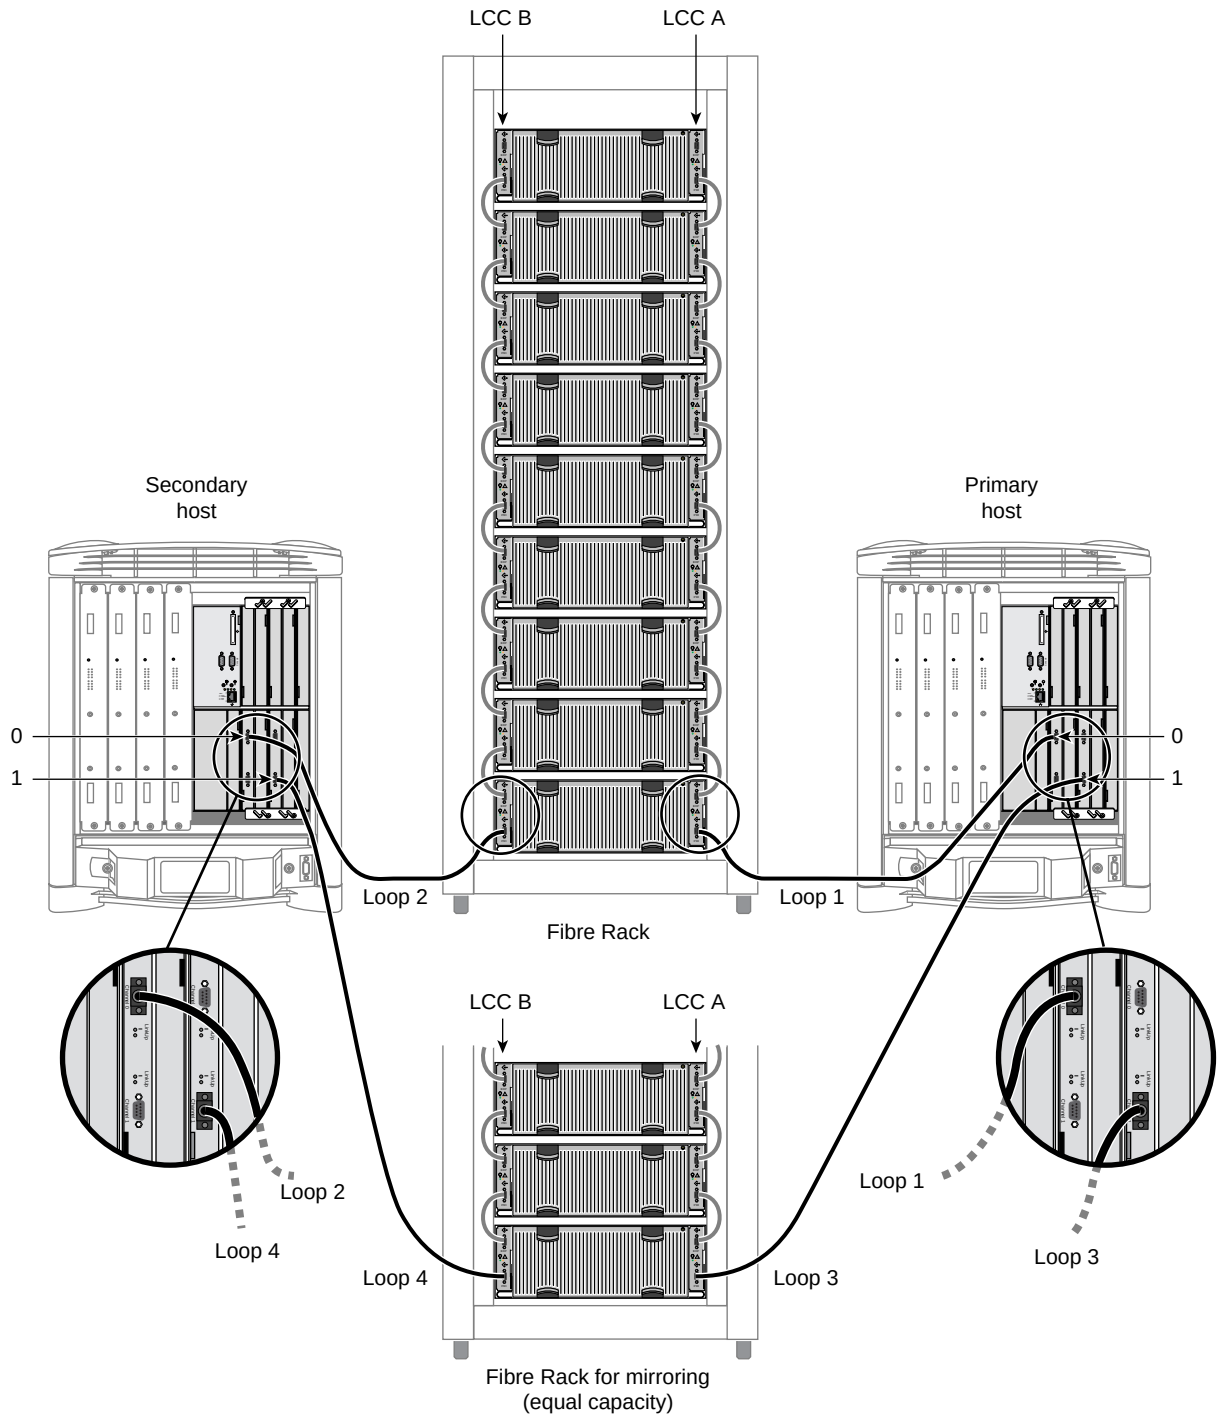

Figure 2-42 Cabling Mirrored JBOD Enclosures

2.10.5.3 Setting Enclosure Addresses on the FibreVault

A drives reads its FC-AL physical address when the drive is powered on, when it is inserted into a powered-on enclosure, or when the drive is reset. Each enclosure address should be set when first installed in the Fibre Channel Rack. Reset enclosure addresses only while the power is off; do not change the address while the enclosure's power is on. Follow these steps:

1. Set the enclosure address for each FibreVault enclosure by using a pencil or ballpoint pen to move the down or up switch. Figure 2-43 shows the address switch.

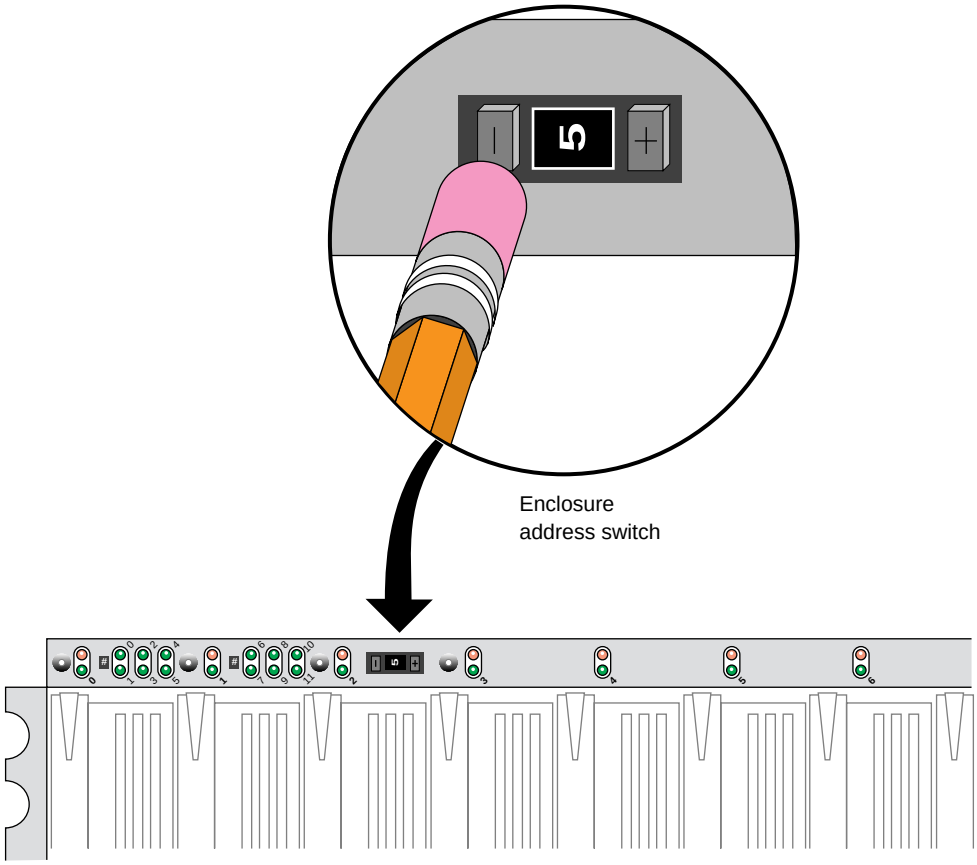

Figure 2-43 Setting the FibreVault ID

The disk slots in a FibreVault enclosure are numbered 0-9 (from left to right). Disk modules are numbered based on the ID of the enclosure. Each enclosure on an FC-AL loop must have a different ID number. Table 2-6 summarizes disk numbering.

Table 2-6 FibreVault Enclosure and Disk Slot Numbering

Enclosure ID	Disk Slot Number
0	0-9
1	10-19
2	20-29

Table 2-6 (continued) FibreVault Enclosure and Disk Slot Numbering

Enclosure ID	Disk Slot Number
3	30-39
...	...
9	90-99
10	100-109

2. When you are finished setting the address, push up the FibreVault front door until it latches into place.
3. Give each SP an address that is unique in the system. To see available addresses, use the scsiha(1M) command.

2.10.5.4 Setting Up Mirroring for FC JBOD Storage

Each FC JBOD enclosure is mirrored with another FC JBOD enclosure in another FC rack with XLV mirroring. The mirroring provides higher availability for the disks. See the latest version of *IRIX Admin: Disks and Filesystems* for information.

2.10.6 Using Optical FC Cables

To connect the Fibre Channel enclosure to the FC XIO or PCI board with optical cabling, the customer must have ordered fiber optic cables and interface adapters (MIAs). There should be one MIA for each end of each cable, totaling four.

Remove the protective caps from each cable's connectors and place an MIA on each connector. This connection is keyed.

Caution: Do not touch the tips of the fiber optic cable connections.

Figure 2-44 shows the MIA and an optical cable connector.

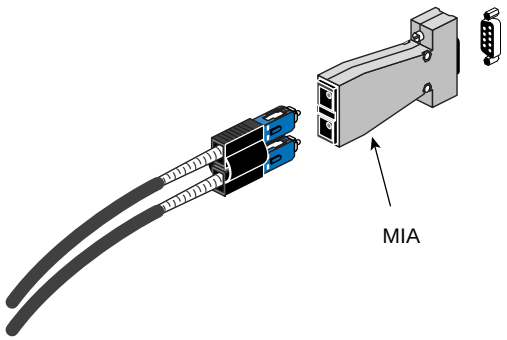

Figure 2-44 Media Interface Adapter (MIA) for Optical Fibre Channel Cables

Note: If the fiber optic cable tips become dirty, clean and dry the tip of each fiber within the cable by gently rubbing the tip with a soft, lint-free cloth that has been moistened with reagent-grade isopropyl alcohol. Do not use prepared cleaning compounds, such as tape-head cleaner or denatured (rubbing) alcohol. If you do not have the proper equipment, skip this step.

Follow steps in Section 2.10.4, "Cabling Fibre Channel RAID (DPEs) for IRIS FailSafe," or Section 2.10.5, "Cabling FibreVault (JBOD) Enclosures for IRIS FailSafe."

2.11 Setting the IRIS FailSafe Host SCSI IDs

If the system has shared SCSI disks, you must set IRIS FailSafe host SCSI IDs. The information in this section also pertains to Fibre Channel disks.

For the IRIS FailSafe system to work, the two servers must have different SCSI IDs. The recommended addresses for an IRIS FailSafe system are 0 for the first IRIS FailSafe host and 7 for the second IRIS FailSafe host, assuming that each node has only one internal drive (SCSI ID 1).

Note: SCSI host address conflicts are a frequent cause of problems getting IRIS FailSafe up and running. Each bus must be free of duplicate SCSI IDs. You might want to diagram the SCSI bus layout of the entire cluster, including the internal buses. The SCSI ID of each host is present on every bus to which it is attached.

Setting the host SCSI IDs consists of these steps:

1. On the first IRIS FailSafe host, enter the command *nvram*. The last line of this command's output should contain the line

   `scsihostid=0`

   This output can also appear without any numeral after the equal sign, which means that the SCSI host ID for this IRIS FailSafe host is 0.

   If this output does not appear, enter **`nvram scsihostid=0`** and reboot this server.

2. On the second IRIS FailSafe host, repeat step 1 in the second window. The last line of the output should be `scsihostid=7`, meaning that the SCSI host ID for the second IRIS FailSafe host is 7.

Note: For IRIS FailSafe configurations using Challenge vaults, SCSI IDs 0 and 2 are suggested. Origin Vault systems, however, have set IDs from 1 to 6 and 9 to 14; thus, it is better to use 0 and 7 as SCSI host IDs so that IDs in the range 1-6 are available for Origin Vault disk drives.

If the second IRIS FailSafe host's SCSI host ID is not 7 (probably the case), follow these steps:

1. Reboot and start the system. Press **`Esc`** to bring up the System Maintenance Menu; choose item 5, the Command Monitor.
2. In the Command Monitor (PROM level), set the AutoLoad variable to Yes:

   `setenv AutoLoad Yes`

3. In the Command Monitor (PROM level), set the SCSI ID for this system:

   ```
   setenv scsihostid 7
   init
   ```

   The command *init* sets the value into the PROM. Choose item 5 again to return to the PROM monitor.

   Caution: No SCSI device on any shared bus should have SCSI ID 0 or 7.

4. To verify that the SCSI ID was set, enter

   ```
   printenv
   ```

   The output should include the line `scsihostid=7`.

5. Exit the System Maintenance Menu and restart the system.

6. Enter the command *hinv* on each server. Check that the disks on the shared SCSI bus appear in the output.

Note: To configure Challenge RAID for the IRIS FailSafe system, see Chapter 6, "Configuring Challenge RAID Storage Systems for IRIS FailSafe."

Chapter 3

Setting Up and Cabling the IRIS FailSafe System With Large Challenge Servers

This chapter explains how to set up and cable the two Challenge DM, L, or XL servers* (IRIS FailSafe hosts) and the Challenge RAID or vault for use as the hardware for the IRIS FailSafe system, in the following sections:

- Section 3.1, “Installing the Software”
- Section 3.2, “Setting Up the Component Systems”
- Section 3.3, “Cabling the Private and Public Networks”
- Section 3.4, “Cabling the Challenge Server Serial Connection”
- Section 3.5, “Testing the Serial Connection”
- Section 3.6, “Cabling the Vaults”
- Section 3.7, “Cabling the Challenge RAID Storage System to the Challenge Servers”
- Section 3.8, “Setting the IRIS FailSafe Host SCSI IDs”
- Section 3.9, “Configuring and Testing the System”

Note: Before installing an IRIS FailSafe system, make sure that the installation site meets the operating limits and AC power requirements as explained in Chapters 3 and 4 of the *CHALLENGE/Onyx Site Preparation Guide.*

The following equipment is required for installation:

- installation guides for the component systems (see “About This Guide” for part numbers)
  - *CHALLENGE/Onyx Site Preparation Guide*
  - *CHALLENGE/Onyx XL Rackmount Installation Instructions*
  - *CHALLENGE/Onyx L Deskside Installation Instructions*
  - *FDDIXpress Mezzanine Board for CHALLENGE and Onyx Installation Instructions*
  - *FDDIXpress Administration Guide*; included in FDDI board shipment

---

* Onyx L and XL systems are also supported as FailSafe nodes, in the same combinations as Challenge L and XL servers.

  - *FDDIXpress Release Notes*; included in FDDI board shipment
  - *CHALLENGE RAID Installation and Maintenance Instructions*
  - *CHALLENGE Vault L Installation Instructions*
  - *CHALLENGE Vault Rack and SCSIBox 2 Installation Instructions*
- laptop, ASCII terminal, or IRISconsole
- Phillips and small flat-blade screwdrivers

Note: Make sure that each Challenge L server used for the IRIS FailSafe system has a Remote System Control port at the extreme right. If it does not, the customer must order an upgrade (013-0624-003) and you must install this port following instructions in the *Remote System Control Port Installation Guide*, included in this shipment.

3.1 Installing the Software

Follow instructions in the latest version of the *IRIS FailSafe Administrator's Guide* (007-3109-003 or later) to install the software needed to run IRIS FailSafe.

For this version of IRIS FailSafe (1.2), the Challenge servers must be running IRIX 6.2.

Depending on the servers and storage in the configuration, you might need to install various patches. For the latest information on compatibility, see http://origin/product/hiavailability/support.html.

3.2 Setting Up the Component Systems

Read through the *CHALLENGE/Onyx Site Preparation Guide* before unpacking equipment at the site. Follow instructions in that manual, specifically:

1. Make sure that the installation site meets the operating limits and AC power requirements as explained in Chapters 3 and 4 of the *CHALLENGE/Onyx Site Preparation Guide*.
2. Make sure required tools and personnel are on hand for unloading and opening crates and for moving large systems.
3. Prepare the physical location to allow for space, air flow, and floor-loading requirements for all component systems, as explained in Chapter 3 of the *CHALLENGE/Onyx Site Preparation Guide*.

   Use only the cables included in the shipment. Plan to situate the servers and vaults fairly close together. Differential SCSI cables, including cabling inside the chassis, should be no longer than 60 feet (18.3 meters).

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

4. Make sure the site meets safety and operating considerations, as explained in Chapter 3 of the *CHALLENGE/Onyx Site Preparation Guide.*
5. Prepare the site for the systems' electrical requirements, as explained in Chapter 4 of the *CHALLENGE/Onyx Site Preparation Guide.*
6. Connect your laptop or ASCII terminal and keyboard to the first Challenge server, as explained in the server installation instructions.
7. Connect the power cord for the first IRIS FailSafe host. Turn on the main power switch on the back of the unit; turn on the laptop or ASCII terminal. If necessary, see the server installation instructions for details.
8. On the IRIS FailSafe host, turn the System Controller key switch to the On position. When power-on diagnostics are completed, a login prompt appears on the console.
9. After the disks are set up, enter **`hinv`** on each IRIS FailSafe host. In the output, look for the installed shared disks and network interfaces.

For the IRIS FailSafe system, you are setting up these networks on each server:

- public network interface(s)

  This network connects the IRIS FailSafe server to clients and the outside world.

- serial connection between Remote System Control port and a serial port

  If one node fails, this connection enables the surviving node to power cycle the other.

- private (heartbeat) network (Ethernet)

  This network, going only between the two nodes, is used exclusively for the keep-alive heartbeat that the two nodes use for monitoring each other's status.

- shared SCSI connection to storage

Figure 3-1 diagrams a basic IRIS FailSafe configuration. Note the four types of connections.

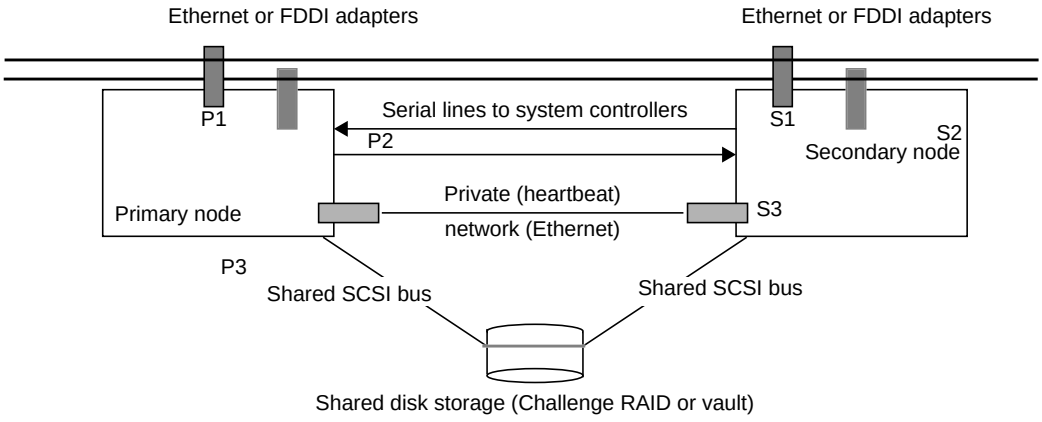

Figure 3-1 IRIS FailSafe Cabling With One Public Network

Consult with the customer on the number and type of public network interface(s) for each server. If a server has two public network interfaces, they can be on the same network or two different networks.

3.3 Cabling the Private and Public Networks

This section explains network cabling and consists of the following:

- Section 3.3.1, “Installing FDDI Boards in the IRIS FailSafe Hosts for a Public Network Connection”
- Section 3.3.2, “Cabling the Private Network”
- Section 3.3.3, “Setting Up a Public Network Ethernet Connection”

  For the public network connection, both connections must be of the same type: both Ethernet or both FDDI. Otherwise, failover does not work. This restriction applies to all Challenge configurations and mixed configurations.

3.3.1 Installing FDDI Boards in the IRIS FailSafe Hosts for a Public Network Connection

This section explains using (optional) FDDI to connect to the network.

Have the latest versions of the following manuals at hand:

- *FDDIXpress Release Notes*
- *FDDIXpress Administration Guide*
- *FDDIXpress Mezzanine Board for CHALLENGE and Onyx Installation Instructions*

Caution: FDDI boards are extremely sensitive and susceptible to damage by electrostatic discharge (ESD), a spark caused by the buildup of electrical static potential on clothing and other material. You must use proper ESD preventive measures.

To install the FDDI boards in the servers, follow these steps:

1. Check the contents of the FDDI board shipment, as explained in Chapter 1 of the *FDDIXpress Mezzanine Board Installation Instructions*.
2. Use the *versions* command to check if the system into which you are installing the FDDIXPress™ mezzanine board is already running FDDIXPress software. If FDDIXPress is not installed, install it following instructions in the *FDDIXpress Release Notes*.
3. Install the FDDIXpress board following instructions in Sections 2.1 and 2.2 in Chapter 2 (Challenge L) or Sections 3.1 and 3.2 in Chapter 3 (Challenge XL) of the *FDDIXpress Mezzanine Board for CHALLENGE and Onyx Installation Instructions*.
4. Cable the connections and complete the installation as explained in Chapter 4 of the *FDDIXpress Mezzanine Board for CHALLENGE and Onyx Installation Instructions*.

3.3.2 Cabling the Private Network

The private network between the Challenge servers supplies the heartbeat of one of the servers to the other. To cable the private network, connect the Ethernet 10-Base-T (**ec0**) ports on the Challenge servers to each other with the Ethernet loopback cable supplied with the IRIS FailSafe system (9290131, 20-foot, or 9290132, 40-foot). For the Challenge L, which does not have any 10-Base-T ports, use a transceiver.

3.3.3 Setting Up a Public Network Ethernet Connection

Connect the public network drop cable to the Ethernet AUI port or to an optional Ethernet port on each Challenge server. For the Challenge L, which does not have any 10-BASE-T ports, use a transceiver if necessary. For information on installing an E-Plex board, see *POWER Channel E-Plex Installation Instructions* (108-0114-001).

3.4 Cabling the Challenge Server Serial Connection

Cross-cable the Remote System Control port to a serial port on the Challenge servers for the serial connection between the two. Follow these steps:

1. Attach a serial cable to a serial port on the first IRIS FailSafe host.
2. Connect the other end of this cable to the 9-pin Remote System Control connector on the second IRIS FailSafe host, as shown in Figure 3-2.

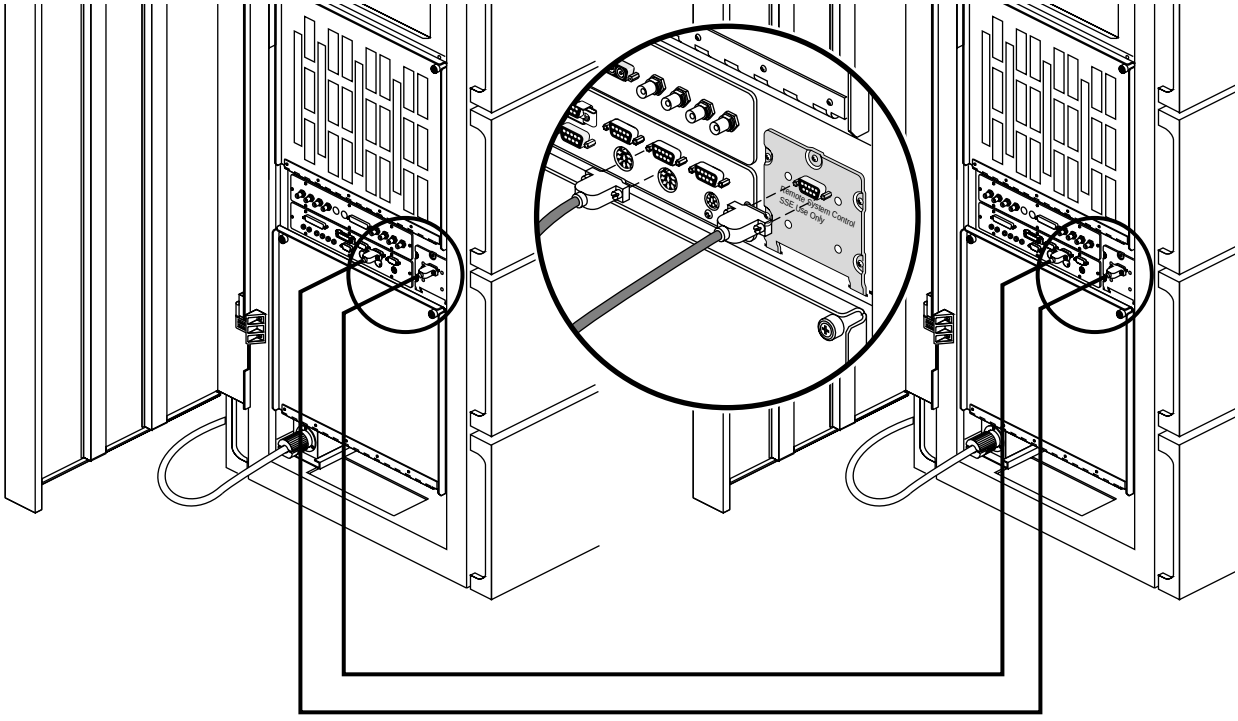

Figure 3-2 Serial Connection: Challenge Servers

3. Attach the other serial cable in the kit to the first host's Remote System Control connector and a serial port (for example, tty_2) on the second host. Figure 3-2 shows the serial cabling.

3.5 Testing the Serial Connection

To test the serial connection between the IRIS FailSafe servers, follow these steps:

1. Make sure the IRIS FailSafe servers are powered on.
2. Enter

   ```
   /etc/init.d/failsafe stop
   ```

3. Test the connection; for example, if the serial cable was connected to tty_2, enter

   ```
   /usr/etc/ha_spng -i 10 -f /dev/ttyd2
   ```

   No output appears; check the return value of the command. If the return value is 0, the connection is good.

   If the return value is 1, perform these checks:

   - Verify that the IRIS FailSafe server is powered on.
   - Verify the cable connections from one server's serial port or remote power control unit and the other server's System Console port.

4. Repeat the process on the second node.

If results are unsatisfactory, make sure that the RJ connector is completely seated at both ends and is making good contact.

3.6 Cabling the Vaults

The numbering of SCSI controllers (ports) on the IRIS FailSafe host IO4 board and mezzanine card(s) (daughter boards), if present, is standardized. Figure 3-3 diagrams the numbering scheme. Note that the numbering is constant regardless of which daughter board is installed.

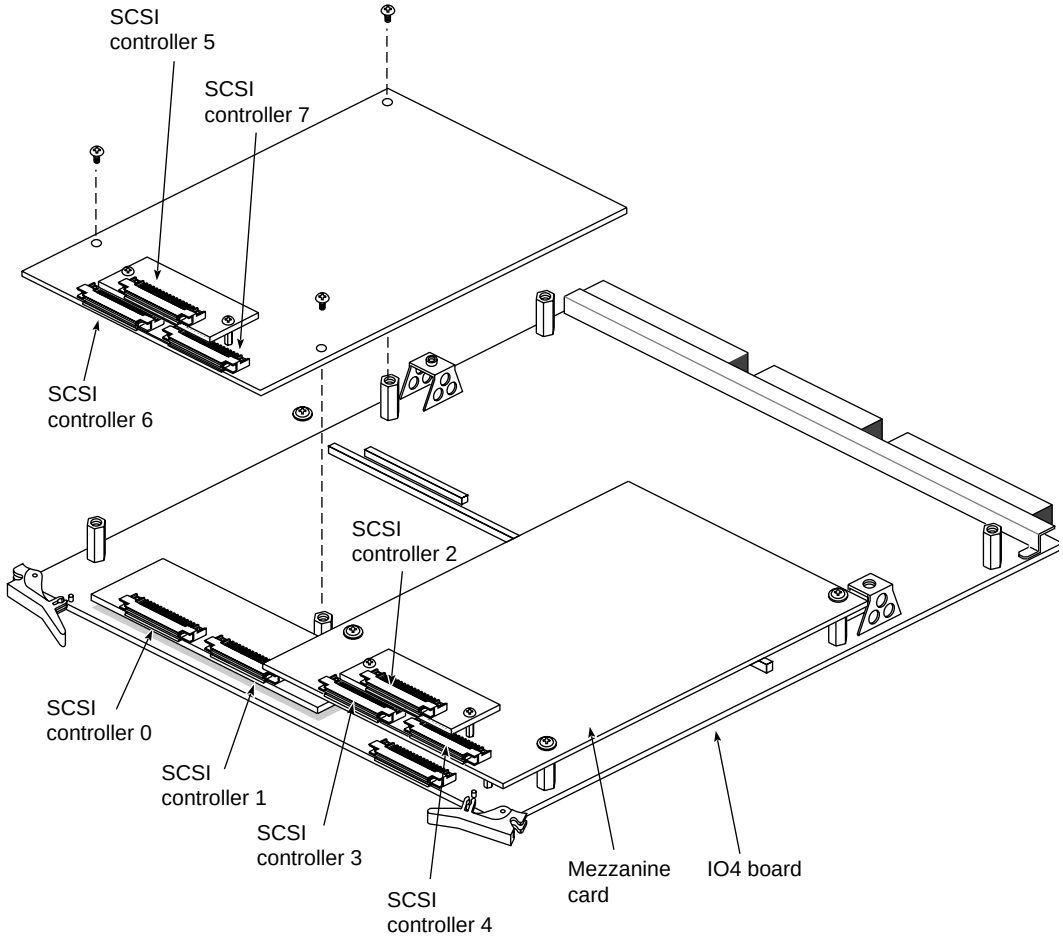

Figure 3-3 SCSI Controller Numbering on IO4 and Mezzanine Boards

In an IRIS FailSafe configuration, each SCSI bus runs from a controller on one Challenge server through a vault (or Challenge RAID) to the second IRIS FailSafe host. Figure 3-4 shows IRIS FailSafe vault cabling.

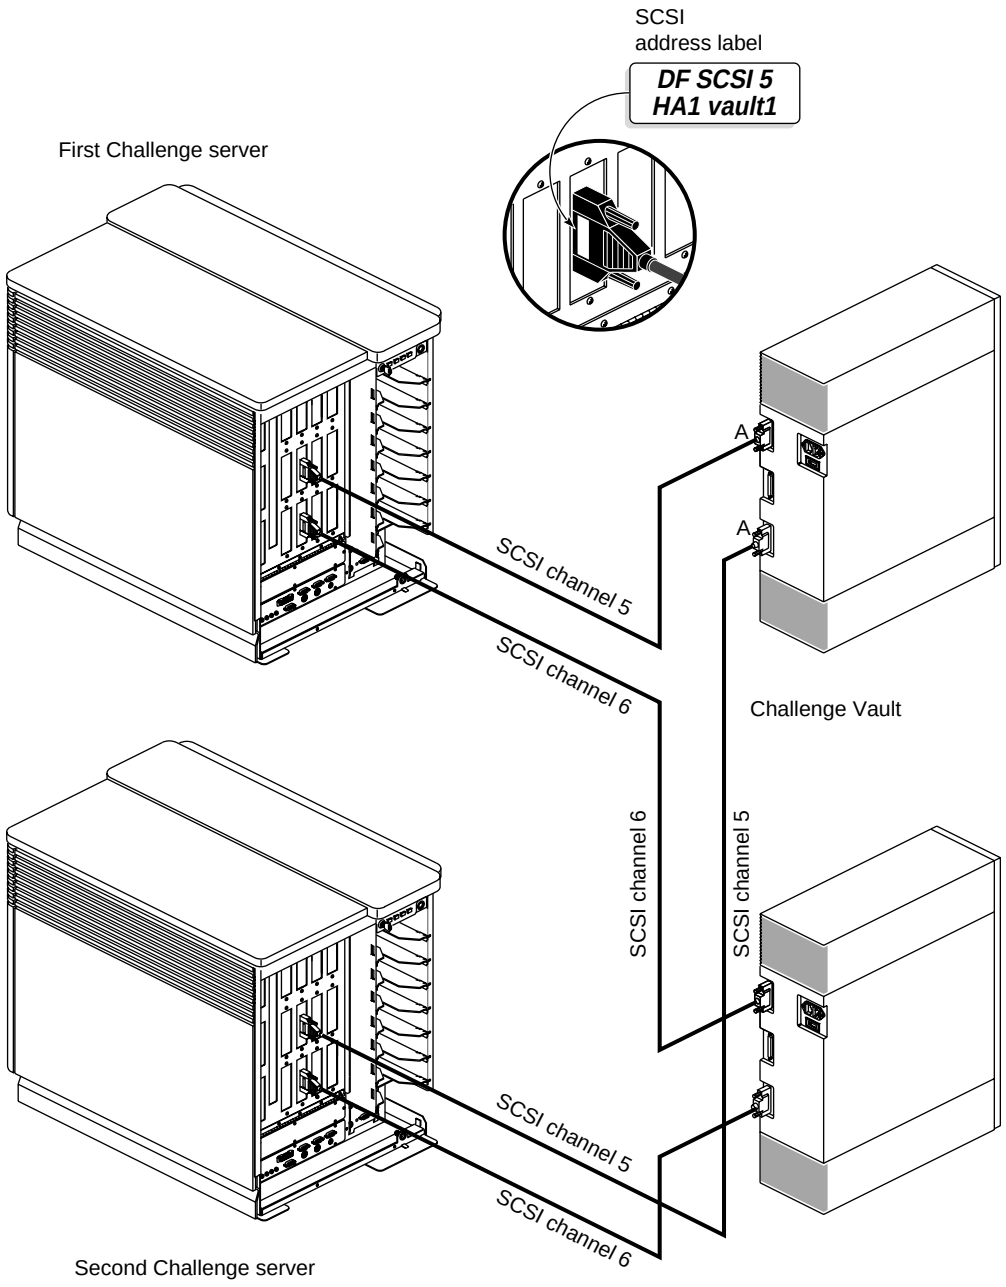

Figure 3-4 Cabling Challenge Vaults and Challenge Servers

To cable the SCSI buses on the vaults to each IRIS FailSafe host, follow these steps:

1. Power off both servers and the vaults.
2. On the back of the vault, insert one of the cables into the top SCSI channel A socket, as shown in Figure 3-4.

3. Connect the other end of this cable to a SCSI bus port on the first IRIS FailSafe host (for example, DF SCSI 5), as shown in Figure 3-4. Label the connector and cable end.
4. Insert a cable into the lower SCSI channel A port on the back of the vault. Label the cable end.
5. Connect the other end of this cable to the same SCSI bus port on the second IRIS FailSafe host that you used on the first host (shown as DF SCSI 5 in Figure 3-4).
6. Repeat steps 2 through 5 for the second vault.
7. Power on the vaults and servers.
8. After the disks are set up, do a dummy I/O to test each link. For example, from the first host, read data on disk 8 on SCSI bus 5 (*dks5d8*) to see if the disk light for disk 8 on the vault lights up. Repeat this test from the second IRIS FailSafe host and check if the same LED is activated.

3.7 Cabling the Challenge RAID Storage System to the Challenge Servers

Table 3-1 charts SCSI ID switch settings for the Challenge RAID storage system.

Note: Challenge RAID SCSI ID switch settings do not conform to frequently used numbering schemes.

Table 3-1 Challenge RAID SCSI ID Switch Settings

SCSI ID Number	Switch Number			
	ID 0	ID 1	ID 2	ID 3
1	Off	On	On	On
2: Do not use	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

Table 3-1 (continued) Challenge RAID SCSI ID Switch Settings

SCSI ID Number	Switch Number			
	ID 0	ID 1	ID 2	ID 3
14	On	Off	Off	Off
15	Off	Off	Off	Off

To cable the Challenge RAID storage system to the Challenge servers, follow these steps:

1. Have ready four 20-foot (or shorter) SCSI cables. Figure 3-5 shows the Challenge RAID SCSI ports.

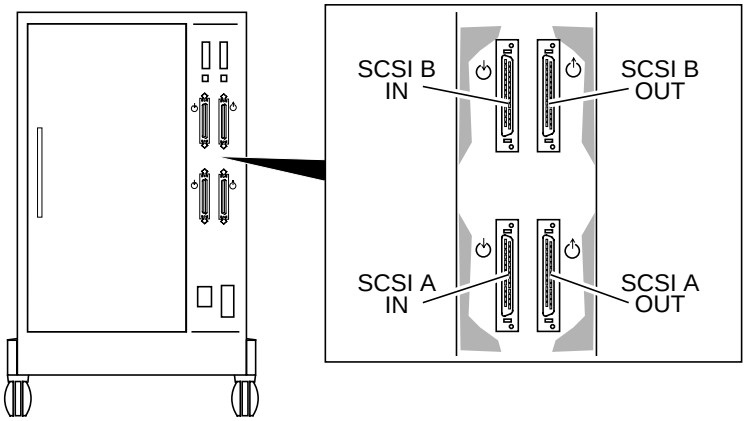

Figure 3-5 SCSI-2 Bus Connectors on Back of Challenge RAID Chassis

2. On the back of the Challenge RAID storage system, attach a SCSI cable to the Challenge RAID **SCSI A** in port, as shown in Figure 3-6.

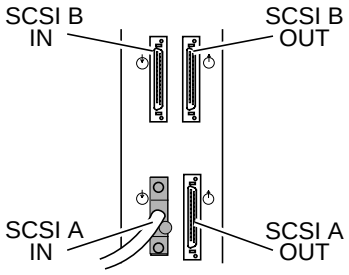

Figure 3-6 Connecting a SCSI Bus Cable to a Challenge RAID SCSI Port

3. Connect the other end of this cable to a SCSI bus (for example, DF SCSI 4) on the first host, as shown in Figure 3-7.

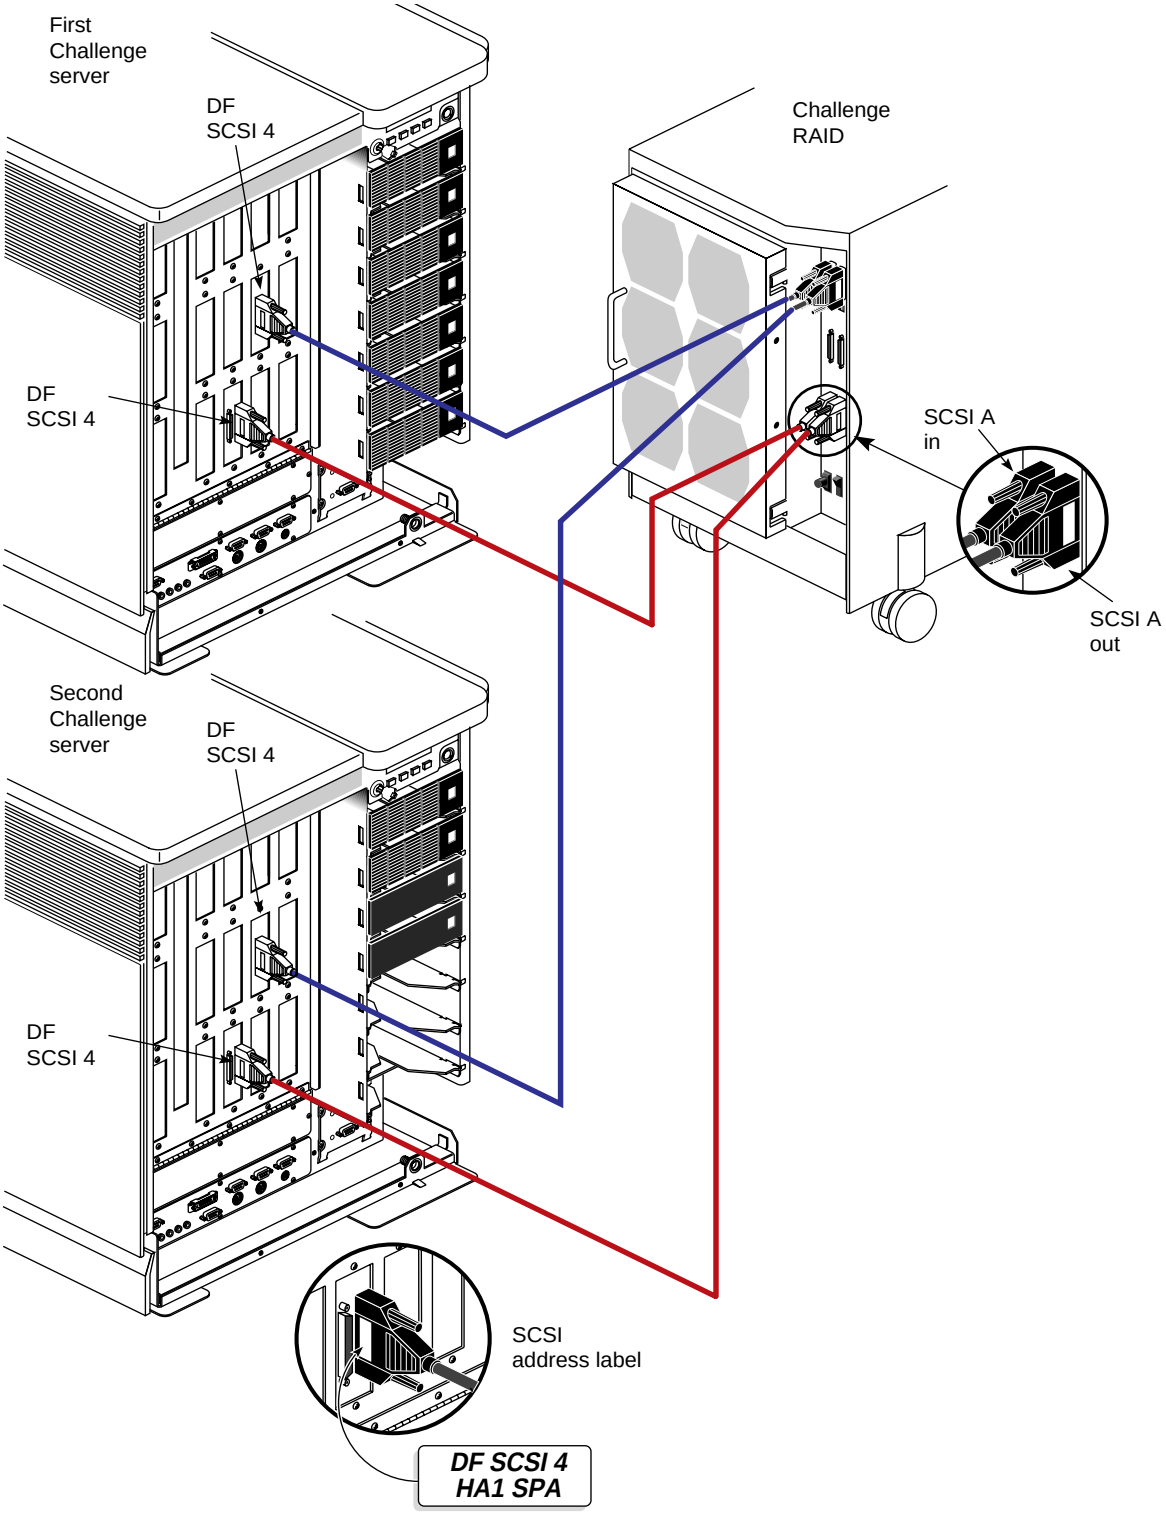

Figure 3-7 Cabling Challenge RAID and Challenge Servers

4. On the back of the Challenge RAID storage system, attach another SCSI cable to the Challenge RAID **SCSI A** out port.
5. Connect the other end of this cable to a SCSI bus (for example, DF SCSI 4) on the second host.
6. Repeat steps 2 through 5 for the SCSI B port on the Challenge RAID storage system.

3.8 Setting the IRIS FailSafe Host SCSI IDs

For the IRIS FailSafe system to work, the two Challenge servers must have different SCSI IDs. The recommended addresses for an IRIS FailSafe system are 0 for the first IRIS FailSafe host and 2 for the second IRIS FailSafe host, assuming that this node has only one internal drive.

Note: SCSI host address conflicts are a frequent cause of problems getting IRIS FailSafe up and running. Each bus must be free of duplicate SCSI IDs. You might want to diagram the SCSI bus layout of the entire cluster, including the internal buses. The SCSI ID of each host is present on every bus to which it is attached.

Setting the host SCSI IDs consists of these steps:

1. On the first IRIS FailSafe host, enter the command *nvram*. The last line of this command's output should contain the line

   `scsihostid=0`

   This output can also appear without any numeral after the equals sign, which means that the SCSI host ID for this IRIS FailSafe host is 0.

   If this output does not appear, enter **`nvram scsihostid=0`** and reboot this server.
2. On the second IRIS FailSafe host, repeat step 1 in the second window. The last line of the output should be `scsihostid=2`, meaning that the SCSI host ID for the second IRIS FailSafe host is 2.

If the second IRIS FailSafe host's SCSI host ID is not 2 (probably the case), follow these steps:

1. Reboot and start the system. Press **`Esc`** to bring up the System Maintenance Menu; choose item 5, the Command Monitor.
2. In the Command Monitor, set the SCSI ID for this system:

   ```
   setenv scsihostid 2
   init
   ```

   The command *init* sets the value into the PROM. Choose item 5 again to return to the PROM monitor.

   Caution: No SCSI device on any bus on the second server should have SCSI ID 2.
3. To verify that the SCSI ID was set, enter

   `printenv`

   The output should include the line `scsihostid=2`.

4. Exit the System Maintenance Menu and restart the system.
5. Enter the command *hinv* on the first server.
6. Enter the command *hinv* on the second server. Compare the output for each server.

3.9 Configuring and Testing the System

Follow instructions in the *IRIS FailSafe Administrator's Guide* to configure and test the newly installed IRIS FailSafe system. If necessary, also consult the latest edition of *IRIX Admin: Disks and Filesystems*.

Caution: You must stop the RAID agent before creating XLV volumes.

To configure Challenge RAID for the IRIS FailSafe system, see Chapter 6, "Configuring Challenge RAID Storage Systems for IRIS FailSafe."

Chapter 4

Setting Up and Cabling the IRIS FailSafe System With Challenge S Servers

This chapter explains how to set up an IRIS FailSafe system for which at least one of the servers is a Challenge S server. The process is described in these sections:

- Section 4.1, “Installing the Software”
- Section 4.2, “Setting Up the Component Systems”
- Section 4.3, “Cabling the Private and Public Networks”
- Section 4.4, “Setting Up the IRIS FailSafe Serial Connection”
- Section 4.5, “Testing the Installed IRIS FailSafe Hardware”
- Section 4.6, “Testing the Serial Connection”
- Section 4.7, “Cabling Storage Systems”
- Section 4.8, “Setting the IRIS FailSafe Host SCSI IDs”
- Section 4.9, “Configuring and Testing the System”

Note: Before installing a IRIS FailSafe system, particularly one involving larger Challenge servers and vaults, make sure that the installation site meets the operating limits and AC power requirements as explained in Chapters 3 and 4 of the *CHALLENGE/Onyx Site Preparation Guide.*

The following equipment is required for installation:

- installation guides for the component systems
- laptop, ASCII terminal, or IRISconsole
- Phillips and small flat-blade screwdrivers

4.1 Installing the Software

Follow instructions in the latest version of the *IRIS FailSafe Administrator’s Guide* (007-3109-003 or later) to install the software needed to run IRIS FailSafe.

4.2 Setting Up the Component Systems

For the IRIS FailSafe system, you must set up

- servers
- shared storage system
- public network interface: one or two per node
- private (heartbeat) network interface (Ethernet 10-Base-T port): one per node
- shared SCSI buses to shared storage: two per node for mirroring
- serial connection through the remote power control unit: one per node

Figure 4-1 diagrams basic IRIS FailSafe cabling.

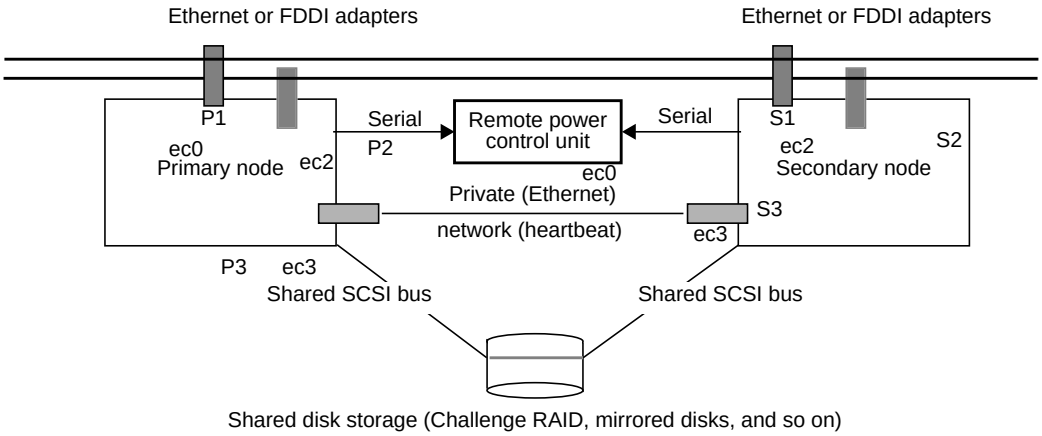

Figure 4-1 IRIS FailSafe Cabling Scheme With One Public Network (Two Challenge S Servers)

To set up the component systems, follow these guidelines:

1. Make sure that the installation site meets the operating limits and AC power requirements of the equipment.

   Note: If an L-class Challenge server is involved, follow guidelines in Section 3.2, "Setting Up the Component Systems," in Chapter 3.

2. Prepare the physical location to allow for space, air flow, and floor-loading requirements for all component systems.

   Use only the cables included in the shipment. Plan to situate the servers and vaults fairly close together. Differential SCSI cables, including cabling inside the chassis, should be no longer than 60 feet (18.3 meters).

   Note: Although the SCSI bus absolute length limit is 80 feet, exceeding 60 feet is not recommended.

3. Make sure the site meets safety and operating considerations.
4. Prepare the site for the systems' electrical requirements.

5. Install the second network adapter (**ec2**) in the Challenge S server if necessary. Follow instructions in the *CHALLENGE S Server Owner's Guide.*
6. Set up the shared storage system.
7. For testing, have ready a laptop or ASCII terminal and keyboard.
8. Make provisions for the network connections, such as obtaining Ethernet drop cables.

Figure 4-2 shows the ports on the Challenge S server for the private and public network interfaces and for the fast and wide SCSI interface (for the Challenge RAID or Challenge Vault storage system).

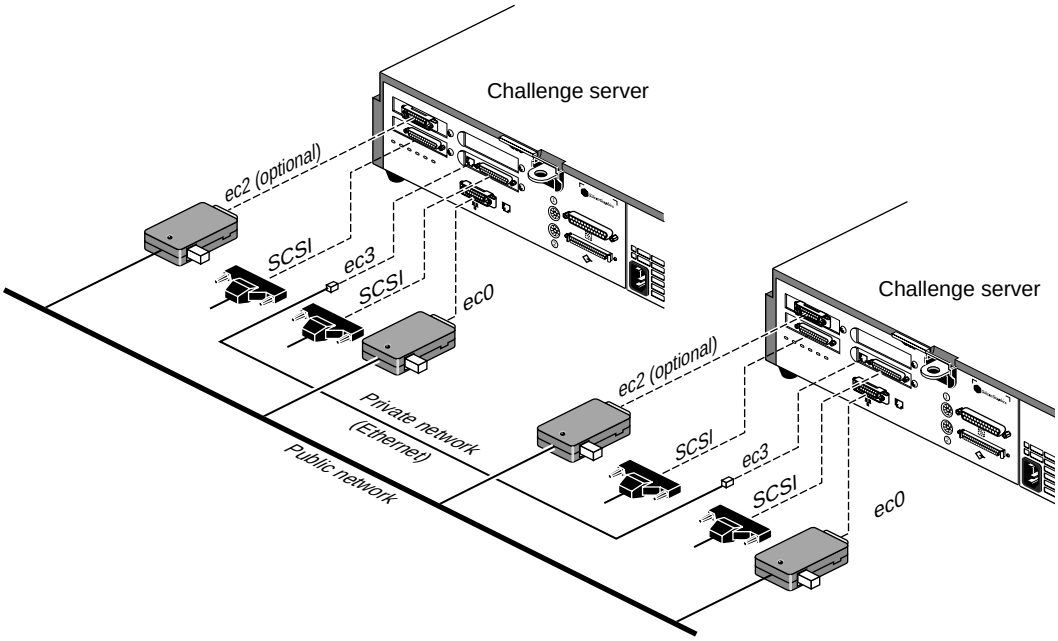

Figure 4-2 Challenge S Server Ports: Private and Public Network Cabling

Figure 4-3 shows an example of serial connection for sites using two Challenge S servers.

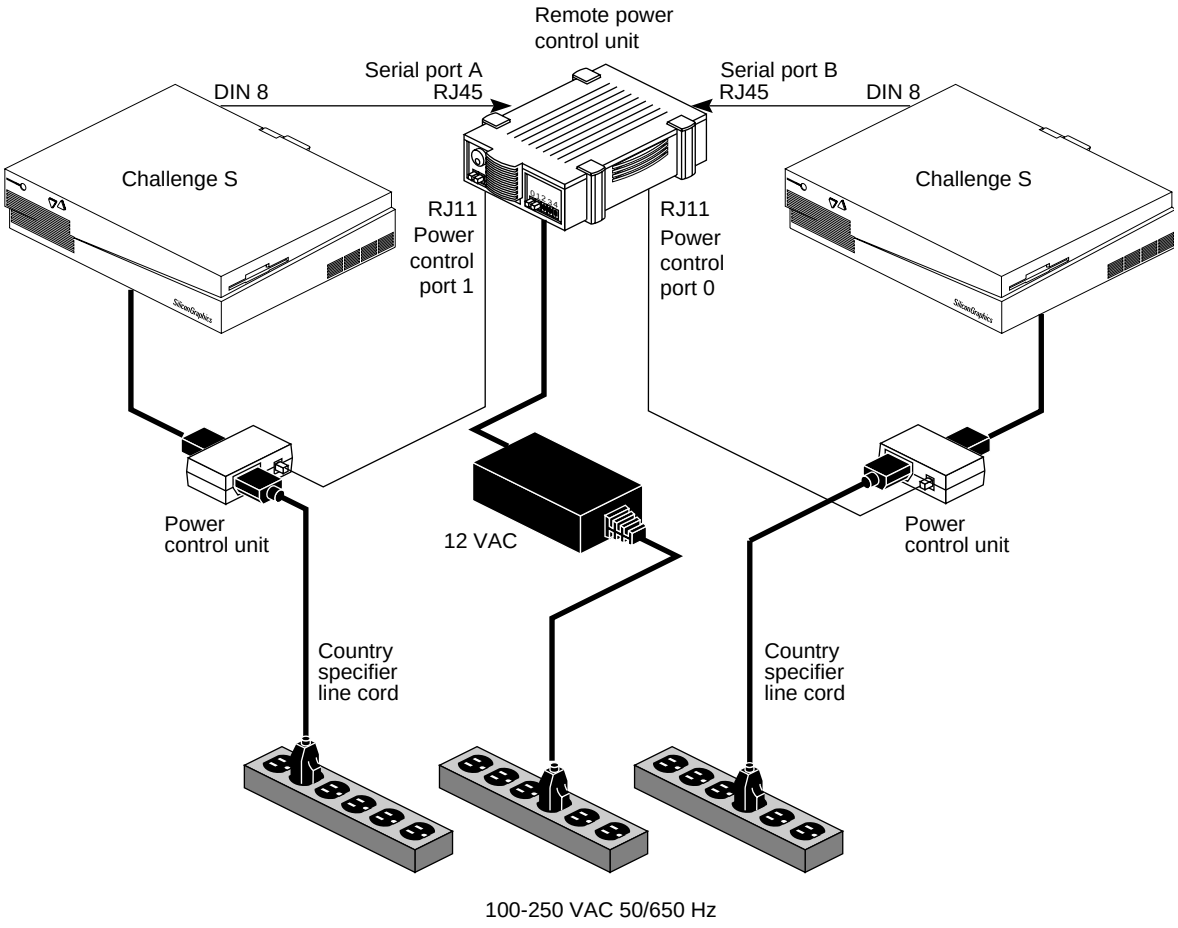

Figure 4-3 Serial and Power Connections: Two Challenge S Servers

4.3 Cabling the Private and Public Networks

This section consists of the following:

- Section 4.3.1, “Cabling the Private Network”
- Section 4.3.2, “Cabling the Public Network”

4.3.1 Cabling the Private Network

The private network between the Challenge servers supplies the heartbeat of one of the servers to the other. To cable the private network, connect the Ethernet 10-Base-T (**ec3**) ports on the Challenge S servers to each other with the Ethernet loopback cable supplied with the IRIS FailSafe product.

This port is between the two SCSI connectors on the back of the Challenge S server, as shown in Figure 4-4. Note that the Ethernet 10-Base-T port and the ISDN port below it look similar. The Ethernet 10-Base-T port is the one on top, immediately next to the SCSI connector, as shown in Figure 4-4.

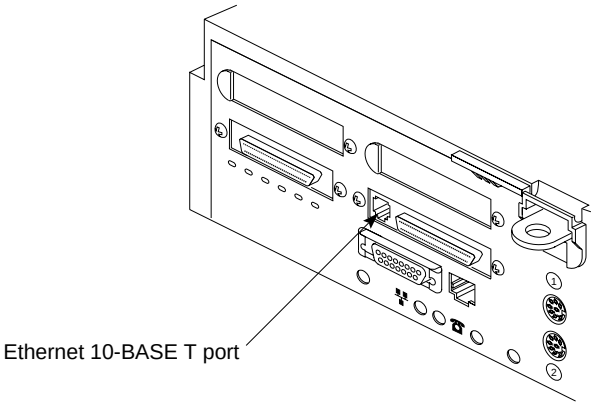

Figure 4-4 Connecting an Ethernet 10-Base-T Cable to the Challenge S Server (S-100 and S-150 Models) for the Private Network (Heartbeat)

If one node in the configuration is a Challenge DM, L, or XL server, an external 10-base T transceiver is required. For example, in a Challenge L to Challenge S configuration, plug the Ethernet null-modem cable into the 10-Base-T port on the Challenge S server and the other end into a 10-Base-T transceiver, which in turn is connected to the Ethernet port on the Challenge L server.

4.3.2 Cabling the Public Network

Connect the public network drop cable to the Ethernet AUI port (**ec0**) on the back of the Challenge S server.

In the Challenge S server, you can configure a single FDDI adapter, which takes up both GIO slots. Thus, if you use FDDI you can have only one connection to the public network.

4.4 Setting Up the IRIS FailSafe Serial Connection

The serial connection between the two Challenge servers in the IRIS FailSafe system makes it possible for one server to reboot the other in case of failure. In the case of two larger servers, the Remote System Control port of one server is cabled to a serial port of the other, and vice versa. Because the Challenge S has no Remote System Control port, the Silicon Graphics remote power control unit and its associated power control units take the place of this port for purposes of the serial connection.

This section explains how to cable the servers to the remote power control unit, consisting of the following:

- Section 4.4.1, “Cabling Two Challenge S Servers”
- Section 4.4.2, “Cabling a Challenge S and a Larger Server”

4.4.1 Cabling Two Challenge S Servers

Figure 4-3 earlier in this chapter shows an example of serial connection for sites using two Challenge S servers.

Note: On the remote power control unit, the leftmost Serial Port (A) controls the leftmost Power Control Port (0), and the rightmost Serial Port (B) controls the next Power Control Port (1). Note that the jacks are different for the two types of ports, as shown in Figure 4-5.

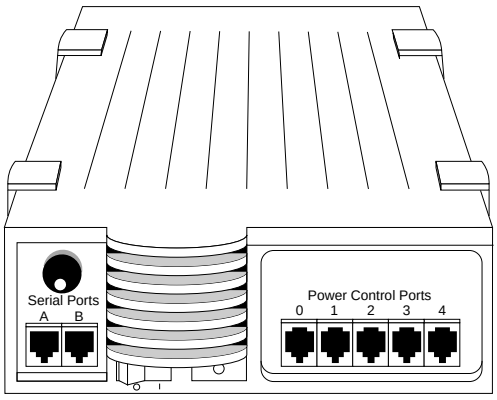

Figure 4-5 Connector Panel on Remote Power Control Unit

Note that Power Control Ports 2 through 4 are not used; they are supplied with terminator plugs to reduce the chance for error.

If both nodes in your IRIS FailSafe configuration are Challenge S servers, follow these steps:

1. Plug one of the two serial cables included with the remote power control unit that has an RJ45 connector into Serial Port A on the remote power control unit, as shown in Figure 4-6. This cable is labeled **CHALLENGE S TO RPCU CABLE**.

   Note: For all RJ45 and RJ11 connections, be sure the connector is properly seated in the jack so that it makes good contact.

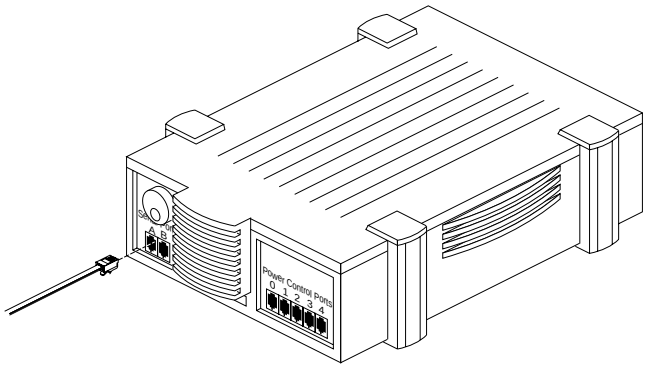

Figure 4-6 Cabling the Serial Port on the Remote Power Control Unit

2. Plug the other end of this serial cable into the serial connector (**tty_2**) on the back of the Challenge S server, as shown in Figure 4-7.

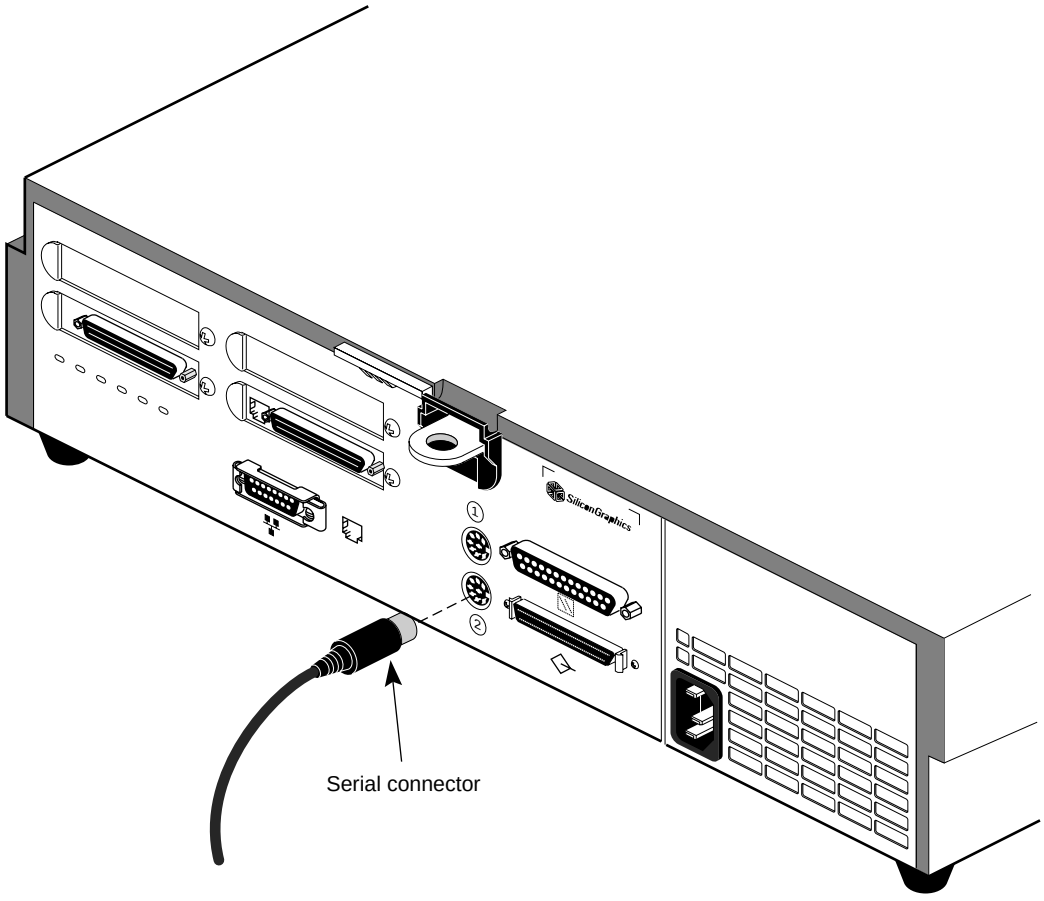

Figure 4-7 Connecting the Serial Cable to the Challenge S Server

3. Plug a serial cable with an RJ11 connector into the Power Control Port 1 on the remote power control unit, as shown in Figure 4-8.

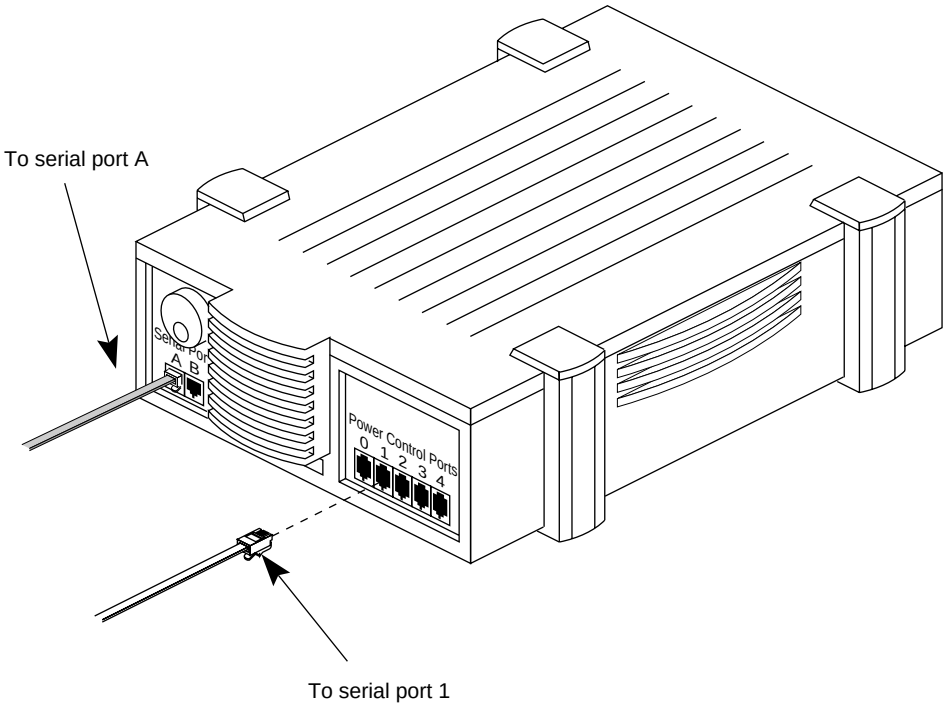

Figure 4-8 Cabling the Power Control Port

4. Attach the RJ11 connector of the serial cable to the RJ11 connector in a power control unit, as shown in Figure 4-9.

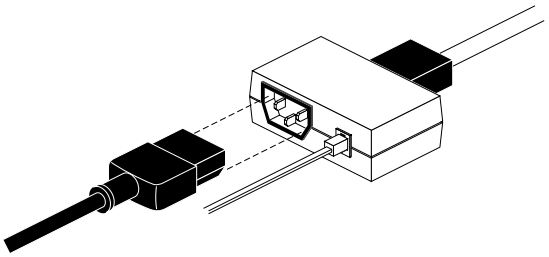

Figure 4-9 Cabling the Power Control Unit

5. Attach the Challenge S power cable to the same server; plug the other end of the power cable into the power control unit.
6. Repeat steps 1 through 5 for the second Challenge S server, using Serial Port B and Power Control Port 0.

   Make sure that serial port A and Power Control Port 1 are connected (through the remote power control unit) to one Challenge S server, and serial port B and Power Control Port 0 are connected to the other.
7. Connect the power for the remote power control unit. Power it on; power on the servers and the ASCII terminal or laptop attached to each.

4.4.2 Cabling a Challenge S and a Larger Server

If you are using a Challenge S and a Challenge DM, L, or XL server, follow steps in this section. Figure 4-10 diagrams the serial interface for this configuration.

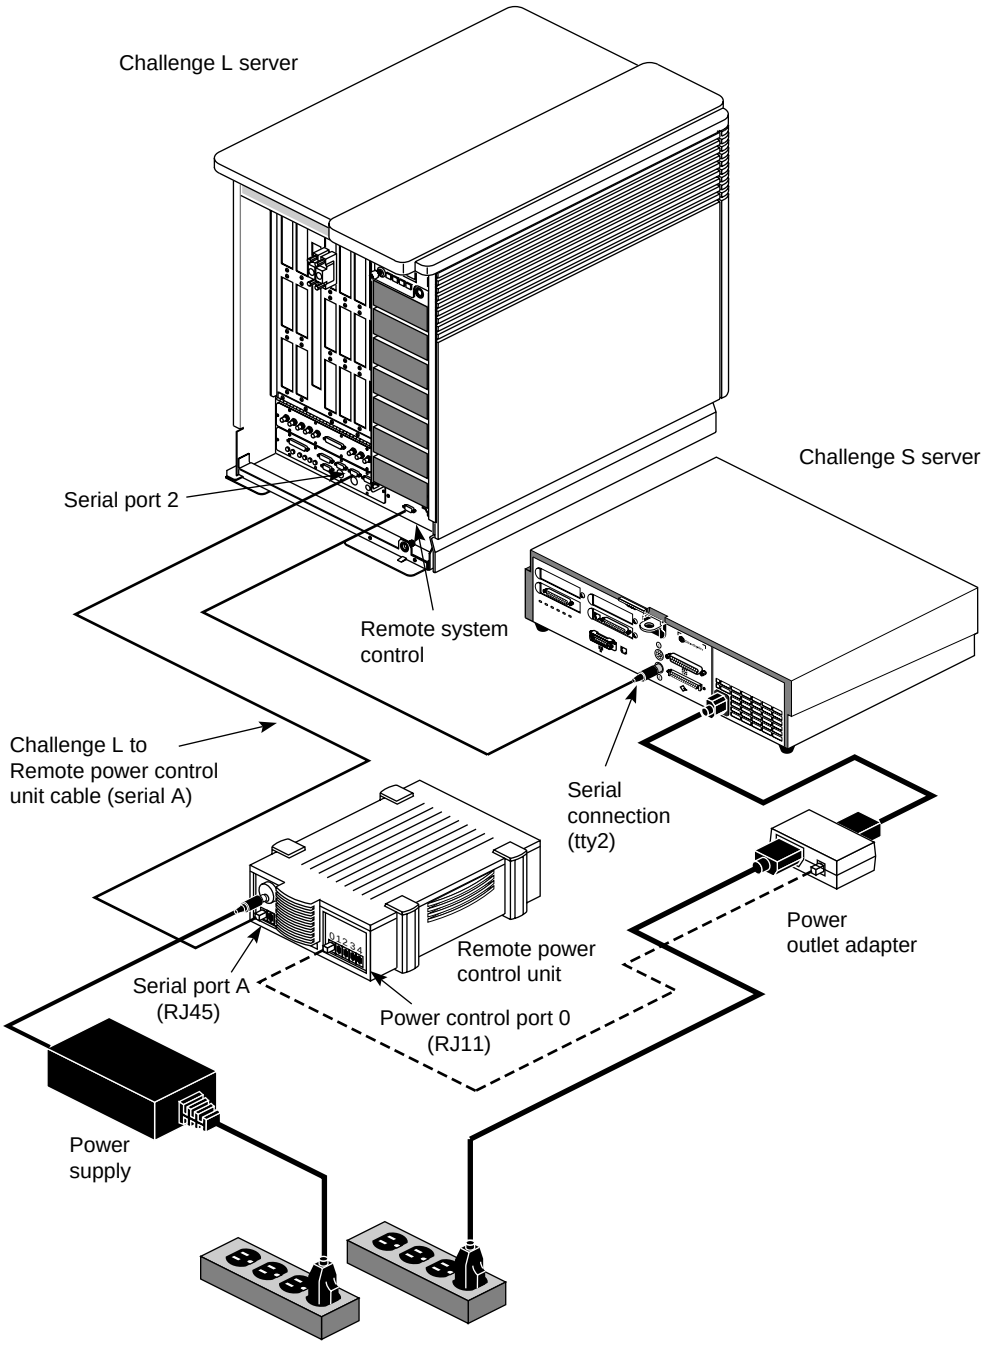

Figure 4-10 Serial Connection: Challenge S Server and Larger Challenge Server

To set up the serial connection for sites using one Challenge S server and another larger Challenge server, follow these steps:

1. Connect the remote power control unit and the larger Challenge server:
   - Attach the serial cable (labeled **CHALLENGE XL/L TO RPCU CABLE**) included with the remote power control unit to RJ45 Serial Port A on the remote power control unit, as shown in Figure 4-6.
   - Attach the other end of this cable to a serial port on the larger Challenge server, shown as serial port 2 in Figure 4-10.
2. Attach one end of a serial cable to Power Control Port 0 on the remote power control unit and the other end of the cable to the RJ45 connector of the power control unit. For details, see Section 4.4.1, "Cabling Two Challenge S Servers."
3. Attach the Challenge S power cable to the server; plug the other end of the power cable into the power control unit.
4. Connect a serial cable to the serial port on the Challenge S server. Attach the other end to the 9-pin Remote System Control port on the larger Challenge server.

4.5 Testing the Installed IRIS FailSafe Hardware

To test the IRIS FailSafe hardware installation, attach an ASCII terminal to each server and run *hinv*. In the output for the Challenge S server, check for

- network interfaces, for example:

  ```
  E++ controller: ec2, version 1
  Integral Ethernet: ec3, version 1
  Integral Ethernet: ec0, version 1
  ```

  This example shows the three network interfaces of a dual-active configuration; an active/standby configuration would lack the first line.

- shared disks, for example:

  ```
  Disk drive: unit 5 on SCSI controller 5
  Integral SCSI controller 4: Version WD33C95A, differential, revision 0
  Integral SCSI controller 0: Version WD33C93B, revision D
  Disk drive: unit 1 on SCSI controller 0
  ```

Type **hinv** on each IRIS FailSafe server; the screen should display output similar to the following:

```
xfs-ha6 14# hinv
1 150 MHZ IP22 Processor
FPU: MIPS R4010 Floating Point Chip Revision: 0.0
CPU: MIPS R4400 Processor Chip Revision: 5.0
On-board serial ports: 2
On-board bi-directional parallel port
Data cache size: 16 Kbytes
Instruction cache size: 16 Kbytes
Secondary unified instruction/data cache size: 1 Mbyte
Main memory size: 64 Mbytes
Integral ISDN: Basic Rate Interface unit 0, revision 1.0
```

```
E++ controller: ec2, version 1
Integral Ethernet: ec3, version 1
Integral Ethernet: ec0, version 1
Integral SCSI controller 5: Version WD33C95A, differential, revision 0
Disk drive: unit 5 on SCSI controller 5
Integral SCSI controller 4: Version WD33C95A, differential, revision 0
Integral SCSI controller 0: Version WD33C93B, revision D
Disk drive: unit 1 on SCSI controller 0
```

4.6 Testing the Serial Connection

Make sure the Remote Power Control unit is powered on.

Test the serial connection as explained in Section 2.6, “Testing the Serial Connection,” in Chapter 2.

4.7 Cabling Storage Systems

Figure 4-11 shows the Challenge S differential SCSI port connection.

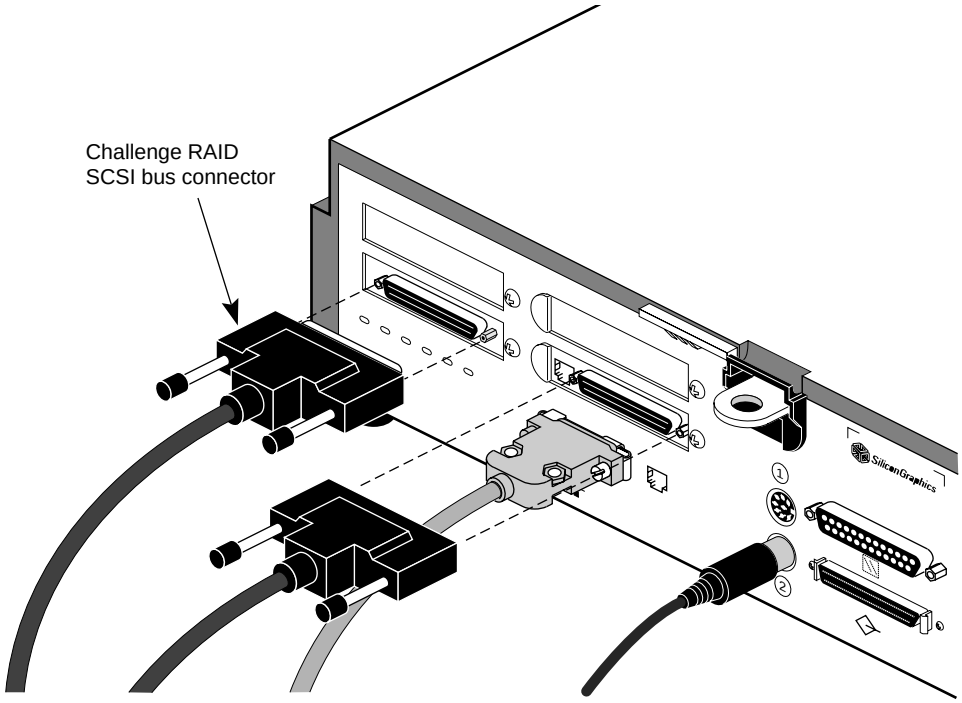

Figure 4-11 Challenge S Differential SCSI Port Connection

To cable the Challenge Vault or Challenge RAID storage system, follow instructions in Section 3.6, “Cabling the Vaults,” and Section 3.7, “Cabling the Challenge RAID Storage System to the Challenge Servers,” in Chapter 3.

4.8 Setting the IRIS FailSafe Host SCSI IDs

Follow instructions in Section 2.11, “Setting the IRIS FailSafe Host SCSI IDs,” in Chapter 2, and Section 3.8, “Setting the IRIS FailSafe Host SCSI IDs,” in Chapter 3.

4.9 Configuring and Testing the System

Follow instructions in the *IRIS FailSafe Administrator’s Guide* to configure and test the newly installed IRIS FailSafe system. If necessary, also consult the latest edition of *IRIX Admin: Disks and Filesystems*.

Caution: You must stop the RAID agent before creating XLV volumes.

To configure Challenge RAID for the IRIS FailSafe system, see Chapter 6, “Configuring Challenge RAID Storage Systems for IRIS FailSafe.”

Chapter 5

Setting Up and Cabling Mixed Configurations (Challenge and Origin Servers)

This chapter explains how to set up and cable an IRIS FailSafe system in which one server (IRIS FailSafe host) is an Origin2000 or Origin200 deskside or rackmount server and the other is a Challenge S, DM, L, or XL server.*

Note: For the latest information on compatibility, see

http://origin/product/hiavailability/support.html

Some configurations require specific patches, firmware versions, versions of storage options, or versions of interface boards.

This chapter consists of these sections:

- Section 5.1, “Installing the Software”
- Section 5.2, “Setting Up the Component Systems”
- Section 5.3, “Installing Interface Boards”
- Section 5.4, “Cabling the Private and Public Networks”
- Section 5.5, “Setting Up the Serial Connection”
- Section 5.6, “Testing the Serial Connection”
- Section 5.7, “Cabling the Storage Systems to the Servers”
- Section 5.8, “Setting the IRIS FailSafe Host SCSI IDs”
- Section 5.9, “Configuring Challenge RAID for IRIS FailSafe”
- Section 5.10, “Configuring and Testing the System”

Note: Before installing an IRIS FailSafe system, make sure that the installation site meets the operating limits and AC power requirements as explained in *Site Preparation for Origin Family and Onyx2* and, if applicable, the *CHALLENGE/Onyx Site Preparation Guide*, chapters 3 and 4.

---

* Onyx and Onyx2 systems are also supported as FailSafe nodes in the same combinations as Challenge L and Origin2000 servers, respectively.

The following equipment is required for installation:

- installation guides for the component systems (see "About This Guide" for part numbers)
  - *Site Preparation for Origin Family, Onyx2, OCTANE, and O2*
  - *CHALLENGE/Onyx Site Preparation Guide*
  - manuals for storage options (vaults, RAID) and interface boards (FDDI, Ethernet, SCSI); see "About This Guide" for names and part numbers
- laptop, ASCII terminal, or IRISconsole
- Phillips and small flat-blade screwdrivers

Note: Make sure that each Challenge L server used for the IRIS FailSafe system has a Remote System Control port at the extreme right. If it does not, the customer must order an upgrade (013-0624-003) and you must install this port following instructions in the *Remote System Control Port Installation Guide*, included in this shipment.

5.1 Installing the Software

Follow instructions in the latest version of the *IRIS FailSafe Administrator's Guide* (007-3109-003) to install the software needed to run IRIS FailSafe.

Depending on the servers and storage in the configuration and the IRIX revision level, you might need to install various patches.

- For the latest information on server and storage compatibility, see

  http://origin.engr.sgi.com/product/hiavailability/support.html
- For information on recommended patches for each platform, see

  http://bits.csd.sgi.com/digest/patches/recommended/

For IRIS FailSafe operation, you must determine the MMSC and MSC passwords on Origin2000 and Origin200 systems used as IRIS FailSafe hosts, if the system administrator has changed them from the default. This password is used when one node reset the other. To specify this password so that the IRIS FailSafe software knows about it, use

ha_spng -w *password* **-d** *dst_sysctlr_type*

For more information on MSC passwords, see Appendix D of the *Origin2000 and Onyx2 Deskside and Rackmount Installation Instructions* or the latest version of the *IRIS FailSafe Administrator's Guide*. For information on the *ha_spng* command, see its man page, ha_spng(1M).

5.2 Setting Up the Component Systems

This section consists of the following:

- Section 5.2.1, “Setting Up the Hardware”
- Section 5.2.2, “Planning the Connections Between IRIS FailSafe Hosts”

5.2.1 Setting Up the Hardware

Read through *Site Preparation for Origin Family, Onyx2, OCTANE, and O2* and the *CHALLENGE/Onyx Site Preparation Guide* before unpacking equipment at the site. Follow instructions in those manuals, specifically:

1. Make sure that the installation site meets the operating limits and AC power requirements for the hardware.
2. Make sure required tools and personnel are on hand for unloading and opening crates and for moving large systems.
3. Prepare the physical location to allow for space, air flow, and floor-loading requirements for all component systems

   Use only the cables included in the shipment. Plan to situate the servers and vaults fairly close together. Differential Fast-20 SCSI (installations with Origin family only), including cabling inside the chassis, should be no longer than 25 meters (82 feet). Differential SCSI-2 cables, including cabling inside the chassis, should be no longer than 18.3 meters (60 feet).

   Note: Although the SCSI-2 bus absolute length limit is 80 feet, exceeding 60 feet is not recommended for Challenge systems. When cable lengths exceed 60 feet, problems can occur on the SCSI mezzanine card, the IO4B board, or both. For the IO6 board, the SCSI-2 bus absolute length limit is 80 feet.
4. Make sure the site meets safety and operating considerations.
5. Prepare the site for the systems’ electrical requirements.
6. Connect your laptop or ASCII terminal and keyboard to the first host, as explained in the server installation instructions.
7. Connect the power cord for the first IRIS FailSafe host. Turn on the main power switch on the back of the unit; turn on the laptop or ASCII terminal. If necessary, see the server installation instructions for details.
8. Turn on the IRIS FailSafe host, following instructions in its manual. When power-on diagnostics are completed, a login prompt appears on the console.
9. Enter **`hinv`** on each IRIS FailSafe host. In the output, look for the installed shared disks and network interfaces.

5.2.2 Planning the Connections Between IRIS FailSafe Hosts

For the IRIS FailSafe system, you are setting up these networks on each server:

- public network interface(s)

  This network connects the IRIS FailSafe server to clients and the outside world. The I/O panel Ethernet port or one from an option board can be used.

  Consult with the customer on the number and type of public network interface(s) for each server. If a server has two public network interfaces, they must be on different networks.

- serial connection between nodes: the MSC (Origin200 deskside system controller), **AUX** (Origin200 deskside tty port), MMSC **ALTERNATE CONSOLE** port (Origin200 rackmount module), or Remote System Control (Challenge) port is cabled to a tty port on the other node

  If one node fails, this connection enables the surviving node to power cycle the other.

- private (heartbeat) network (Ethernet)

  This network, going only between the two nodes, is used exclusively for the keep-alive heartbeat that the two nodes use for monitoring each other's status and IRIS FailSafe control messages. The I/O panel Ethernet port or one from an option board can be used.

- shared fibre channel or differential SCSI connection to storage

Figure 5-1 diagrams an example IRIS FailSafe system with an Origin2000 deskside server and a Challenge L server.

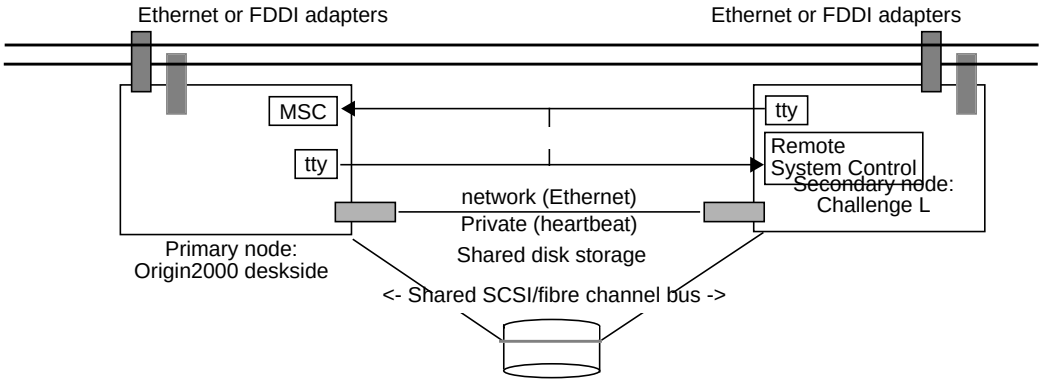

Figure 5-1 Origin2000 Deskside and Challenge L Servers

Figure 5-2 diagrams an example IRIS FailSafe system with an Origin2000 rackmount server and a Challenge L server.

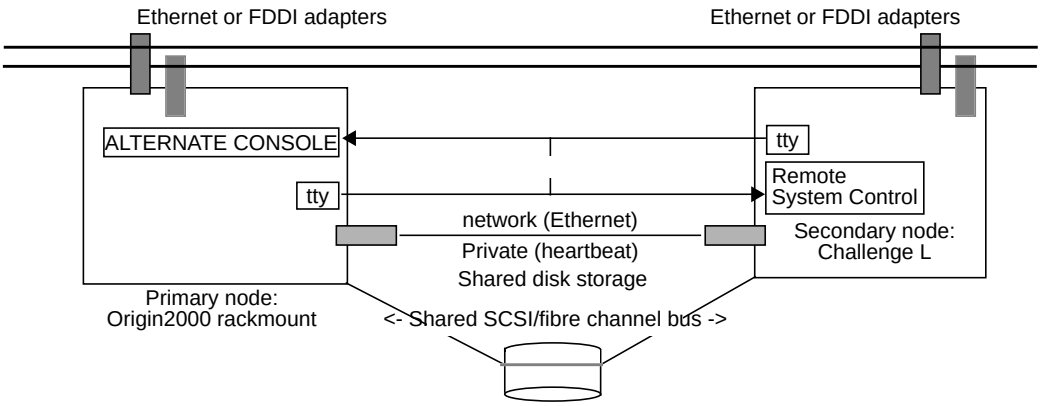

Figure 5-2 Origin2000 Rackmount and Challenge L Servers

Figure 5-3 diagrams an example IRIS FailSafe system with an Origin200 server and a Challenge L server.

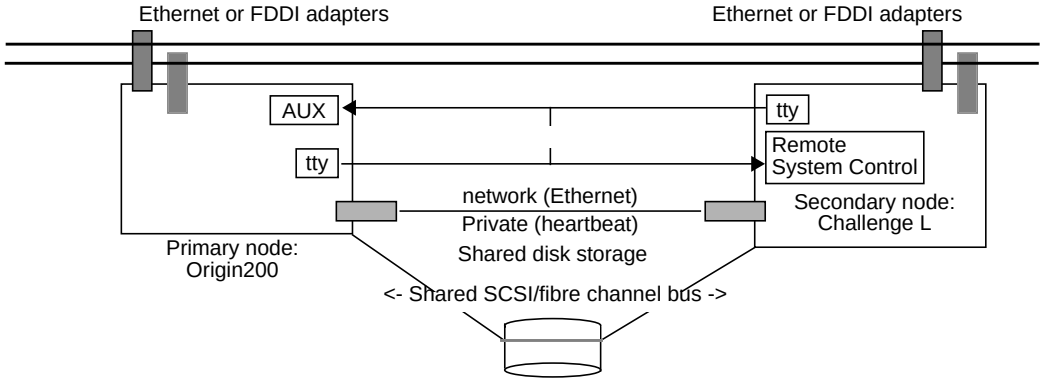

Figure 5-3 Origin200 and Challenge L Servers

Figure 5-4 diagrams an example IRIS FailSafe system with an Origin200 server and a Challenge S server.

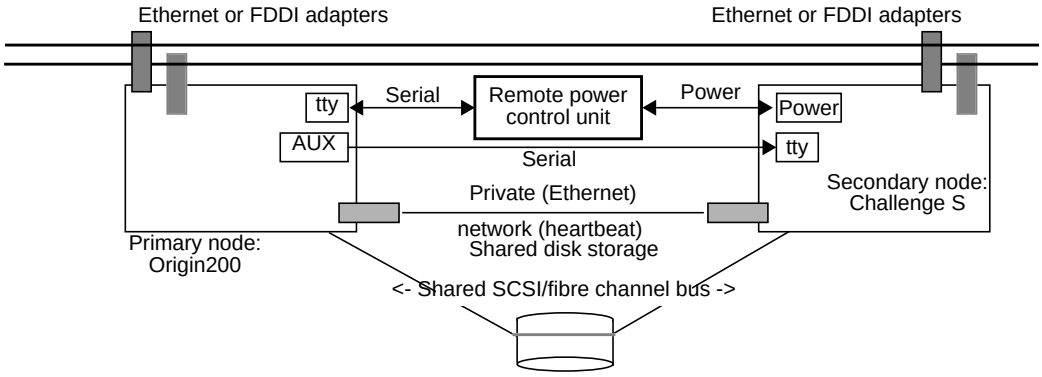

Figure 5-4 Origin200 and Challenge S Servers

Note: For currently supported combinations of servers and requirements for the various combinations, see

http://origin.engr.sgi.com/product/hiavailability/support.html

5.3 Installing Interface Boards

Before you cable the networks, install any required interface boards:

- Ethernet:

  The Origin2000 IO6 panel has one 10-Base-T/100-Base-T Ethernet port. The optional ENET XIO board holds four additional Fast Ethernet ports and six additional asynchronous serial ports. (For information on installing the ENET board, see *IRIS 4-Port Fast Ethernet with Serial XIO Board Installation Instructions*.)

  The Origin200 has one 10-Base-T/100-Base-T Ethernet port. The optional ENET PCI board holds an additional Fast Ethernet port. For information on this board, see its documentation (*Fast Ethernet PCI Option Installation Instructions*).

  For systems with a large Challenge server, install a 100-Base-T board in the Challenge server (see *100Base-T VME Board Installation Instructions*), or use the Ethernet (10-Base-T) port on the IO4 panel.

  For information on installing an E-Plex board in a Challenge server, see *POWER Channel E-Plex Installation Instructions*.

- FDDI: If the customer is using FDDI for the private network, see
  - Origin200 or Origin2000 server FDDI PCI board: see *Origin200 and Origin Vault Installation Instructions* or Origin2000 and Onyx2 installation manuals for instructions on installing PCI boards

    Note: As of this writing, no FDDI XIO board is available from Silicon Graphics, although one is planned for calendar 1998.
  - Challenge L or XL server: Section 3.3.1, “Installing FDDI Boards in the IRIS FailSafe Hosts for a Public Network Connection,” in Chapter 3 of this manual, and FDDI manuals for Challenge servers
- SCSI: If the customer has ordered Ultra SCSI XIO board(s) or PCI SCSI board(s), see
  - Origin2000 server XIO board: *Ultra SCSI XIO Board Installation Instructions*
  - Origin200 or Origin2000 server PCI board: *Origin200 and Origin Vault Installation Instructions*
- Fibre Channel: If the customer has ordered a Fibre Channel XIO or PCI board, see
  - XIO board: *Origin FibreVault and Fibre Channel RAID Installation Instructions*
  - Origin200 or Origin2000 server PCI board: *Origin200 and Origin Vault Installation Instructions*

5.4 Cabling the Private and Public Networks

This section consists of the following:

- Section 5.4.1, “Setting Up a Private Network Connection”
- Section 5.4.2, “Setting Up a Public Network Ethernet Connection”

5.4.1 Setting Up a Private Network Connection

The private Ethernet network between the servers supplies the heartbeat of each server to the other. This section explains cabling the private network for various combinations of servers.

Figure 5-5 shows the Ethernet connector on the IO6 panel. The server might also have one or more Ethernet boards.

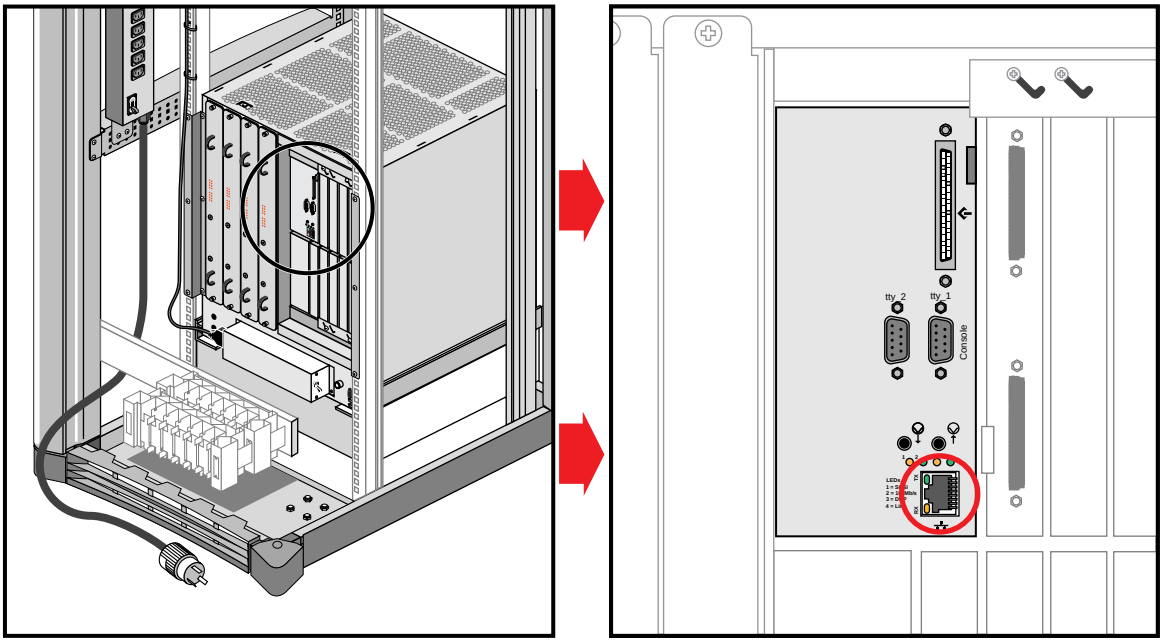

Figure 5-5 Origin2000 Rackmount Server Ethernet Connector

The Origin200 rear panel has one 10-Base-T / 100-Base-T Ethernet port. The optional ENET PCI board holds one additional Fast Ethernet port. Figure 5-6 shows these ports.

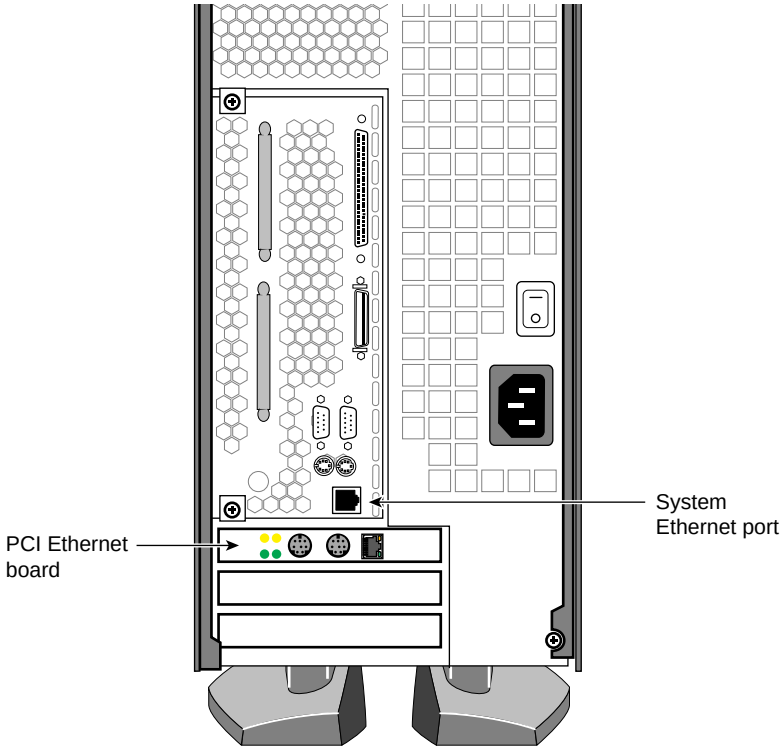

Figure 5-6 Origin200 Deskside Server Ethernet Ports

Figure 5-7 shows the Ethernet connector for a Challenge XL rackmount system.

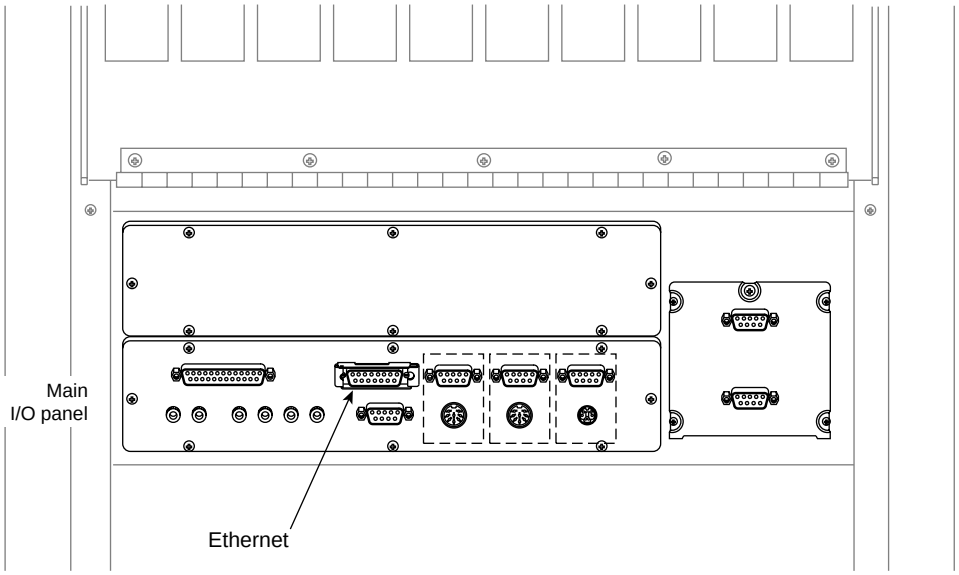

Figure 5-7 Challenge XL Rackmount Server Ethernet Port

Figure 5-8 shows Ethernet connectors for a Challenge S system.

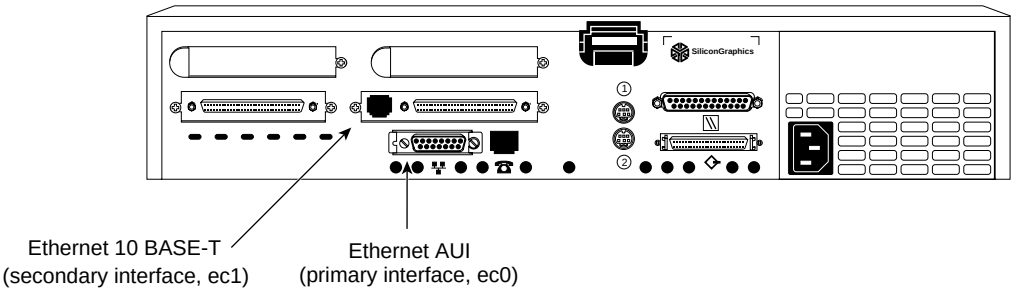

Figure 5-8 Challenge S Server Ports

To cable the private network, attach an end of the null modem Ethernet cable supplied with the IRIS FailSafe system (p / n 9290131: 20-foot, or p / n 9290132: 40-foot) to an Ethernet port on each host module.

Note: If the configuration includes a Challenge server with no 10-Base-T port, attach a transceiver.

Figure 5-9 shows an example of private network cabling for Origin2000 and Challenge L servers.

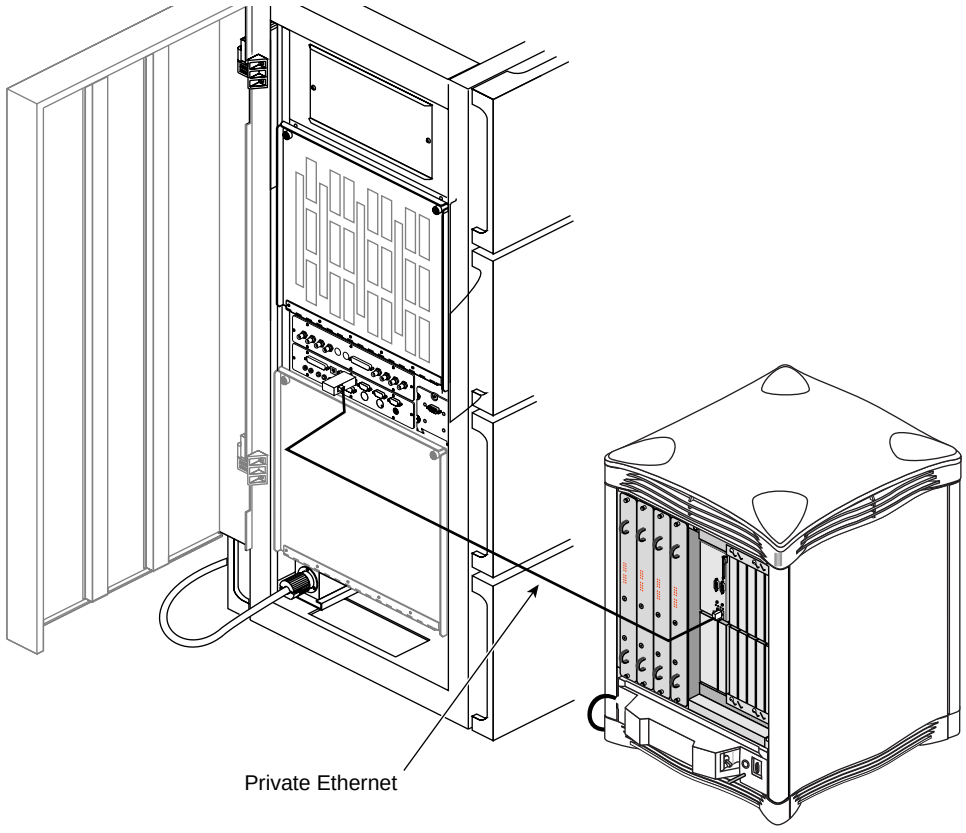

Figure 5-9 Origin2000 and Challenge L Private Network Ethernet Cabling

Figure 5-10 shows an example or private network cabling for Origin200 and Challenge S servers.

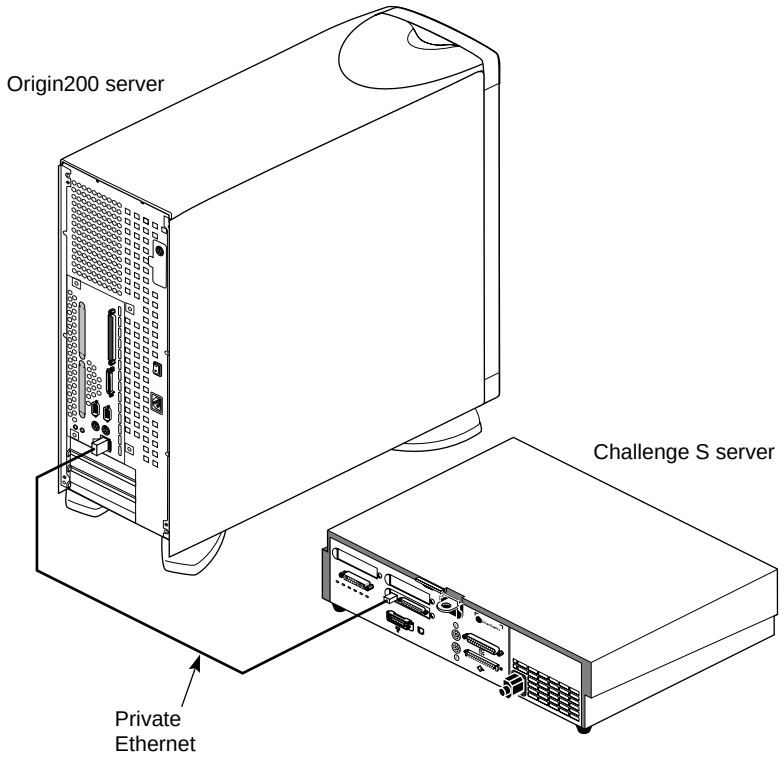

Figure 5-10 Origin200 and Challenge S Private Network Ethernet Cabling

5.4.2 Setting Up a Public Network Ethernet Connection

On each server, connect the public network drop cable to a 10-Base-T or 100-Base-T Ethernet port. For a larger Challenge server, which does not have any 10-Base-T ports, use a transceiver if necessary, or connect to an Ethernet port on the 100-Base-T option board for a Challenge server.

5.5 Setting Up the Serial Connection

Cross-cable the system controller port on one server to a serial (tty) port on the other server's IO board panel for the serial connection between the two so that one can power off the other in case of failure. This section explains this cabling for various combinations of servers, in these sections:

- Section 5.5.1, "Ports for the Serial Connection"
- Section 5.5.2, "Cabling the Serial Connection"
- Section 5.5.3, "Cabling the Serial Connection With a Challenge S Server"

5.5.1 Ports for the Serial Connection

Table 5-1 summarizes the ports to use on each type of server.

Table 5-1 Ports for Serial Connection

Server	Controller (Serial) Port	Serial Port (9-Pin Sub-D)
Origin2000 deskside	MSC serial port on rear (9-pin sub-D)	tty_2 on IO6 board
Origin2000 rackmount	MMSC **ALTERNATE CONSOLE** port (8-pin mini-DIN)	tty_2 on IO6 board
Origin200	**AUX** port on rear (8-pin mini-DIN)	tty_2 on rear
Challenge XL / L / DM	Remote System Control port (9-pin sub-D)	tty_2
Challenge S	tty_2 to Silicon Graphics remote power control unit (9RJ45)	tty_2

Figure 5-11 shows the serial port on the rear of the Origin2000 module.

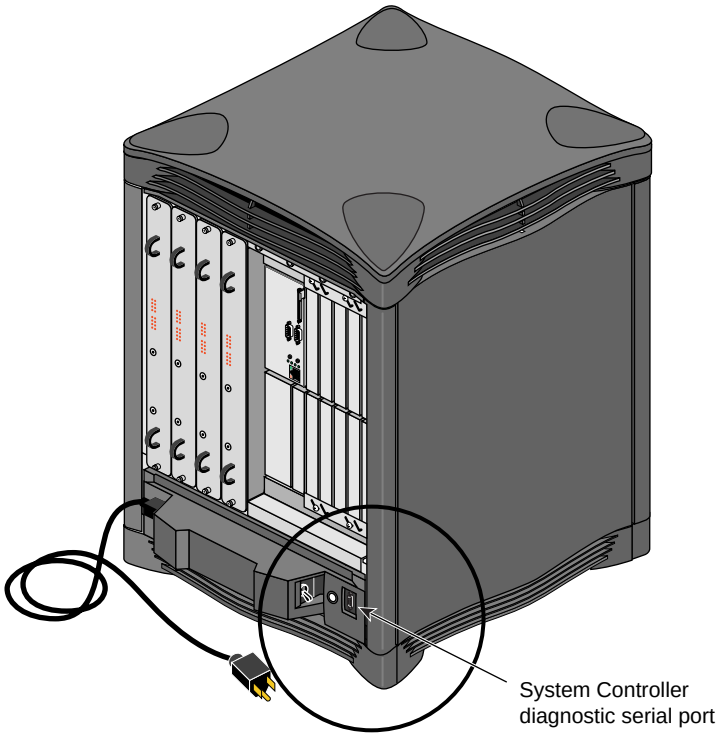

Figure 5-11 Origin2000 Deskside Server Serial Port (Rear)

Caution: The 9-pin sub-D serial port on the rear of the system is the same electrically as the 8-pin mini-DIN serial port on the front of the Origin2000 server. If anything is attached to either port before you begin installing the IRIS FailSafe system, it must be removed.

Figure 5-12 shows the **ALTERNATE CONSOLE** port on the Origin Rack MMSC.

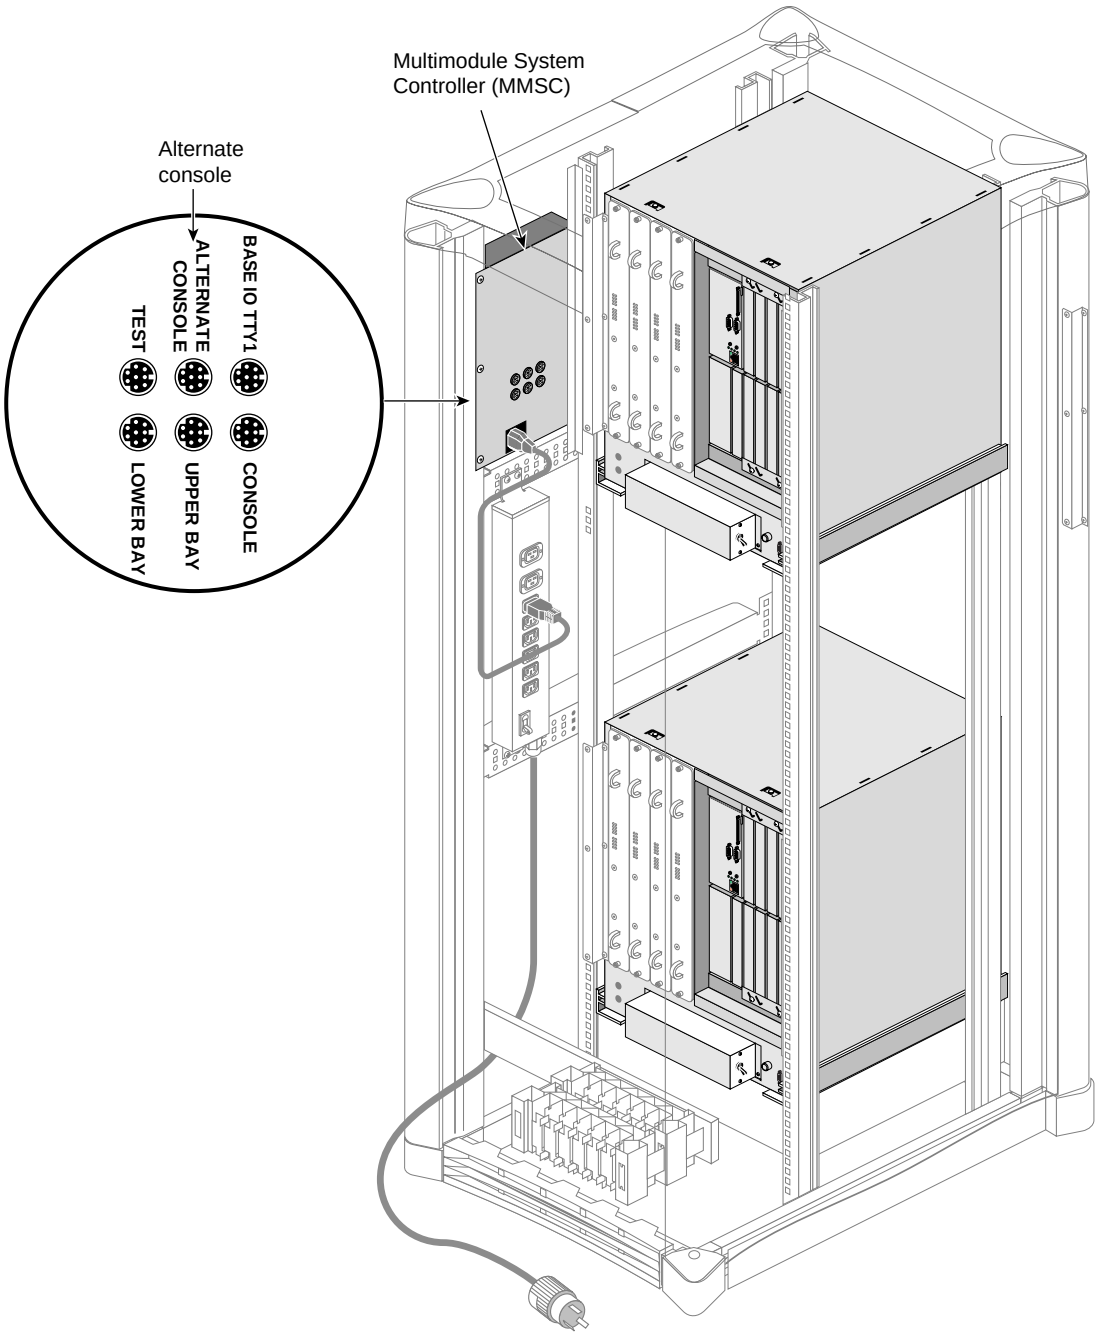

Figure 5-12 Origin2000 Rack MMSC ALTERNATE CONSOLE Port

Figure 5-13 shows the tty ports on the Origin2000 server IO6 panel. The right-hand one is the console port, to which the system console is connected. The IRIS FailSafe serial connection uses the left-hand tty port.

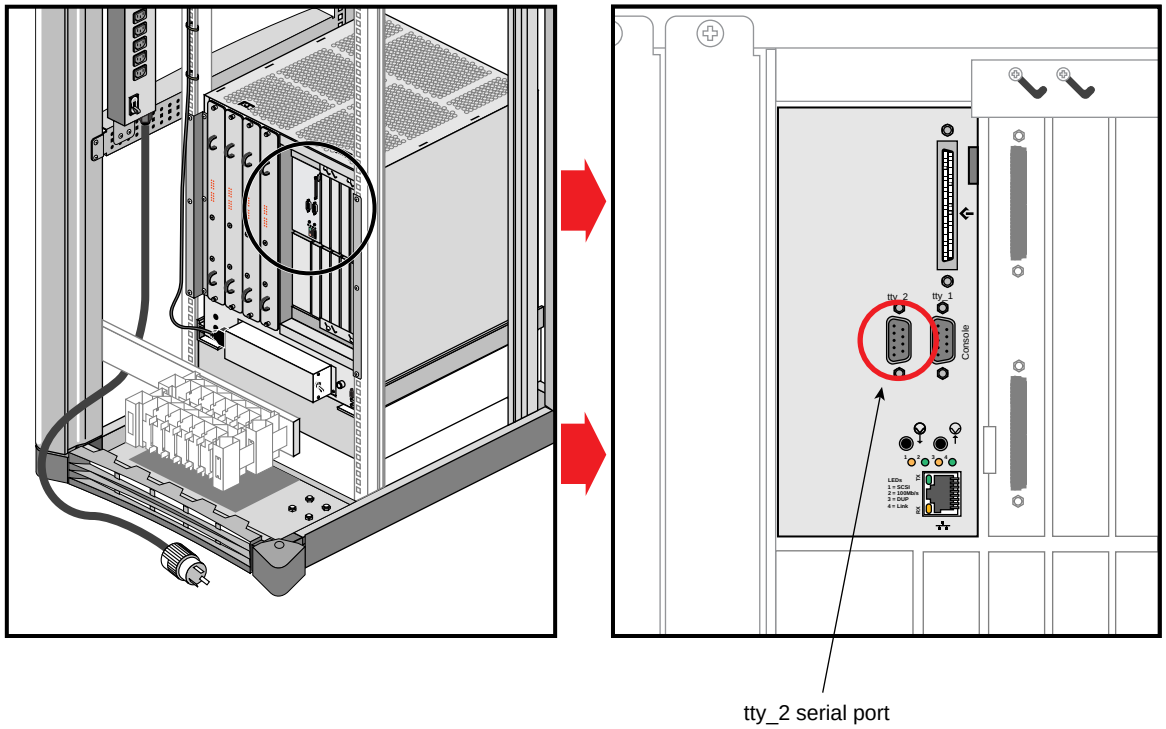

Figure 5-13 Origin2000 Rackmount Server tty Ports

Figure 5-14 shows the serial ports on the Origin200 server.

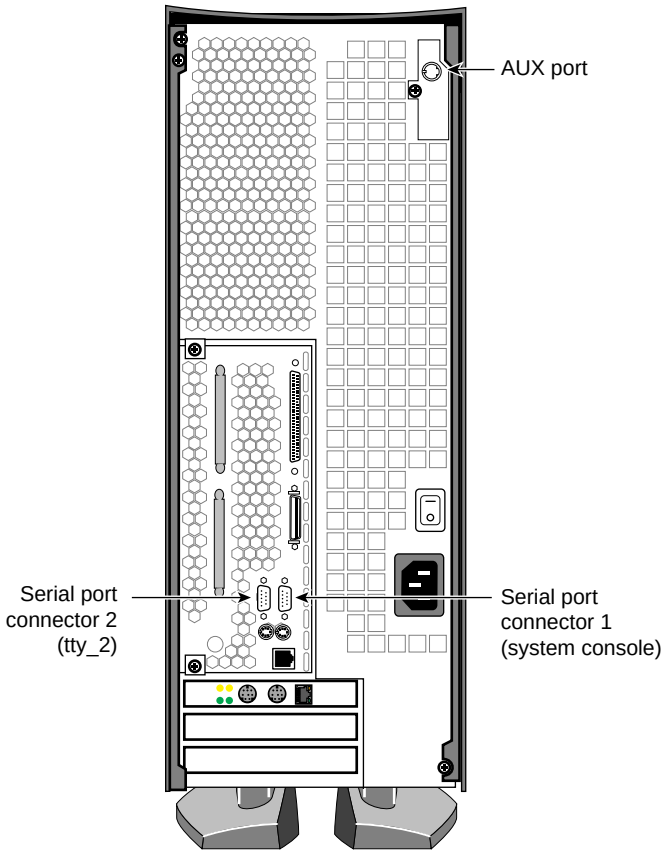

Figure 5-14 Origin200 Deskside Server Serial Ports

Figure 5-15 shows the Challenge L remote system control port and tty2 connector.

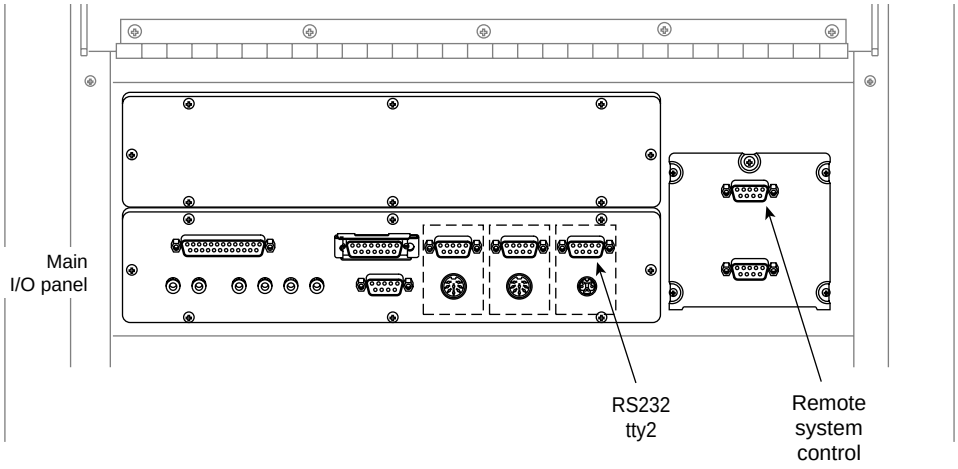

Figure 5-15 Challenge L Server Remote System Controller Port

For serial connectors on a Challenge S system, see Figure 5-8.

5.5.2 Cabling the Serial Connection

Table 5-2 summarizes cables and connections for various combinations of servers.

Table 5-2 Serial Connection, Origin2000 or Origin200 Server and Challenge Server

Servers	Cable	Connections
Origin2000 deskside; Challenge XL/L/DM	018-0669-001	Origin2000: MSC serial port on rear to tty_2 on Challenge I/O panel
	018-0669-001	Challenge: Remote System Control port to tty_2 on Origin2000 IO6 board
Origin2000 rackmount; Challenge XL/L/DM	018-0690-001	Origin2000: MMSC **ALTERNATE CONSOLE** port to tty on Challenge I/O panel
	018-0669-001	Challenge: Remote System Control port to tty_2 on Origin2000 IO6 board
Origin200; Challenge XL/L/DM	018-0690-001	Origin200: **AUX** port to tty_2 on Challenge
	018-0669-001	Challenge: Remote System Control port to tty_2 on Origin200
Origin200; Challenge S	018-8223-001	Origin200: **AUX** port to tty_2 on Challenge S
	018-0668-001	Challenge S: power cable to remote power control unit and serial cable to tty_2 on Origin200

Follow these steps:

1. For an Origin2000 server, make sure that nothing is connected to the serial port on the front or the back.
2. In the Origin2000 Rack, make sure that the cabling to the MMSC is correct.
3. Attach a 10-Base-T transceiver to the Ethernet port on the Challenge XL/L/DM server.
4. Cable the serial connection following the information in Table 5-1 and Table 5-2.

Figure 5-16 shows the serial connection between an Origin2000 deskside server and a Challenge L server.

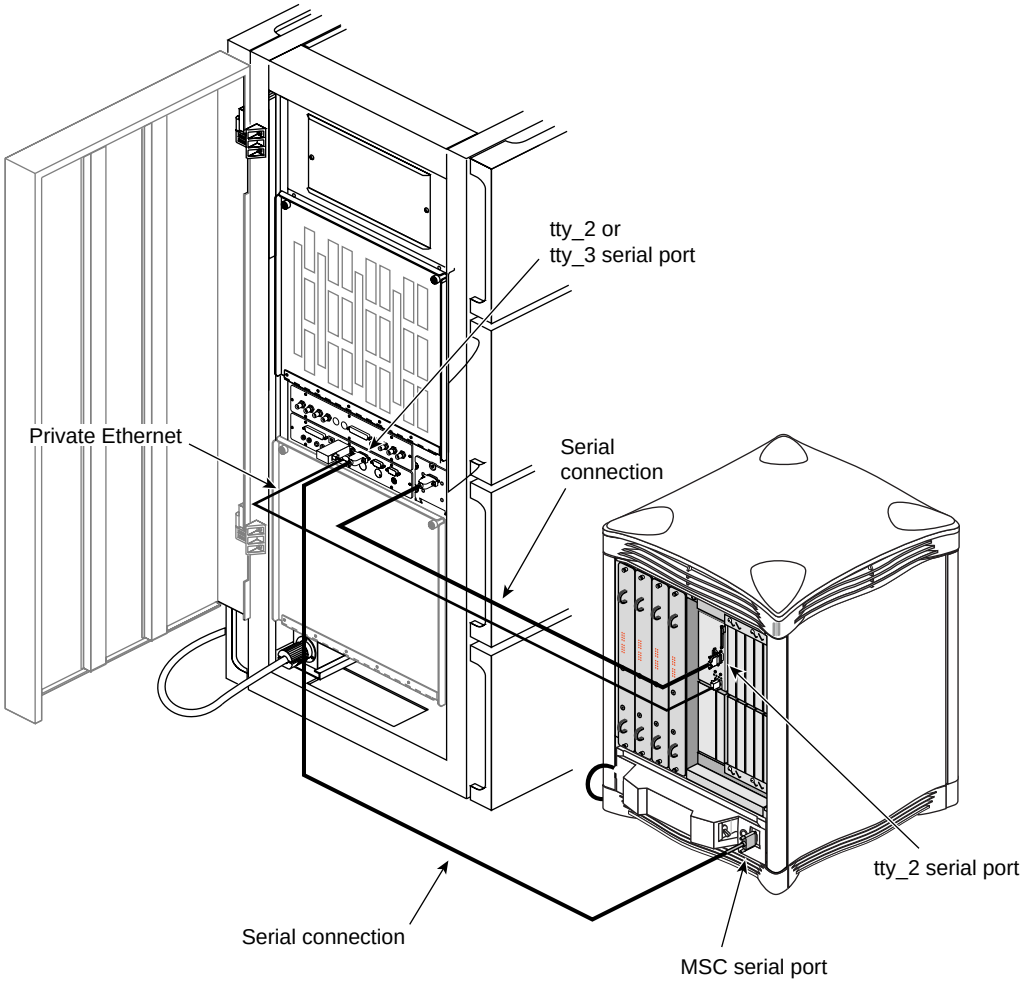

Figure 5-16 Serial Connection and Private Ethernet Connection, Origin2000 Deskside Server and Challenge L Server

Figure 5-17 shows the serial connection between an Origin2000 rackmount server and a Challenge L server.

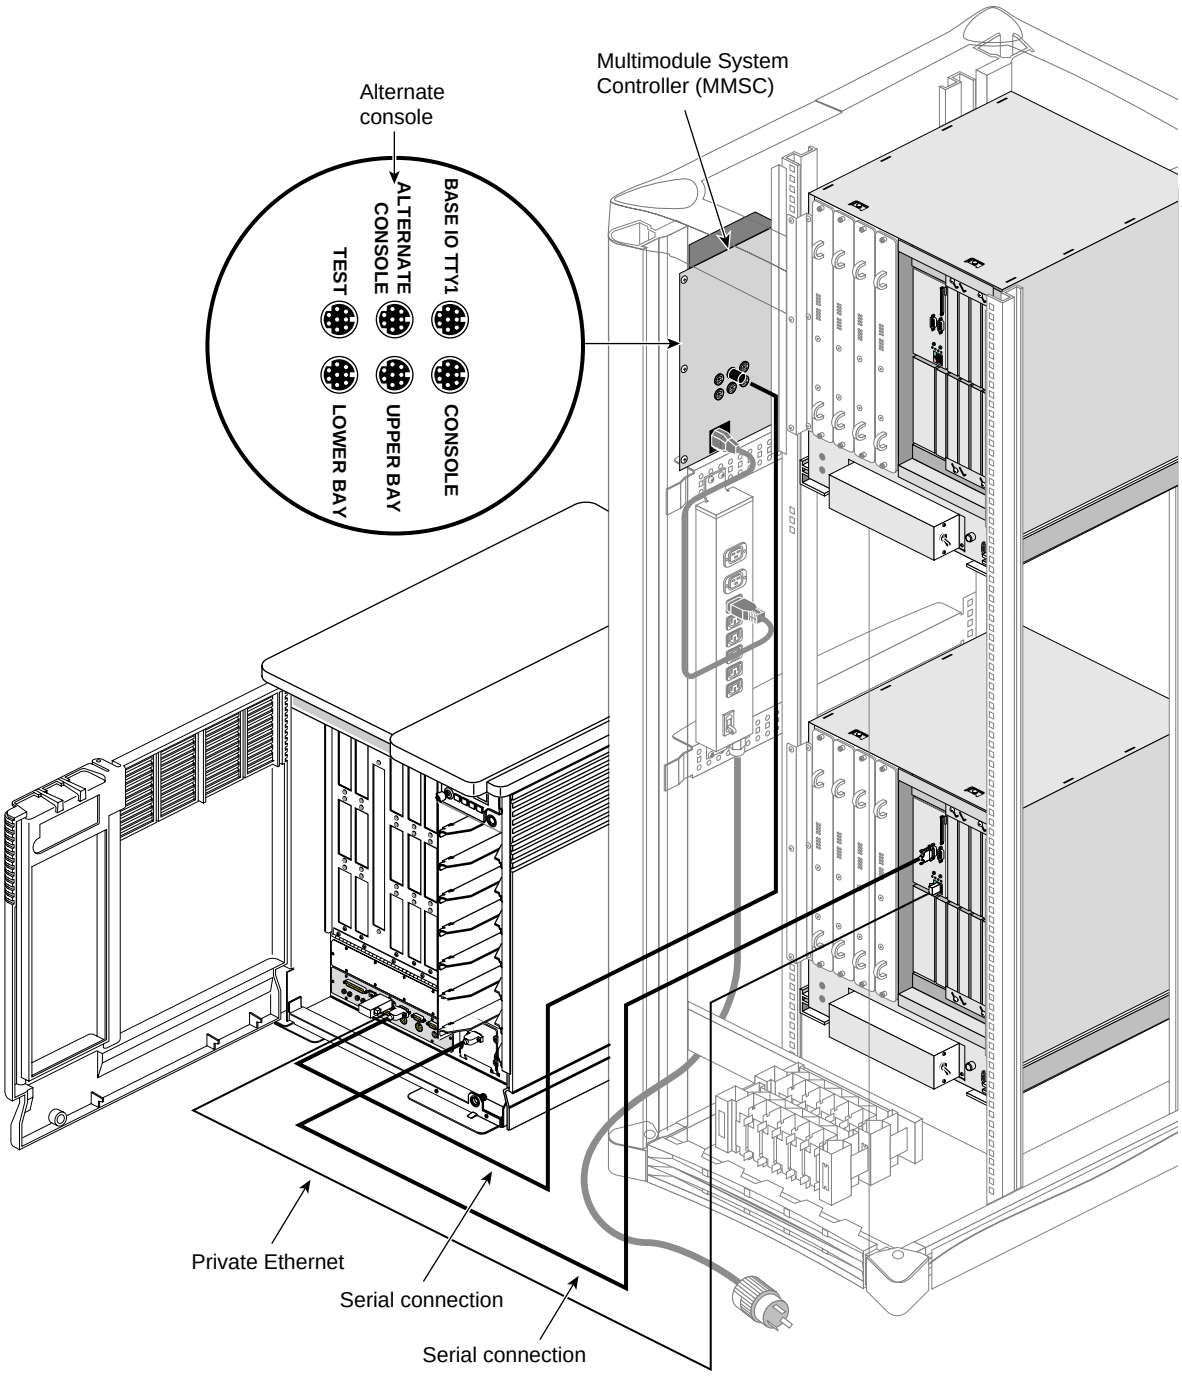

Figure 5-17 Serial Connection and Private Ethernet Connection, Origin2000 Rackmount Server and Challenge L Deskside Server

Figure 5-18 shows the serial connection between Origin200 server and Challenge L deskside servers.

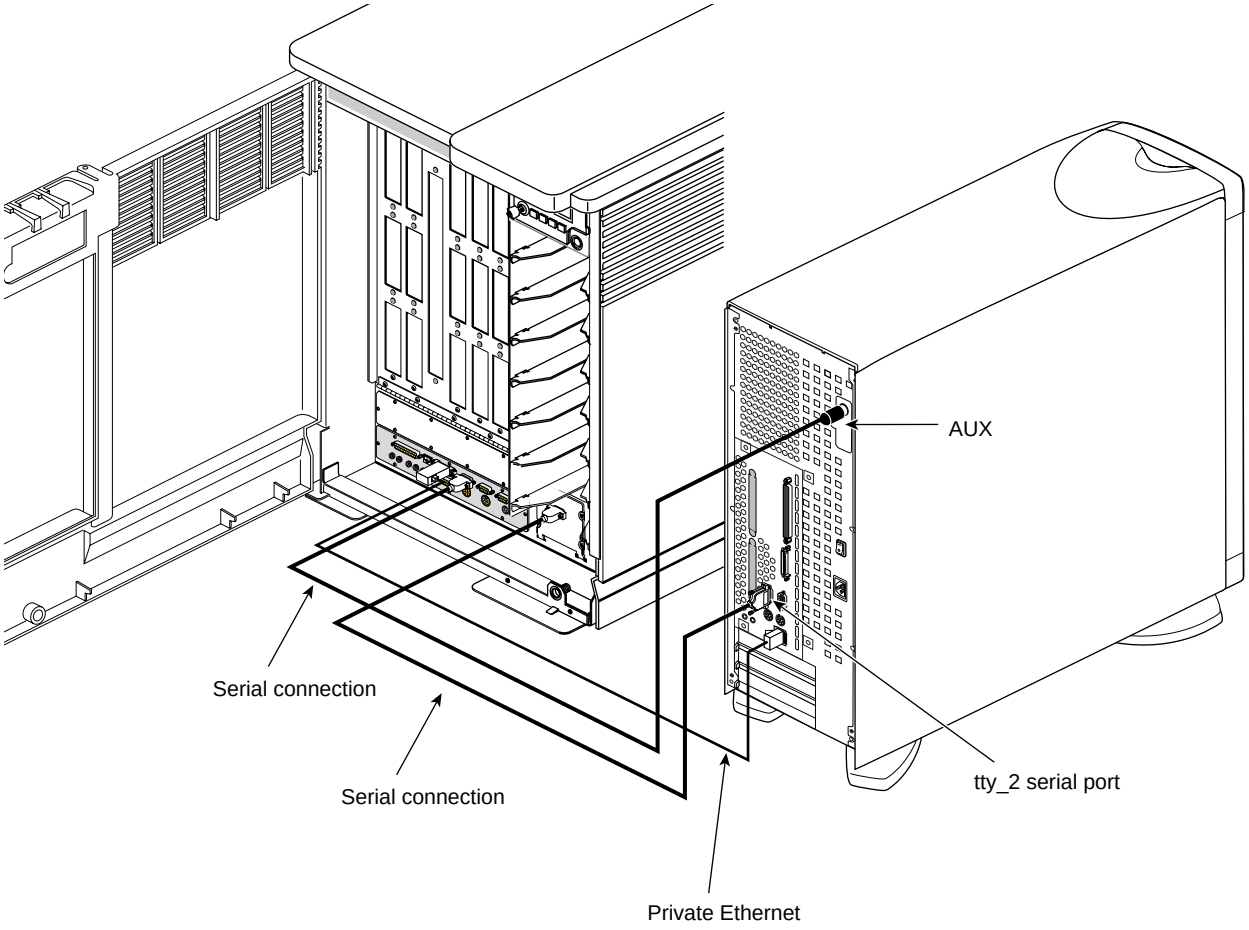

Figure 5-18 Serial Connection and Private Ethernet Connection, Origin200 and Challenge L Deskside Servers

5.5.3 Cabling the Serial Connection With a Challenge S Server

For the Challenge S server, which has no system control unit, the Silicon Graphics remote power control unit is required. Figure 5-19 diagrams this configuration.

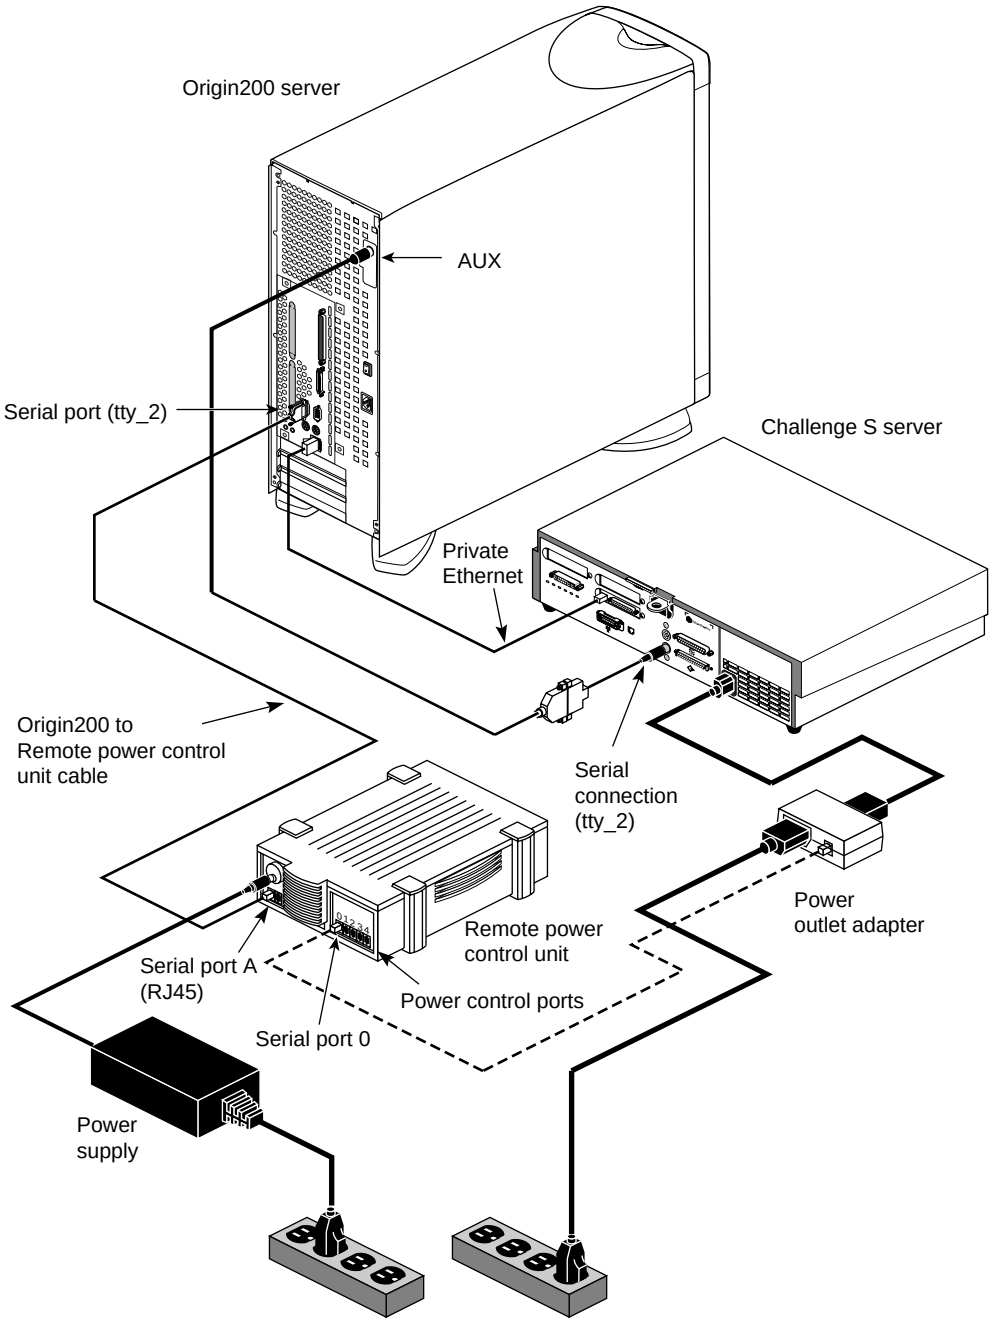

Figure 5-19 Serial Connection, Origin200 Server and Challenge S Server

To set up the serial connection between an Origin200 server and a Challenge S server, follow these steps:

1. Connect the Challenge S remote power control unit and the Origin200 server: attach the serial cable (labeled **CHALLENGE XL/L TO RPCU CABLE**) included with the remote power control unit to RJ45 Serial Port A on the remote power control unit, as shown in Figure 5-20.

   Note: For all RJ45 and RJ11 connections, be sure the connector is properly seated in the jack so that it makes good contact.

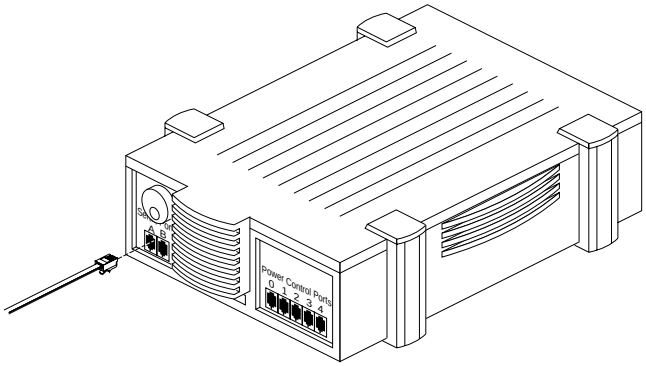

Figure 5-20 Cabling Serial Port A on the Remote Power Control Unit

   Note that power control ports 2 through 4 are not used; they are supplied with terminator plugs to reduce the chance for error.

2. Attach the other end of this cable to tty_2 on rear of the Origin200 server; see Figure 5-14.

3. Attach one end of a serial cable to power control port 0 on the remote power control unit, as shown in Figure 5-21.

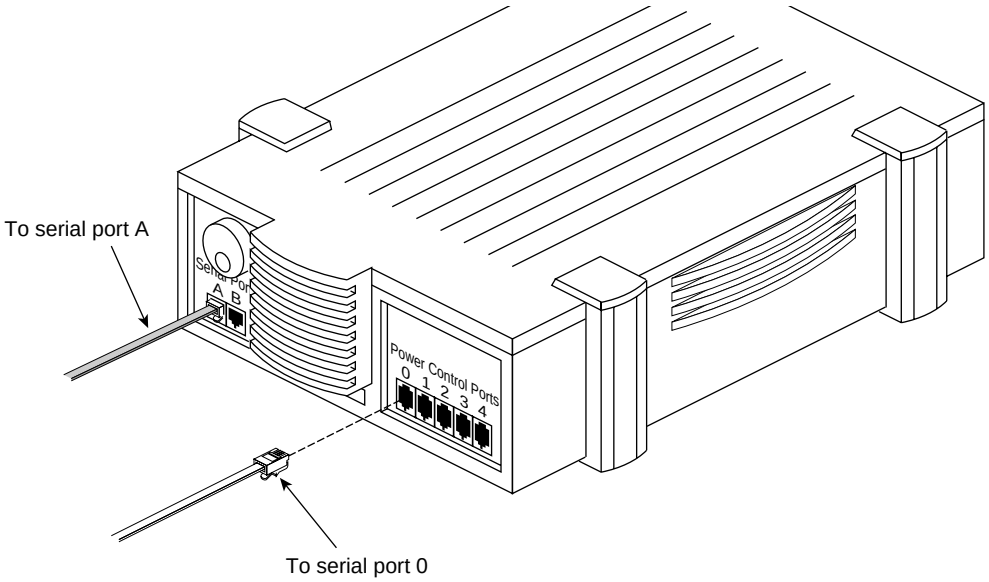

Figure 5-21 Cabling the Power Control Port

4. Attach the RJ11 connector of the serial cable to the RJ11 connector in a power control unit, as shown in Figure 5-22.

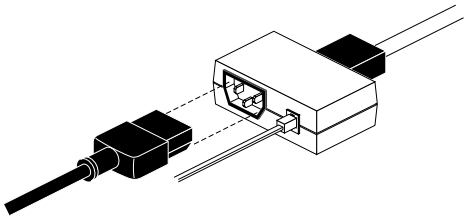

Figure 5-22 Cabling the Power Control Unit

5. Attach the Challenge S power cable to the server; plug the other end of the power cable into the power control unit.
6. Connect a serial cable to serial port tty_2 on the Challenge S server. Attach the other end to the **AUX** port on the Origin200 server.

Figure 5-23 shows the serial connection between an Origin200 server and a Challenge S server.

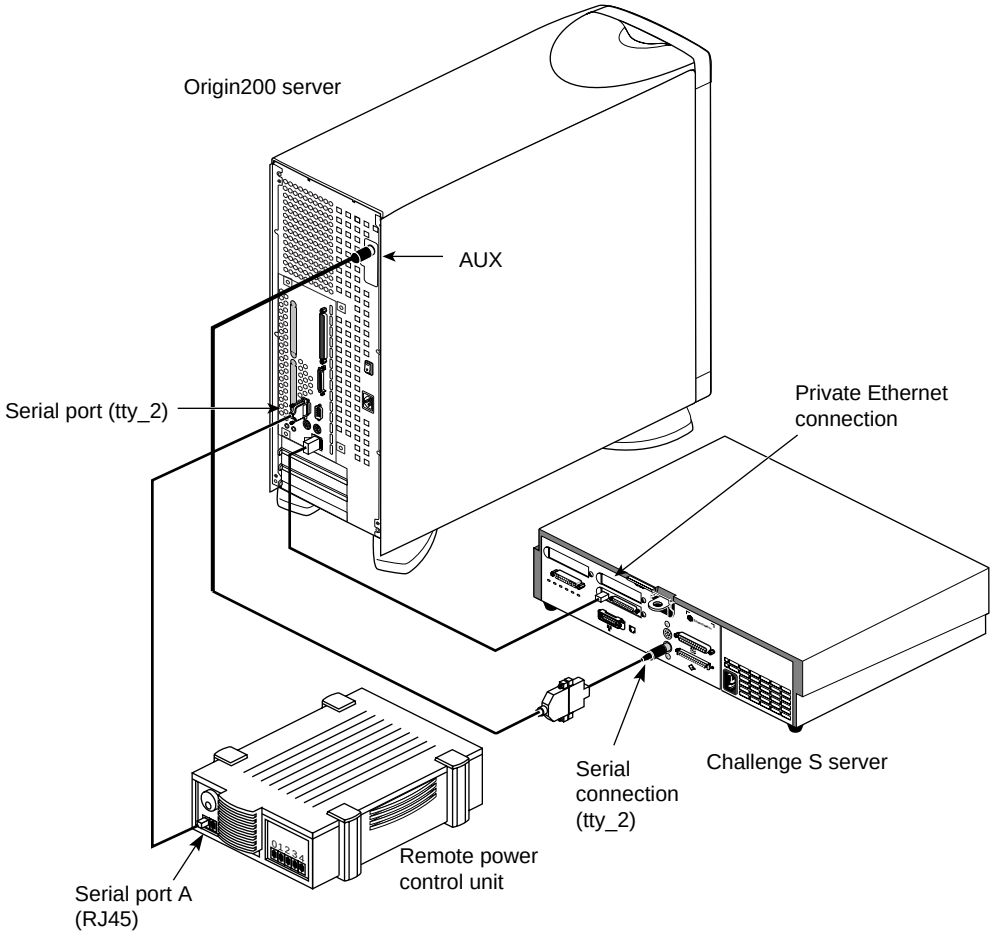

Figure 5-23 Serial Connection, Origin200 Deskside Server and Challenge S Server

5.6 Testing the Serial Connection

To test the serial connection between the IRIS FailSafe servers, follow these steps:

1. Make sure the IRIS FailSafe servers are powered on.
2. To stop IRIS FailSafe on both nodes, enter

   /etc/init.d/failsafe stop

3. Test the connection. For example, if the serial cable was connected to tty_2 and the other host is a Challenge server, enter

   /usr/etc/ha_spng -i 10 -f /dev/ttyd2

   If the other host is an Origin200 server, enter

   /usr/etc/ha_spng -i 10 -f /dev/ttyd2 -d MSC

   If the other host is a deskside or rackmount Origin2000 server, enter

   /usr/etc/ha_spng -i 10 -f /dev/ttyd2 -d MMSC

   In this command, **/dev/ttyd2** is the tty of the node on which you are entering this command and **MMSC** or **MSC** is the system controller of the other node.

   No output appears; check the return value of the command. If the return value is 0, the connection is good.

   If the return value is 1, perform these checks:

   - Verify that the IRIS FailSafe server is powered on.
   - Verify the cable connections from one server's serial port or remote power control unit and the other server's system controller port.

4. Repeat step 3 on the second node.

Note: If the system administrator has changed the MMSC or MSC password from the default, the IRIS FailSafe software must be notified; see Section 5.1, "Installing the Software," or the ha_spng(1M) man page.

If *ha_spng* fails, make sure all the cables are seated properly and rerun the command with a higher verbosity level using the **-v** option. This level shows all the commands and responses to and from the MMSC / MSC. To do more debugging, or if the messages from *ha_spng* are not clear, run the MMSC / MSC commands by hand directly after connecting to the MMSC / MSC through **cu -l <tty_dev>**.

Note: Precede all MMSC and MSC commands with Ctrl+T.

The commands *ha_spng* and *ha_killd* send the following commands to the MSC for pings and resets:

- The following command determines the version.

```
MSC> ver
        ok VER 3.0
```

- The following commands replace the password (if set) with none, and reset the MSC.

```
MSC> pas none
        ok
```

```
MSC> rst
        ok
```

For MMSC pings:

- The **ALTERNATE CONSOLE** port (COM5) is in RAT mode by default, so the MMSC prompt might not be visible. To get to the MMSC prompt, *ha_spng* and *ha_killd* send Cut-rate to the tty port, followed by

  R . MMSC

  ```
  MMSC> ^U
        CANCEL
  MMSC> ver
        R1:MMSC 1.0.8
  ```

- The following commands reset the MMSC:

  ```
  MMSC> ^U
        CANCEL
  MMSC> authority service <password_if_present>
        R1:OK
  MMSC> ^U
        CANCEL
  MMSC> R all pwr c 5
        R1U:ok
        R1L:ok      ok
  ```

5.7 Cabling the Storage Systems to the Servers

You can use Challenge Vault, Origin Vault, Challenge RAID, and Fibre Channel RAID and non-RAID storage systems within certain parameters; see the following sections for more information:

- Origin 2000 and Origin200, in Chapter 2:
  - Section 2.8, “Cabling the Challenge or Origin Vaults”
  - Section 2.9, “Cabling the Challenge RAID Storage System to the Origin Servers”
- Challenge XL, L, or DM, in Chapter 3:
  - Section 3.6, “Cabling the Vaults”
  - Section 3.7, “Cabling the Challenge RAID Storage System to the Challenge Servers”
- Challenge S, in Chapter 4: Section 4.7, “Cabling Storage Systems”

Figure 5-24 shows Challenge vault cabling for a mixed configuration.

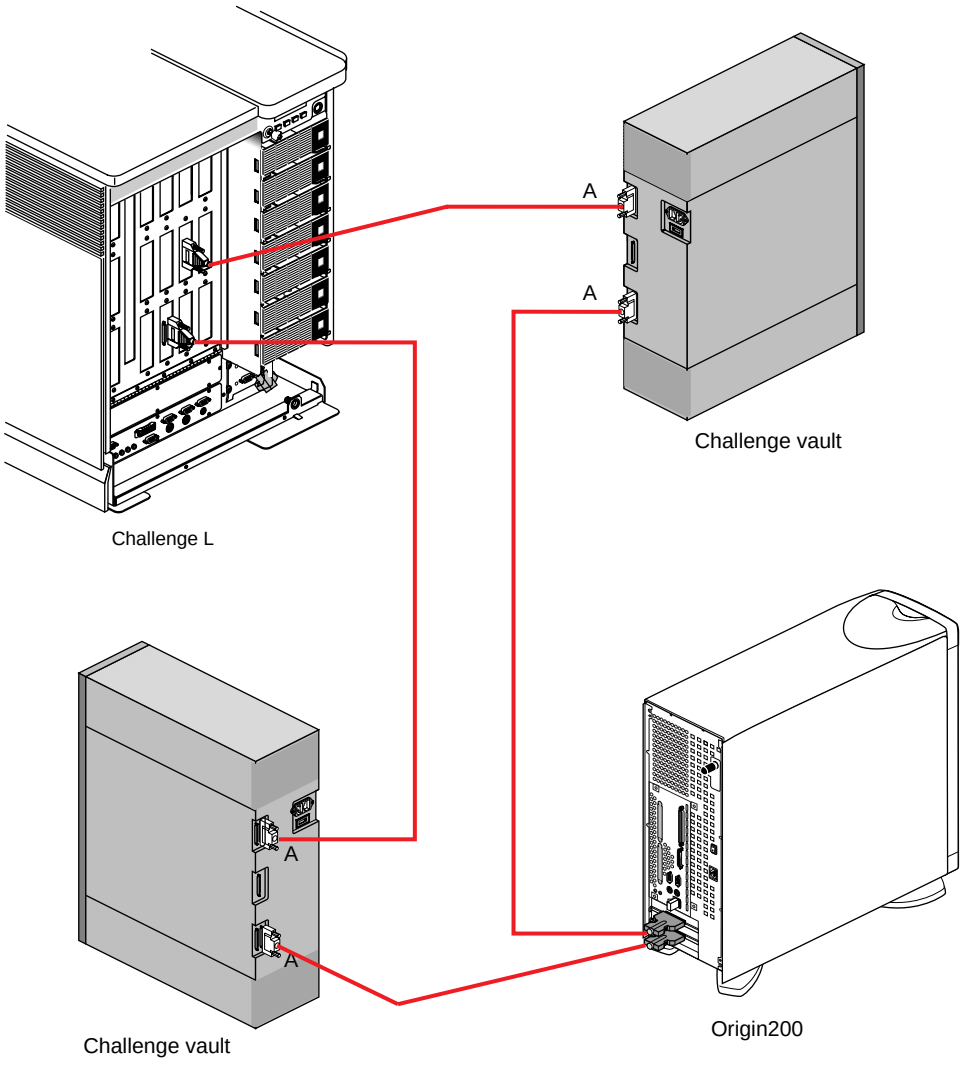

Figure 5-24 Example Challenge Vault Cabling for Mixed Configuration

Figure 5-25 shows Challenge vault cabling for a mixed configuration that includes an Origin200 server.

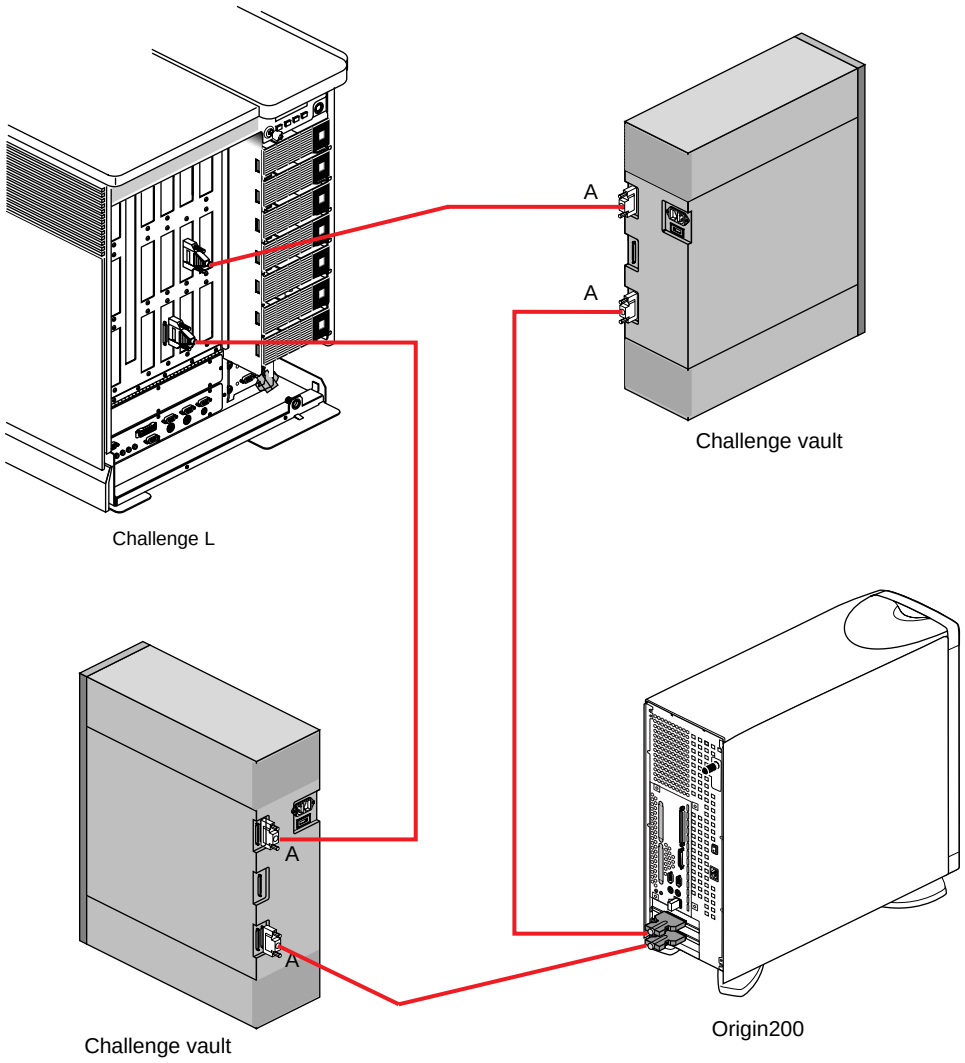

Figure 5-25 Example Challenge Vault Cabling for Mixed Configuration Including Origin200 Server

Figure 5-26 shows Challenge RAID cabling for a mixed configuration.

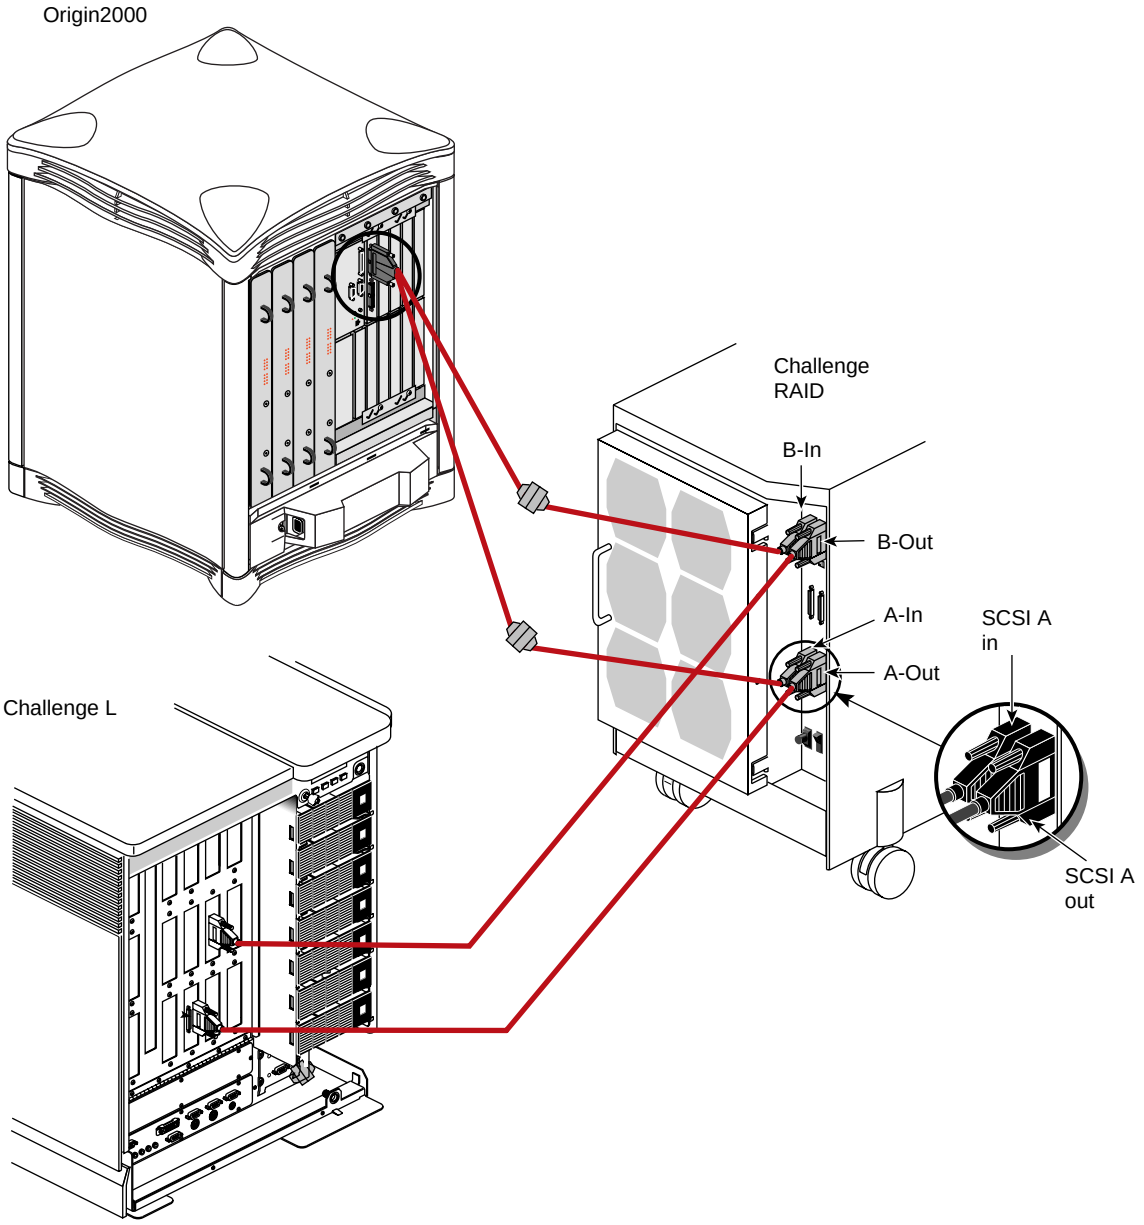

Figure 5-26 Example Challenge RAID Cabling for Mixed Configuration

Figure 5-27 shows Challenge RAID cabling for a mixed configuration that includes an Origin200 server.

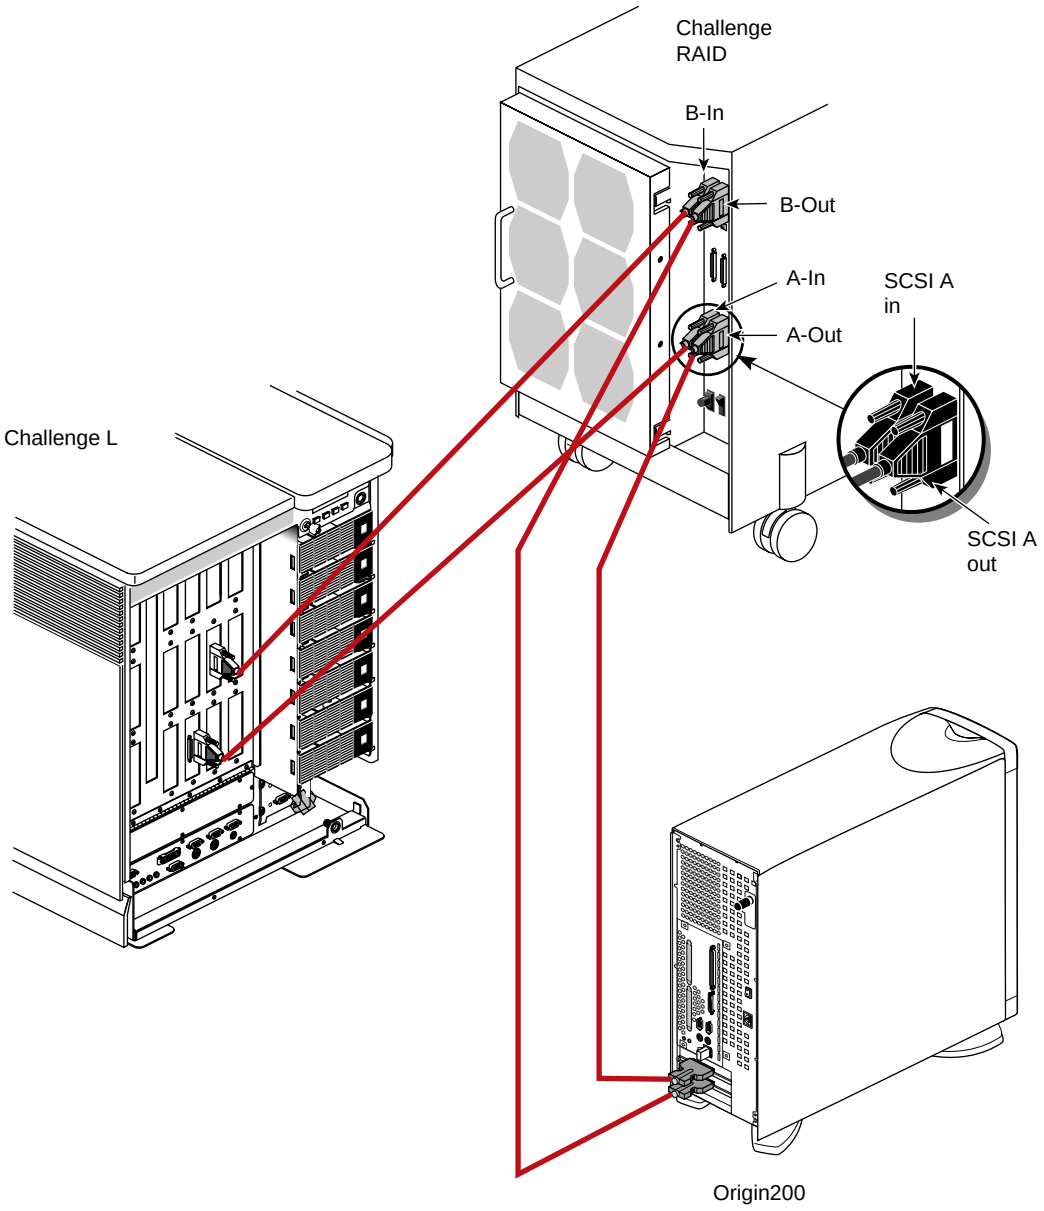

Figure 5-27 Example Challenge RAID Cabling for Mixed Configuration Including Origin200 Server

5.8 Setting the IRIS FailSafe Host SCSI IDs

Follow instructions in Section 2.11, “Setting the IRIS FailSafe Host SCSI IDs,” in Chapter 2.

5.9 Configuring Challenge RAID for IRIS FailSafe

Follow instructions in Chapter 6, “Configuring Challenge RAID Storage Systems for IRIS FailSafe.”

5.10 Configuring and Testing the System

Follow instructions in the latest version of the *IRIS FailSafe Administrator’s Guide* to configure and test the newly installed IRIS FailSafe system. If necessary, also consult the latest edition of *IRIX Admin: Disks and Filesystems*.

Caution: You must stop the RAID agent before creating XLV volumes.

Chapter 6

Configuring Challenge RAID Storage Systems for IRIS FailSafe

If your IRIS FailSafe system uses a Challenge RAID system for shared storage, configure it following the instructions in this chapter.

Note: For information on configuring Silicon Graphics Fibre Channel vaults and disks, see Appendices B and C in the *Origin FibreVault and Fibre Channel RAID Installation Instructions*.

Load the Challenge RAID software from the CD-ROM. The Challenge RAID software images are created in *inst* format; thus, they require that the *inst* utility be used. Installation software and instructions are included with the standard Silicon Graphics software products.

This chapter consists of these sections:

- Section 6.1, “Vault Configuration for IRIS FailSafe”
- Section 6.2, “Challenge RAID Configurations for IRIS FailSafe”
- Section 6.3, “Creating the Challenge RAID Configuration File”
- Section 6.4, “Restarting the Agent and Checking the Configuration File”
- Section 6.5, “Configuring LUNs”
- Section 6.6, “Creating SCSI Device Nodes”
- Section 6.7, “Making the LUN 0 Device in /dev”
- Section 6.8, “Enabling Command-Tagged Queuing on the LUNs”

Note: For complete information on the Challenge RAID software, see the *CHALLENGE RAID Owner’s Guide*.

6.1 Vault Configuration for IRIS FailSafe

These configurations have been designed to ensure availability of services in the following situations:

- disk failures: by plexing shared storage using XLV
- during repair of failed disks: by placing each plex of a shared volume on a separate vault (in multivault systems)
- failure of a single SCSI channel: by placing each vault on a separate SCSI channel

In a single-bus plexed configuration, each volume element on Vault 1-Disk 1 is plexed on Vault 2-Disk 1, as diagrammed in Figure 6-1.

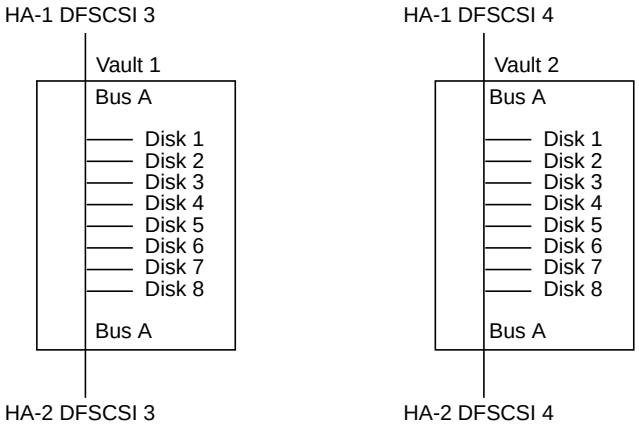

Figure 6-1 Single-Bus Plexed Configuration

In a split-bus plexed configuration, each volume element on Vault 1-Disk 1 is plexed on Vault 2-Disk 1, as diagrammed in Figure 6-2.

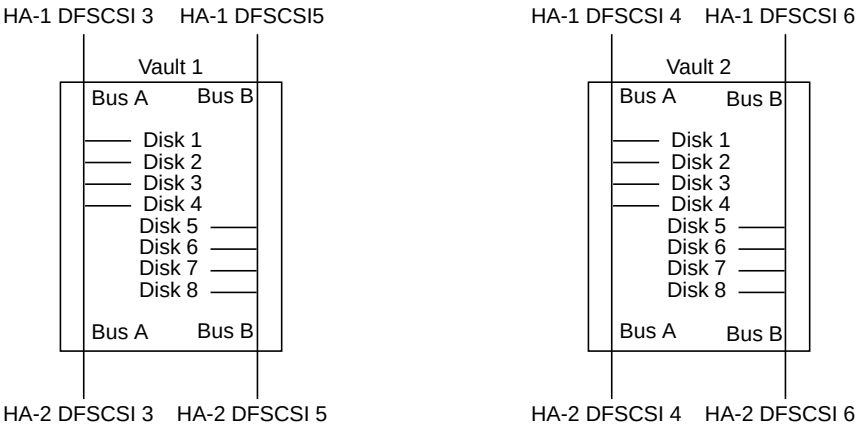

Figure 6-2 Split-Bus Plexed Configuration

The configuration in Figure 6-2 uses more SCSI buses than that in Figure 6-1, but is easier to service because there are fewer disks per SCSI bus. Replacing a faulty disk while the cluster is up requires that you manually stop all I/O to all disks on the same bus as the faulty drive.

6.2 Challenge RAID Configurations for IRIS FailSafe

The only configuration of RAID disks supported is the dual-bus, dual-initiator configuration, which provides the highest availability. Each host has two SCSI-2 adapters, each of which connects by a separate SCSI-2 bus to a separate storage-control processor (SP) in the Challenge RAID storage system. Since this configuration protects against a SCSI-bus cable failure, SP failure, or SCSI-2 adapter failure, it provides higher availability than the dual-initiator configuration for enterprises requiring the highest level of availability.

This configuration was designed for these goals:

- The Challenge RAID storage system allows the online hot replacement of failed disk drives (by owner or SSE) and other redundant system components, such as SPs and power supplies (by SSE only).
- By placing each storage-control processor (SP) of each Challenge RAID storage system on a separate controller, system availability is ensured if a single SCSI channel, a single SP, or a SCSI-2 adapter fails. See Figure 6-3.

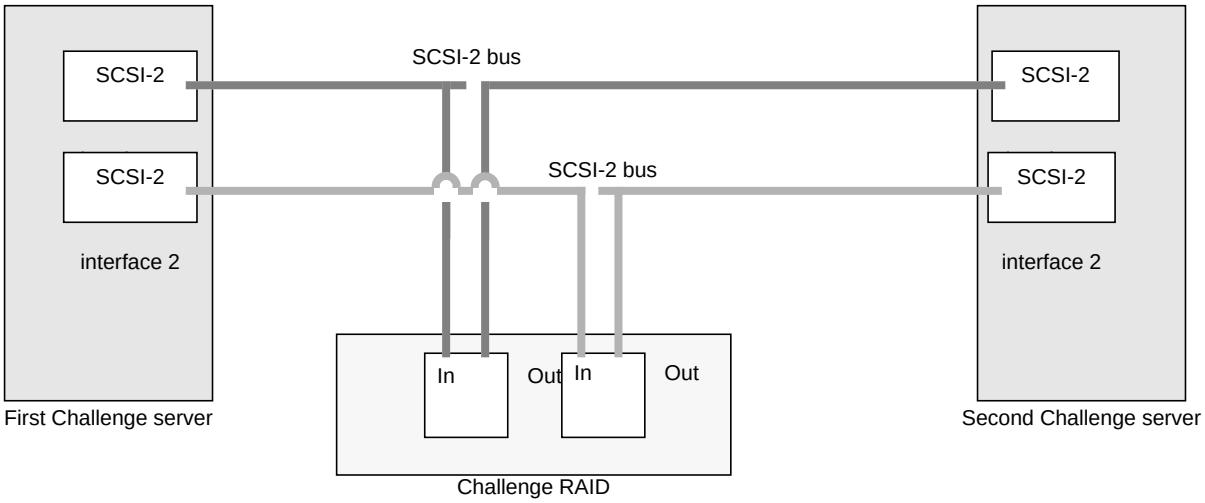

Figure 6-3 Dual-Bus/Dual-Initiator Configuration Example

Caution: Because both hosts have access to all disk modules and their data in this configuration, you must configure the filesystem appropriately and install the IRIS FailSafe software to prevent one host from overwriting the other's data. Furthermore, configuring these systems requires following the guidelines in the *IRIS FailSafe Administrator's Guide*.

Figure 6-4 shows supported Challenge RAID configurations.

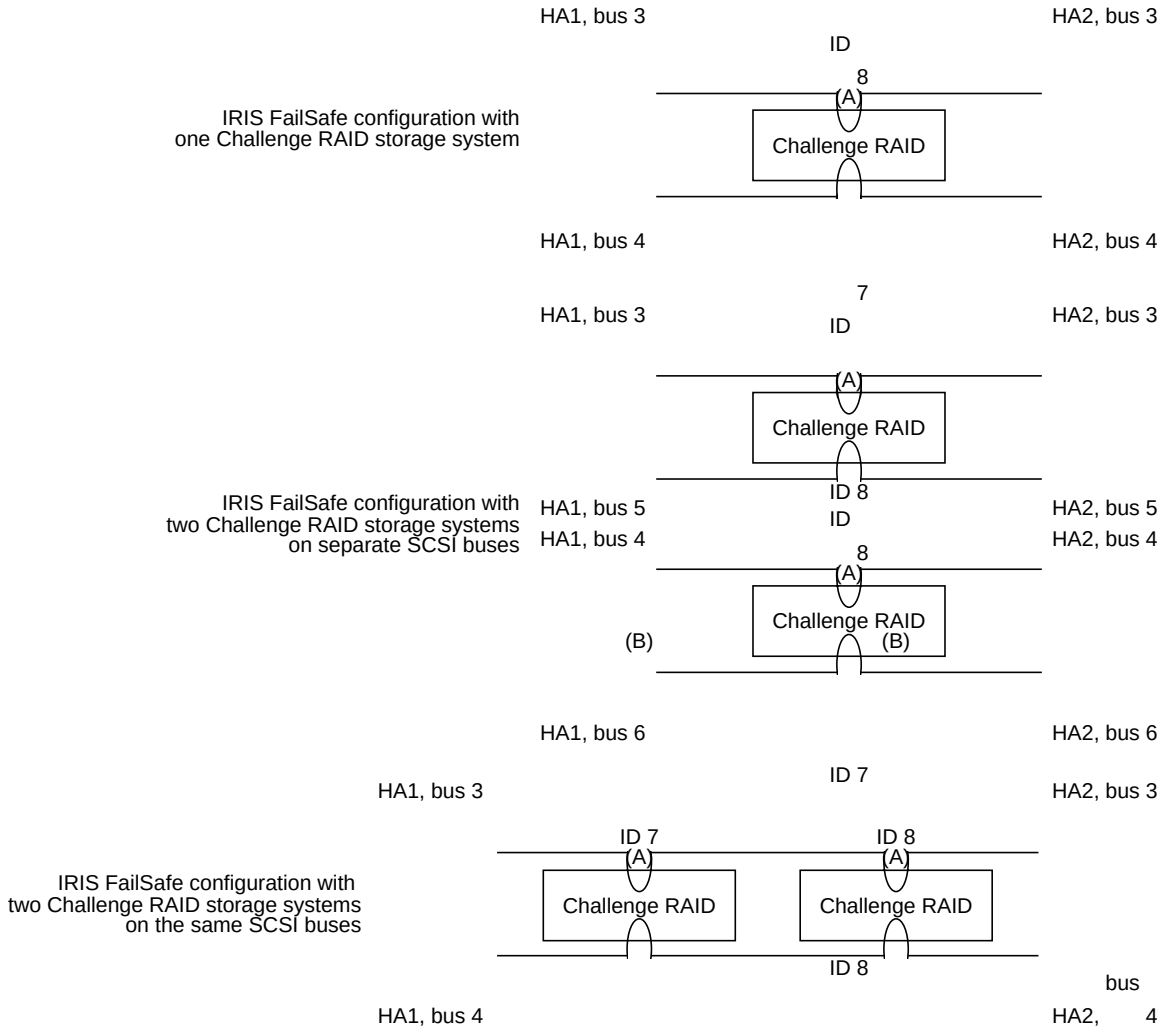

Figure 6-4 Challenge RAID Configurations for IRIS FailSafe Systems

6.2.1 Challenge RAID Configuration With XLV

Using XLV on the Challenge RAID storage system requires certain considerations on the part of the system administrator. Without these considerations, failover functionality can be compromised.

When a Challenge RAID storage system is configured correctly, it can provide an alternate path to peripheral data. XLV takes advantage of this capability and automatically switches to the alternate path if the primary path becomes unavailable, for example, because of a SCSI host adapter failure, a loose cable, or a failing SP. When failure occurs, the XLV software issues a command to the SP on the alternate SCSI channel, causing it to take ownership of the LUN containing the desired data. Subsequent references to this data use the alternate channel.

This XLV path failover mechanism requires the hardware configuration be determined at boot time. The XLV command *xlv_assemble* uses the */dev/scsi* path to determine the hardware configuration.

Challenge RAID administration is done through a daemon called the RAID agent, which also communicates to the RAID through */dev/scsi* device nodes. When it is active, the RAID agent opens the */dev/scsi* device nodes in exclusive mode, so that no other process (including *xlv_assemble*) can open them. As a consequence, *xlv_assemble* does not correctly assemble the failover paths to a RAID device if the RAID agent is active. At boot time, the system takes care to assure that *xlv_assemble* is allowed to complete prior to launching the RAID agent.

The administrator can also run *xlv_assemble* manually. In this case it is important to be sure that the RAID agent has been turned off first:

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

After *xlv_assemble* has completed, the RAID agent must be restarted:

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

6.2.2 Challenge RAID Configuration With the Objectserver

When disks are bound in a LUN, ownership of that LUN goes to the first SP that does an I/O to it after the bind operation is complete. If the objectserver is running, ownership of the newly created LUN is most likely given to the SP connected to the lowest-numbered SCSI channel. If both SPs are on the same SCSI channel, ownership goes to the SP with the lowest TargetID.

Once ownership is assigned, access to that LUN through the *dksc* interface (*/dev/dsk* device nodes) is available only through the SP that has ownership. Therefore, the objectserver must not be running when a LUN is being bound.

To stop the objectserver, enter the following:

```
# /etc/init.d/cadmin stop
```

To start the objectserver, enter the following:

```
# /etc/init.d/cadmin start
```

In a server configuration, leave the objectserver off. If it is on, enter the following command to configure it off permanently:

```
# /etc/chkconfig objectserver off
```

6.3 Creating the Challenge RAID Configuration File

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

The Challenge RAID agent 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

  These lines specify the primary LUN for each SP (0 for SP A and 1 for SP B). Entries for any other LUNs are not needed. The file must have an entry for each SP; for example:

  ```
  device sc40d1l0 sgi-clariion-SPA “SP-A on SCSI 40”
  device sc41d1l1 sgi-clariion-SPB “SP-B on SCSI 41”
  ```

- 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

To create the Challenge RAID configuration file, follow these steps:

1. Run *hinv*.
2. Open */usr/raid5/raid5-agent.config.proto* in your favorite editor. This file is self-documenting and contains samples.
3. Save */usr/raid5/raid5-agent.config.proto* as */usr/raid5/raid5-agent.config*.
4. In the RAID specification section, add RAID devices, using the controller numbers displayed in the *hinv* output.
5. Add the storage-control processors (SPs); for example:

   ```
   device sc6d6l0 sgi_clariion_SPA “SGI Agent Development System”
   device sc6d5l0 sgi_clariion_SPB “SGI Agent Development System”
   ```

   Note: In */usr/raid5/raid5-agent.config*, only the SCSI device number for LUN 0 is required.
6. In the User Specification section, enter the user IDs of those who can perform write operations.
7. In the Poll Specification section, enter the poll specification.
8. Save and exit the configuration file. If the configuration file is correct, the following message appears in *syslog*:

   ```
   RAID5 Agent [pid]: Starting up
   ```

   Otherwise, this message appears in */var/adm/SYSLOG*:

   ```
   RAID5 Agent [pid]: Is the device configured correctly?
   ```

9. Restart the agent as explained in Section 6.4, “Restarting the Agent and Checking the Configuration File.”

If the customer is using RAID level 5, you should not have to reconfigure LUNs. For information on configuring LUNs, see the *Challenge RAID Installation and Maintenance Instructions*.

6.4 Restarting the Agent and Checking the Configuration File

To restart the Challenge RAID software after editing */usr/raid5/raid5-agent.config*, follow these steps:

1. Enter as root:

   ```
   /etc/init.d/raid5 stop
   /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 the *CHALLENGE RAID Owner’s Guide*.

   Note: Error messages from the agent are not reported to the console, but are stored in the *syslog* file.

2. Use the */usr/raid5/raid5 getagent* command to display information on RAID devices. The output of this command covers all devices. Since the disks have not been bound into LUNs, information for only the storage control processors (SPs) is shown; for example:

   ```
   Name: sgi-clariion-SPA
   Desc: SGI Agent Development System+SGI Agent Development System
   Node: sc6d6l0
   Signature: 0xcde20300
   Peer Signature: 0x11151f00
   Revision 7.57.0
   SCSI Id: 6
   Prom Rev: 0x00076300
   SP Memory: 64
   Serial No: 94-2555-445

   Name: sgi-clariion-SPB
   Desc: SGI Agent Development System
   Node: sc7d7l0
   Signature: 0x11151f00
   Peer Signature: 0xcde20300
   Revision 7.57.0
   SCSI Id: 7
   Prom Rev: 0x00076300
   SP Memory: 64
   Serial No: 94-2555-445
   ```

   The output lists the SCSI ID numbers and SCSI node numbers for each SP.

6.5 Configuring LUNs

A usable RAID unit is known by a *logical unit number*, or LUN. The LUN is a decimal (actually octal) number between 0 and 7, even though no more than four LUNs can be specified for the Challenge RAID storage system.

To group physical disks (disk modules) into units, use the */usr/raid5/raidcli bind* command. This command has several required and optional parameters; for details, see Appendix B of latest version of the *CHALLENGE RAID Installation and Maintenance Instructions*. For steps on binding LUNs, see Chapter 4 in that manual.

6.6 Creating SCSI Device Nodes

Note: This section applies only to systems using IRIX 6.2. For later versions, skip these instructions.

To create SCSI device nodes, follow these steps:

1. As superuser, enter this command in the */dev* directory:

   `MAKEDEV dks scsi`

   The result of this command is that a device is created for each LUN:

   - LUN 0: */dev/{r}dsk/dks<controller-#>d<drive-#>{s<partition-#> | vh | vol*

     Note that the device for LUN 0 is different from that for the other LUNs.

   - LUN 1 through 7:
     /dev/{r}dsk/dks<controller-#>d<drive-#>l<lun-#>{s<partition-#> | vh | vol}

   Note: If the Challenge RAID storage system is not connected to a Challenge server during the software installation process, no device nodes are created.

2. Use *ls -l* in the */dev/scsi* directory to display the SCSI device nodes. The following is a fragment of an example output for the first SP:

```
crw-r--r--    1 root     sys      195,387 Mar  6 12:30 /dev/scsi/sc6d6l0
crw-------    1 root     sys      195,403 Mar 21 15:40 /dev/scsi/sc6d6l1
crw-------    1 root     sys      195,419 Mar 21 15:40 /dev/scsi/sc6d6l2
crw-------    1 root     sys      195,435 Mar 21 15:40 /dev/scsi/sc6d6l3
crw-------    1 root     sys      195,451 Mar 21 15:45 /dev/scsi/sc6d6l4
crw-------    1 root     sys      195,467 Mar 21 15:45 /dev/scsi/sc6d6l5
crw-------    1 root     sys      195,483 Mar 21 15:45 /dev/scsi/sc6d6l6
crw-------    1 root     sys      195,499 Mar 21 15:45 /dev/scsi/sc3d3l7
```

   In this example, the first line lists LUN 0 and the next seven lines show LUN 1-7. In this address:

   - The digit after `sc` is the SCSI controller number.
   - The digit after the letter `d` in this address is the number of the RAID disk module.
   - The digit after the letter `l` in this address (the last character in the address) is the LUN number.

6.7 Making the LUN 0 Device in /dev

Note: This section applies only to systems using IRIX 6.2. For later versions, skip these instructions.

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. Also, in a system with two SPs that are used for primary and secondary paths, both SPs must be made aware of the new disks. This section explains how to accomplish this task without rebooting.

For each SP, you must create the LUN 0 device, that is, create the device nodes for the new disks. In the */dev* directory, enter as root for each SP:

./MAKE_VLUNS *controller_number target_number*

In this command line, *controller_number* is the SCSI bus to which an SP is connected and *target_number* is the SCSI ID of the SP. For example:

```
./MAKE_VLUNS 5 2
```

To bind the newly installed disk modules into one or more physical disk units, see the *Challenge RAID Installation and Maintenance Instructions*.

6.8 Enabling Command-Tagged Queuing on the LUNs

Once the LUNs are bound, use *fx* to partition the LUNs as needed. The syntax of the *fx* program is as follows:

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

For more information on this command, consult its reference page, fx(1M).

Follow these steps:

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

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

   Output such as the following appears:

   ```
   fx version 5.3, Jan 3, 1995
   ...opening dksc(6,6,2)
   ...controller test...OK
   Scsi drive type == SGI RAID 5 0767
   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, enter 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, enter **`enable`**:

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

4. At the `CTQ depth` prompt, enter **`12`**:

   ```
   fx/label/set/parameters: CTQ depth = (2) 12
   ```

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

5. At the prompt in the last line in step 4, enter **`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. To repartition the disk, enter **`r`**.

7. Enter **`exit`** to exit *fx*. The following message appears:

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

8. Enter **`y`** to write the changes to disk.

Chapter 7

Maintaining and Upgrading the High-Availability System

This chapter consists of these sections:

- Section 7.1, “Replacing the Serial Connection”
- Section 7.2, “Isolating a Node”
- Section 7.3, “Replacing a Challenge Vault Disk Module”
- Section 7.4, “Replacing Origin Vault Disk Drive Modules”
- Section 7.5, “Warm Swapping Origin Vault Disk Modules”
- Section 7.6, “Exchanging Challenge RAID Disk Modules”
- Section 7.7, “Replacing Fibre Channel Disk Modules”
- Section 7.8, “Restoring LUN Ownership on a Challenge RAID System”
- Section 7.9, “Replacing Batteries in the Remote Power Control Unit”

Note: For troubleshooting information and *ha_cfverify* error messages, see the *IRIS FailSafe Administrator’s Guide.*

7.1 Replacing the Serial Connection

To replace the serial connection to the system controller port (Challenge DM/L/XL) or to the remote power control unit (Challenge S), follow these steps:

1. Turn off monitoring the serial connection to a server:

   `ha_admin -m stop` *servername*

   This command returns

   ```
   ha_admin: Stopped monitoring the serial connection to xfs-ha5
   ```

2. Replace the defective serial cable or remote power control unit.
3. To verify that the cable is functioning, enter

   `ha_spng -i10 -f <reset-tty>`

   and check the return code.

Caution: Do not run *ha_spng* unless you have stopped the IRIS FailSafe software from monitoring, as directed in step 1. Omitting this step confuses the IRIS FailSafe software.

For example, if you are running *csh*, and *reset-tty = dev/ttyd2*, type

```
# ha_spng -i 10 -f /dev/ttyd2
# echo $status
0
```

The zero at the end of this output indicates normal operation.

4. To turn on monitoring the serial connection to a server, enter

   `ha_admin -m start` *servername*

   This command returns

   ```
   ha_admin: Started monitoring the serial connection to
   <servername>serial
   ```

7.2 Isolating a Node

This section explains how to isolate a node (for example, node 1) with a failed component from the other node (in this case, node 2), while the system continues to run the IRIS FailSafe software. This technique allows several maintenance procedures to be applied. For example, you can replace the IO4 or the SCIP mezzanine card that affects the SCSI chains joining the two nodes.

Isolating a node (Challenge server) is explained in these sections:

- Section 7.2.1, “Stopping IRIS FailSafe and Halting the Node”
- Section 7.2.2, “Disconnecting a Node From a SCSI Bus”
- Section 7.2.3, “Replacing SCSI Cables”
- Section 7.3, “Replacing a Challenge Vault Disk Module”

Note: For this process, you need one differential SCSI terminator for each shared SCSI bus (Challenge RAID or vault system) to which the node is attached.

7.2.1 Stopping IRIS FailSafe and Halting the Node

Follow these steps:

1. To stop IRIS FailSafe on the node to be isolated, enter as root:

   # **`/etc/init.d/failsafe stop`**

2. Wait until the node has successfully transitioned to standby state and the failsafe processes have exited. This process can take a few minutes.

3. To prevent accidental reintegration of the node, *chkconfig* the failsafe software off:

   # **`/sbin/chkconfig failsafe off`**

Note: If the remaining node were to be restarted after this node is isolated (removed), it would move to standby state. To move it from standby state to degraded state, enter **`ha_admin -G`**. The remaining node stays in degraded state, because its peer is not available.

4. To halt the node you are isolating, enter

   `halt`

5. On a Challenge S server, disconnect its power cord from the outlet. On a larger Challenge server, turn the key on the front panel to the off position.

7.2.2 Disconnecting a Node From a SCSI Bus

In a dual-hosted environment, the two nodes provide termination for the SCSI bus. Since an unterminated SCSI bus causes system failures, the procedure for disconnecting SCSI buses must ensure correct termination.

For example, nodes A and B share two SCSI buses, DF SCSI4 and DF SCSI5:

- DF SCSI4 has three devices: RAID SP A, vault 1A, and vault 2A.
- DF SCSI5 has three devices: RAID SP B, vault 1B, and vault 2B.

Each volume on a vault in this example is a plex composed of a volume element from vault 1A and 1B, or vault 2A and 2B. Figure 7-1 shows this configuration.

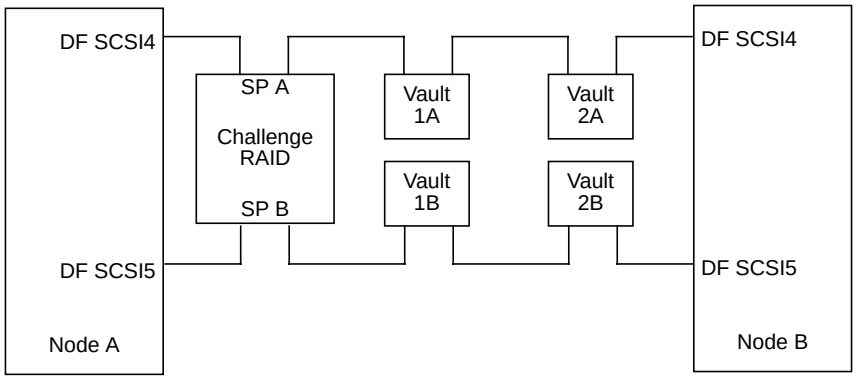

Figure 7-1 Example Shared Storage

In this example, to isolate node B, you would take these steps:

- Isolate host A from SCSI bus DFSCSI-4:
  - Stop I/O to RAID SP A, vault 1A, and vault 2A.
  - Disconnect and terminate the SCSI bus at vault 2A (since it is closest to node B, which is being isolated).
  - Enable I/O to RAID SP A, vault 1A, and vault 2A.

- Isolate host A from SCSI bus DFSCSI-5:
  - Stop I/O to RAID SP B, vault 1B, and vault 2B.
  - Disconnect and terminate the SCSI bus at vault 2B (since it is closest to node B, which is being isolated).
  - Enable I/O to RAID SP B, vault 1B, and vault 2B.

To disconnect the node from a SCSI bus, follow these steps:

1. Quiesce the bus by stopping I/O to all devices connected to the bus. To make this suspension of I/O transparent to applications, you must ensure that an alternate path to the data exists.

   On a Challenge RAID storage system, this alternate path exists because the LUNs can be accessed via the SP on another SCSI bus. For a Challenge vault, the data is plexed using XLV on another Challenge vault on a different SCSI bus.

   - To stop I/O to a vault that has been mirrored using the XLV plexing option, detach each plex on the vault from its logical volume, as explained in "Creating and Administering Volumes" in *IRIX Admin: Disks and Filesystems*.

     Note: You must perform the detach procedure on the node that owns the XLV volume. In the IRIS FailSafe configuration, this node is the one that is currently mounting filesystems on the volume.

   - To stop I/O to an SP on the Challenge RAID storage system, disable power to the SP, as explained in Chapter 5 of the *CHALLENGE RAID Installation and Maintenance Instructions* (108-0128-006 or later).

2. Disconnect the SCSI cable at the device end (Challenge RAID or vault system) of the bus closest to the node you are isolating.

3. Terminate the device with a differential SCSI terminator:

   - For a vault, remove the SCSI cable connecting the vault to the node you are isolating. Place a differential SCSI terminator on the vault's SCSI connector, as diagrammed in Figure 7-2.

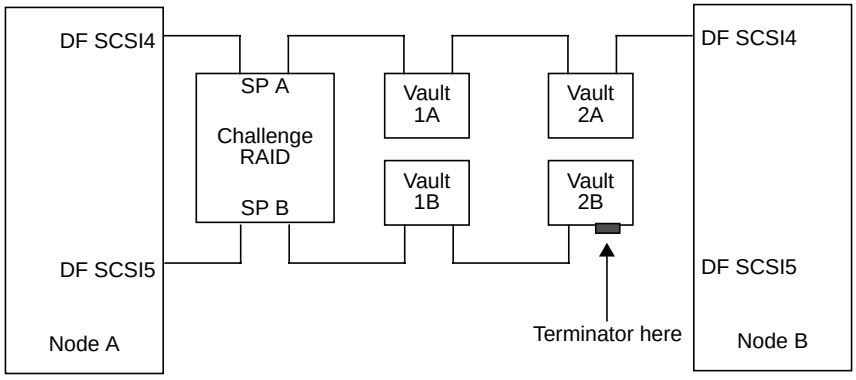

Figure 7-2 Terminating a Vault (Isolating Node B)

For a Challenge RAID storage system, remove the SCSI cable connecting the node being isolated to the RAID SP. Attach a differential SCSI terminator to the SP's SCSI connector, as diagrammed in Figure 7-3.

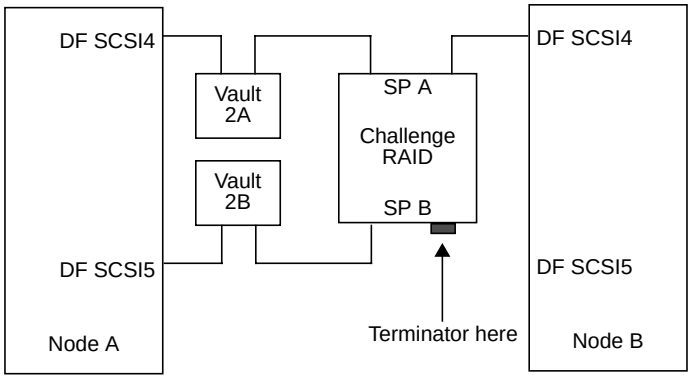

Figure 7-3 Terminating a Challenge RAID Storage System (Isolating Node B)

4. Enable the alternate path that was broken in step 1 so that the system can resume activity on the SCSI bus:

   - For a plexed vault on the SCSI bus, follow the XLV procedures for attaching a plex to a logical volume, as explained in "Creating and Administering Volumes" in *IRIX Admin: Disks and Filesystems*. Do this for each volume element on the vault.

     Caution: The plex revive operation can take up to a few hours, depending on the volume of data being revived, and the activity on the system. Do not proceed beyond this step until the revive operation has completed on all plexed volumes.

     You must perform the plex attach procedure on the node that owns the XLV volume. In the IRIS FailSafe configuration, this node is the one that is currently mounting filesystems on the volume.

     When this procedure is used for purposes other than isolating a node (for example, for replacing a nonfunctional vault disk), the node on which the revive operation should be performed might be different from the node on which the earlier detach plex procedure was performed, if a failover of filesystems occurred between these two events.

   - For a Challenge RAID device on the SCSI bus, turn on the SP to enable I/O to the bus. The SP takes about five minutes to come online. After five minutes, enter the *raidcli* command **`/usr/raid5/raidcli getcrus`**. The output should contain these lines to show that both SPs are online:

     ```
     SPA State: Present
     SPB State: Present
     ```

5. If the SCSI chain has more than one device connected to it, repeat steps 1 through 4 for each device.

7.2.3 Replacing SCSI Cables

To replace a faulty SCSI cable, follow these steps:

1. Stop I/O to all devices on the SCSI bus that has the faulty cable, as explained in step 1 in Section 7.2.2, "Disconnecting a Node From a SCSI Bus."

   For RAID devices, if the faulty cable has already resulted in I/O errors, XLV has masked these errors by switching to an alternate path. The bus is probably already quiesced, but it is still safer to follow the procedures documented in step 1 in Section 7.2.2, "Disconnecting a Node From a SCSI Bus."

   Note: If an XLV volume cannot be detached (for example, the detach operation returns an error), continue stopping I/O on all other devices and proceed with steps 2, 3, and 4. If all XLV volumes can be detached, proceed to steps 2 and 4, skipping step 3.

2. Replace the faulty cable.
3. Retry the XLV detach operation on all volumes that returned an error in step 1. They should all succeed.
4. Enable I/O to all devices on the SCSI bus as documented in step 4 in Section 7.2.2, "Disconnecting a Node From a SCSI Bus."

7.3 Replacing a Challenge Vault Disk Module

Caution: This procedure causes the IRIS FailSafe software to power cycle the node that currently owns the volume with the failed disk. For a short period, services provided by the node will not be available. It is recommended that you perform this procedure during periods of light load, or during scheduled downtime.

To replace the disk, the entire vault must be powered off. Because powering off a vault can potentially disrupt I/O on the SCSI bus, follow these steps:

1. Identify all devices on the same SCSI bus as the faulty vault. Stop I/O to all these devices, as documented in step 1 in Section 7.2.2, "Disconnecting a Node From a SCSI Bus."
2. Use the *xlv_mgr detach -force* option to detach the plex containing the faulty disk.
3. Wait for HA1 to reboot. Wait for the IRIS FailSafe software to move into normal state, as displayed by *ha_admin -i*.
4. Power off the vault with the faulty disk, replace the disk, and power on the vault. Wait for all the disks to spin up.
5. Remake the plex object that contained the faulty disks, using *xlv_make* as documented in *IRIX Admin: Disks and Filesystems*.
6. Enable I/O to all devices on the SCSI bus as documented in step 4 in Section 7.2.2, "Disconnecting a Node From a SCSI Bus." This operation requires reattaching all detached plexes from step 1 to their respective subvolumes.

7.4 Replacing Origin Vault Disk Drive Modules

Replace an Origin Vault disk drive module with another such module from Silicon Graphics. Disk drives ordered from Silicon Graphics are preconfigured; you can install them in the bays with no additional configuration. If you are installing third-party peripherals, or wish to change the configuration of a Silicon Graphics supplied drive, follow instructions in the *Origin Vault Owner's Guide* or the *Origin200 and Origin Vault Installation Instructions.*

Caution: Origin Vault drive modules are not hot-pluggable. However, you can exchange a 3.5-inch disk drive module (except for the system disk) in the Origin Vault option without powering it off; see Section 7.5, "Warm Swapping Origin Vault Disk Modules."

To replace an Origin Vault disk drive module, follow the steps below:

1. Make sure no users are using the data stored in the expansion option. Check the status LEDs for all drives to make sure they are not active.
2. Unlock and open the Origin Vault front door using the key included with the option.
3. Remove each new disk drive module and place it on the antistatic packaging. Attach a label identifying the slot into which it is to go.
4. Turn off the power button at the front; the LED goes dark. Turn off the power on the back of the Origin Vault expansion option.
5. Attach the grounding strap to ground yourself to a metal part of the Origin Vault.
6. For the disk drive you are removing, pull the drive handle down (tower) or leftward (rackmounted enclosure) as shown in Figure 7-4 and Figure 7-5.

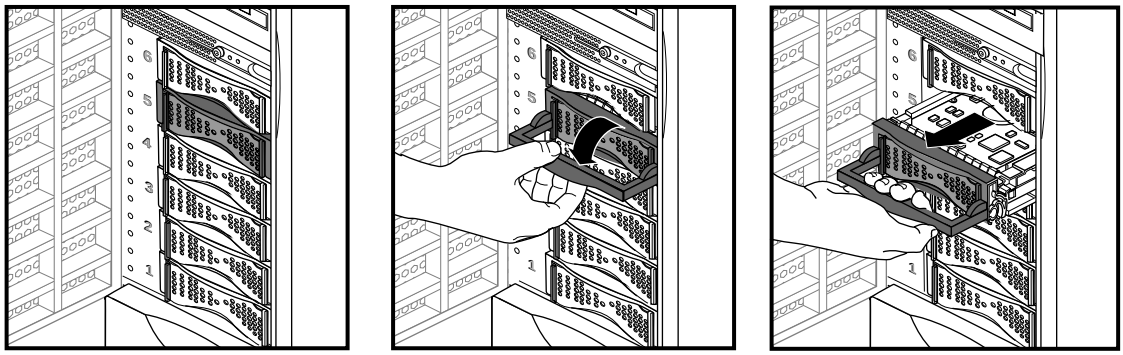

Figure 7-4 Removing a 3.5-Inch Disk, Standalone Tower

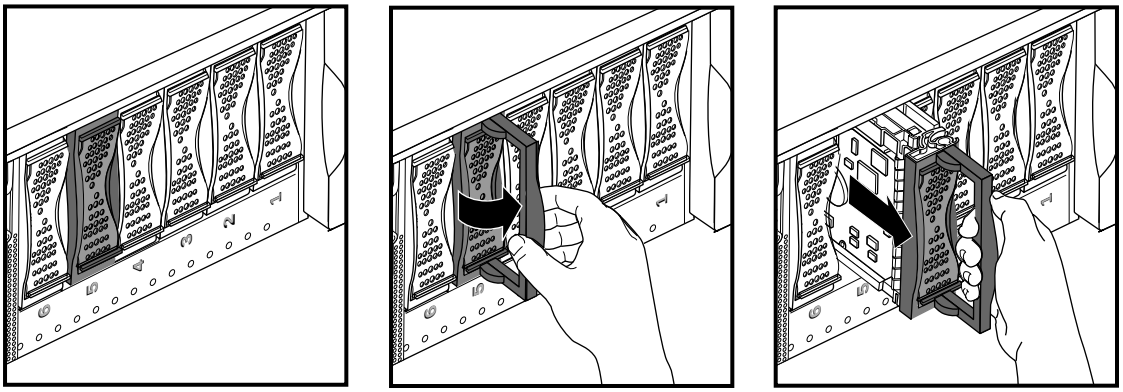

Figure 7-5 Removing a 3.5-Inch Disk, Rackmountable Enclosure

7. Insert each drive sled into a bay until it is almost flush with the edge of the drive bay, as shown in Figure 7-6 and Figure 7-7.

   Latch the drive: For a tower, pull the handle up to latch the drive; for an enclosure in a rack, pull the handle to the left.

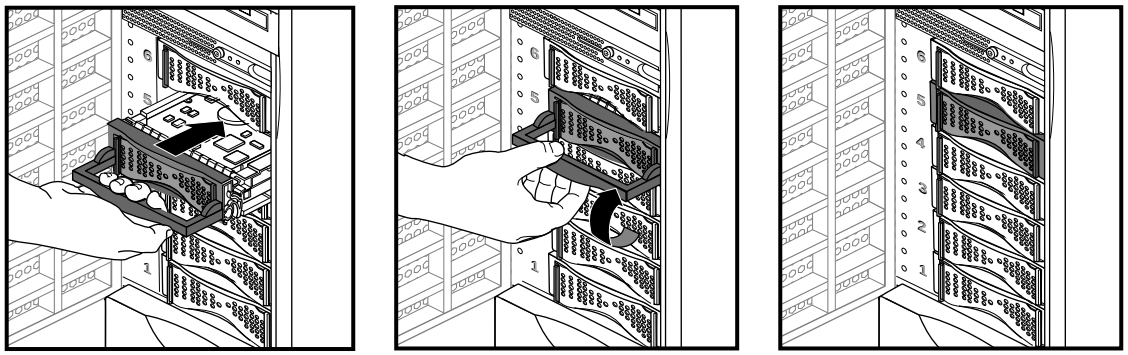

Figure 7-6 Inserting a 3.5-inch Drive, Standalone Tower

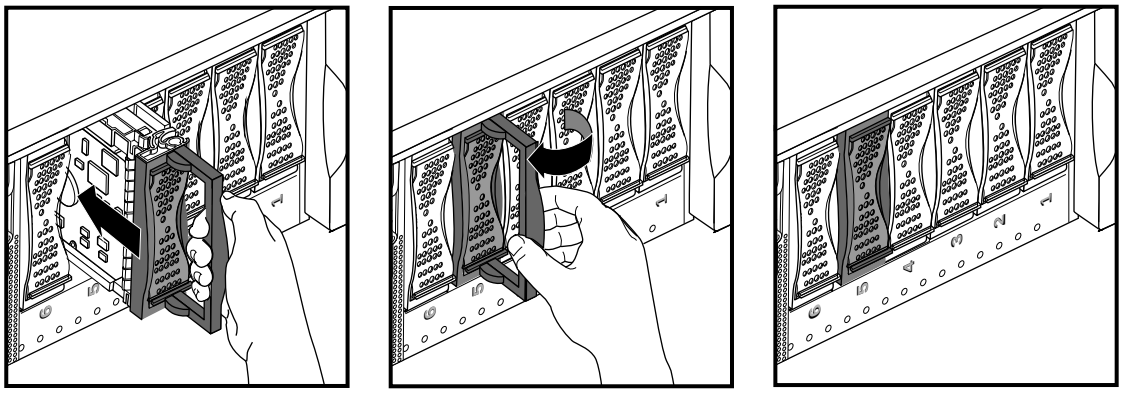

Figure 7-7 Inserting a 3.5-inch Drive, Rackmountable Enclosure

8. Remove the grounding strap.

9. Turn on the circuit breaker at the back of the Origin Vault option; turn on the power button on the front.
10. To update the system software to recognize a new drive, reboot. Consult the latest edition of *IRIX Admin: Disks and Filesystems* if necessary.

7.5 Warm Swapping Origin Vault Disk Modules

Origin Vault 3.5-inch disk modules are not hot-pluggable, but they can be warm swapped: you can exchange a disk drive module (except for the system disk) in the Origin Vault option without powering it off. This procedure is known as administrative warm swap. In the Origin Vault option, only 3.5-inch drives can be warm swapped.

Administrative warm swap utilizes *scsiadminswap*. When this command is in effect, the SCSI bus for the drive you are swapping is quiesced; that is, all commands for this bus are held off. The default bus-quiesce time is 300 seconds (5 minutes), but you can reset this time limit.

Note: Besides *scsiadminswap*, this procedure uses system commands that are fully explained in *IRIX Admin: Disks and Filesystems*; have this publication available before you start the procedure. If the disk is part of an XLV volume, see that publication or man pages for fuser(1M), xlv_mgr(1M), xlv_make(1M), xlv_assemble(1M), and xlv_shutdown(1M). See scsiadminswap(1M) for information on that command.

The administrative warm swap process is explained in these sections:

- Section 7.5.1, “Removing and Replacing the 3.5-Inch Disk Drive Module”
- Section 7.5.2, “Configuring Filesystems on the Replacement Disk Drive Module”

7.5.1 Removing and Replacing the 3.5-Inch Disk Drive Module

Follow these steps:

1. Make sure that the disk drive module you want to swap is not the system disk. To replace the system disk, you must power down the Origin Vault system, as explained in Section 7.4, “Replacing Origin Vault Disk Drive Modules.”
2. Have ready a grounding strap and the replacement disk drive module.

   Disk drive modules ordered from Silicon Graphics are preconfigured; you can install them in the bays with no additional configuration. If you are installing third-party peripherals, or wish to change the configuration of a Silicon Graphics supplied drive, follow guidelines in the *Origin Vault Owner’s Guide* or the *Origin200 and Origin Vault Installation Instructions*.
3. Make sure no users are using any filesystems on the disk.
4. Determine the SCSI bus for this drive.

Note: Once the command for the administrative warm swap is issued, the bus for this drive is quiesced; that is, all commands on this bus are held off for the default bus-quiesce time of 300 seconds (5 minutes) or the time you set. If appropriate, inform users that the bus is not available for this period.

5. As root, unmount all filesystems of the disk you wish to replace; see umount(1M). Use *fuser* to find and kill processes that use a filesystem.
6. For XLV volumes, enter commands as follows:
   - If a filesystem on the disk is an unmirrored XLV volume, enter *xlv_shutdown*.
   - If a filesystem on the disk is a mirrored XLV volume, enter *xlv_mgr* and perform a detach of the plex that uses this disk.

   Note: See *IRIX Admin: Disks and Filesystems* or man pages for xlv_mgr(1M), xlv_make(1M), xlv_assemble(1M), or xlv_shutdown(1M).
7. Enter

   `scsiadminswap -u -b # -d # -t #`

   where

   -p indicates that this operation is an unplug

   -b *#* specifies the SCSI bus for the device

   -d *#* specifies the drive's slot number

   -t *seconds* changes the time to quiesce the bus from the default 300 seconds to the number of seconds specified; there are no lower or upper limits

   Caution: This command must be allowed to complete. Once you have pressed Enter for this command, use of the bus for this disk is held off. If the disk drive module cannot be replaced within five minutes (or the time set with the **-t** option), do not kill the *scsiadminswap* command, which leaves the SCSI bus in a bad state. Let the command finish, and retry the process later.

   The following example sequence of commands is used for swapping a disk in slot 2, with partition 7 mounted as */x*:

   ```
   umount /x
   scsiadminswap -p -u -b 0 -d 2
   mount -t xfs /dev/dsk/dks0d2s7 /x
   ```
8. Attach the grounding strap to ground yourself to a metal part of the Origin Vault.

   Note: Be prepared to remove the disk drive module and insert its replacement within five minutes, or the time set with the **-t** option in step 7.
9. Remove the disk drive module by pulling the drive handle down (tower) as shown in Figure 7-4, or leftward (rackmounted enclosure) as shown in Figure 7-5, earlier in this chapter.
10. With the grounding strap attached, insert the replacement disk drive module until it is almost flush with the edge of the drive bay.
11. Latch the drive: For a tower, pull the handle up to latch the drive; for an enclosure in a rack, pull the handle to the left. Remove the grounding strap.

12. Enter

    `scsiadminswap -p -b # -d #`

    where **-p** indicates that this operation is a plug operation. If desired, you can use the **-t** option to change the default bus-quiesce time.

7.5.2 Configuring Filesystems on the Replacement Disk Drive Module

Set up the filesystems on the replacement disk depending on the type of filesystem that goes on it:

- non-XLV volume:

  Mount the filesystem using *mount*.

- unmirrored XLV volume with no prior XLV information on the drive:
  1. Use *xlv_make*.
  2. Use *xlv_mgr* and perform an attach.
  3. Mount the filesystem using *mount*.
- mirrored XLV volume with no prior XLV information on the drive:
  1. Mount the filesystem using *mount*.
  2. Use *xlv_assemble*.
- unmirrored XLV volume with prior XLV information on the drive:
  1. Use *xlv_mgr*.
  2. Delete the old XLV label information.
  3. Use *xlv_make*.
  4. Use *xlv_mgr* and perform an attach.
  5. Mount the filesystem using *mount*.
- mirrored XLV volume with prior XLV information on the drive:
  1. Use *xlv_mgr*.
  2. Delete the old XLV label information.
  3. Mount the filesystem using *mount*.
  4. Use *xlv_assemble*.

Note: After you have set up the filesystems, notify users that they can write to the disk and the SCSI bus.

The following example sequence of commands is used to swap a disk containing old XLV information in slot 2 with partition 7 mounted as an unmirrored XLV volume named */lvl*:

```
umount /lv1
xlv_shutdown -n sample_vol
scsiadminswap -p -u -b 0 -d 2
xlv_mgr -xc “delete label /dev/rdsk/dks0d2vh”
```

```
xlv_make (see xlv_make man page for examples)
xlv_mgr (perform attach)
mount -t xfs /dev/dsk/dks0d2s7 /lv1
```

7.6 Exchanging Challenge RAID Disk Modules

This section reproduces information from the *CHALLENGE RAID Installation and Maintenance Instructions* for convenience. For complete instructions on replacing field-replaceable units, see that manual.

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

- 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 special system firmware for Challenge RAID storage systems. Do not remove disk modules from bus 0 (the top five modules on the right) 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. The replacement 2 GB drive part number is 9410113; the replacement 4.3 GB drive part number is 9410114. 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 storage-control processor slot.
- SCSI terminators must be in place at all times during operation. Removing them causes the storage system to crash.
- If you replace a disk module, you must update the firmware as explained in Section 6.5.2.4, “Updating the Disk Module Firmware.”

Figure 7-8 diagrams location of disks, which can be replaced by owners or SSEs.

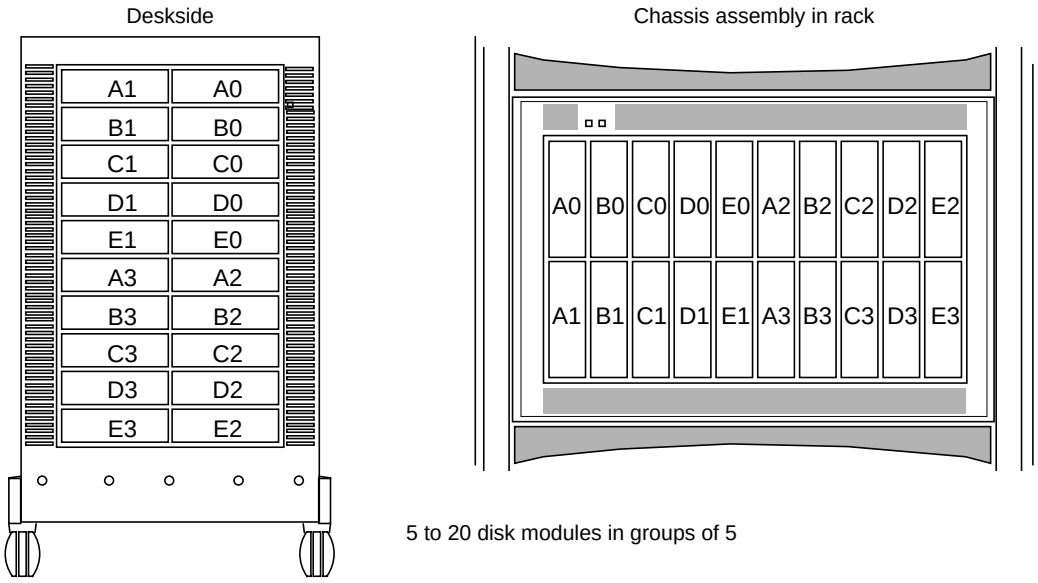

Figure 7-8 Location of Disks (Front of Challenge RAID)

Note: For information on replacing other FRUs (field-replaceable units) in the Challenge RAID storage system, see the *CHALLENGE RAID Installation and Maintenance Guide*. For information on replacing SCSI-2 adapters in the Challenge server, see the Challenge server documentation.

This section consists of the following:

- Section 7.6.1, “Opening and Closing the Fan Module”
- Section 7.6.2, “Replacing or Adding a Disk Module”
- Section 7.6.3, “Installing an Add-On Disk Module Array”

7.6.1 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 7-9.

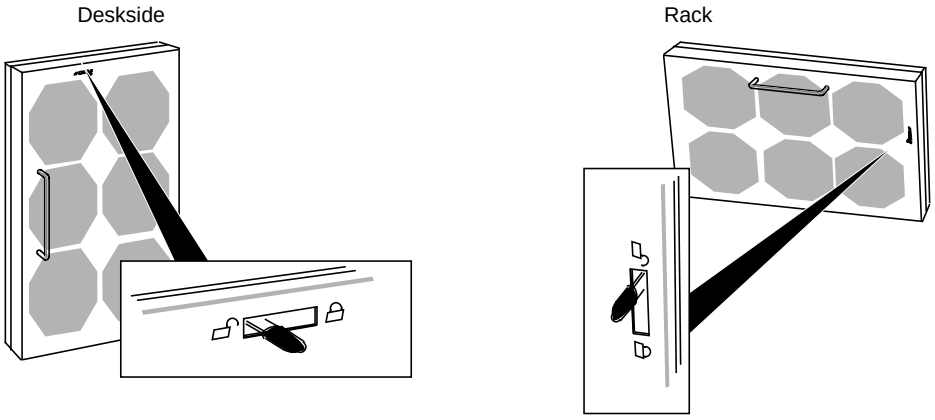

Figure 7-9 Unlocking the Fan Module

2. Swing open the fan module, as shown in Figure 7-10.

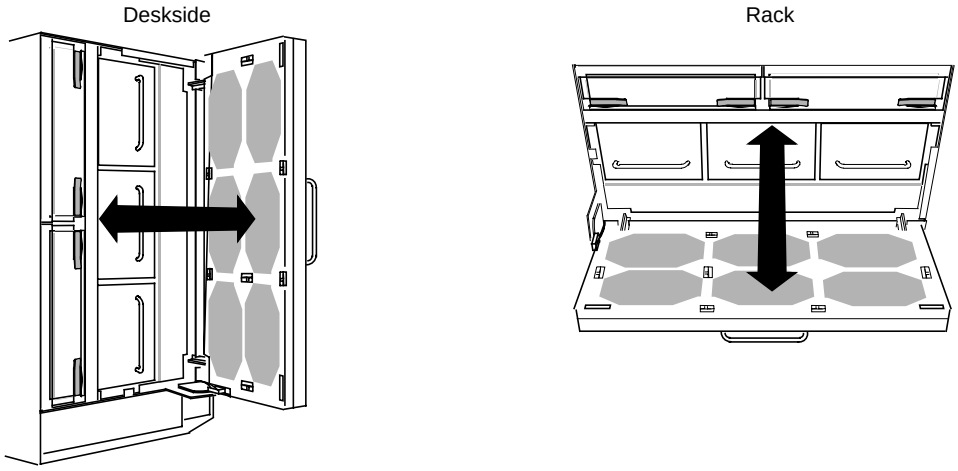

Figure 7-10 Opening the Fan Module

3. 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.

7.6.2 Replacing or Adding a Disk Module

This section consists of the following:

- Section 7.6.2.1, “Identifying and Verifying a Failed Disk Module”
- Section 7.6.2.2, “Removing a Failed Disk Module”
- Section 7.6.2.3, “Installing a Replacement Disk Module”
- Section 7.6.2.4, “Updating the Disk Module Firmware”

7.6.2.1 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, 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.

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

1. Look for the module with its amber fault light on. Figure 7-11 shows the fault indicator light and other lights on a disk module.

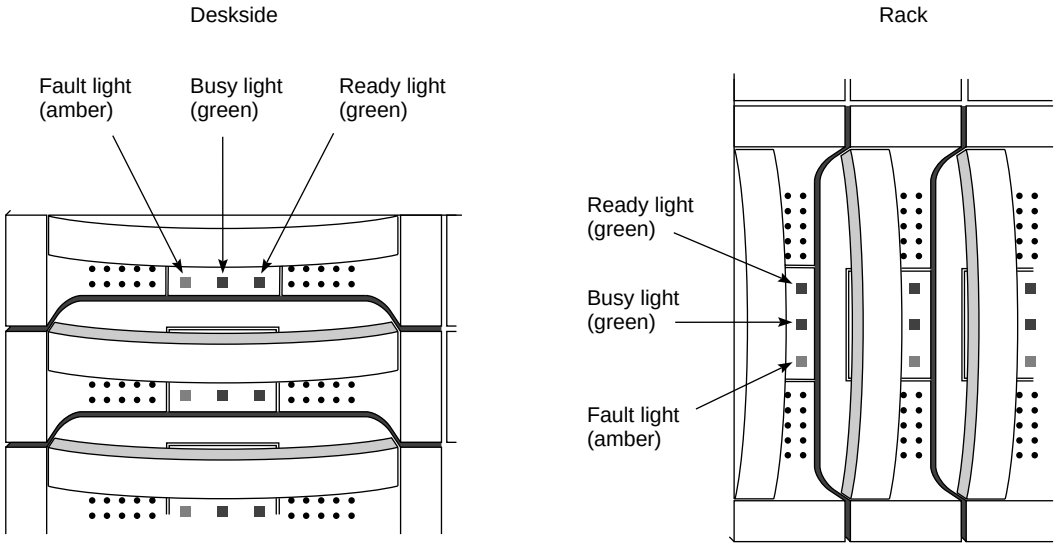

Figure 7-11 Disk Module Status Lights

2. Determine the failed module’s ID; see Figure 7-8.

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.
4. If you have not already checked the unsolicited error log with *raidcli getlog* for a message about the disk module, 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, which might mean the disk module itself has not failed.

   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 Chapter 7 of the *CHALLENGE RAID Owner's Guide*.

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.

7.6.2.2 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 storage-control processor 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; see Section 7.6.2.1, "Identifying and Verifying a Failed Disk Module."
2. When a disk fails, the storage-control processor 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*.
3. Locate the disk module that you want to remove; see Figure 7-8 if necessary.

4. Position the new disk module in its antistatic packaging within reach of the storage system.
5. 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 7-12. Put the wrist band around your wrist with the metal button against your skin.

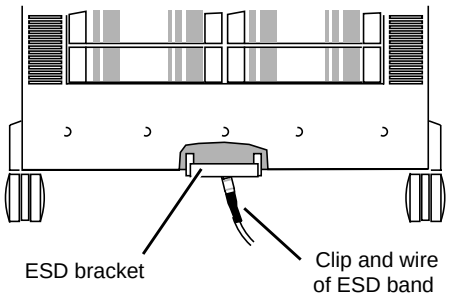

Figure 7-12 Attaching the ESD Clip to the ESD Bracket on the Storage System

Figure 7-13 shows where to attach the clip on a rack storage system.

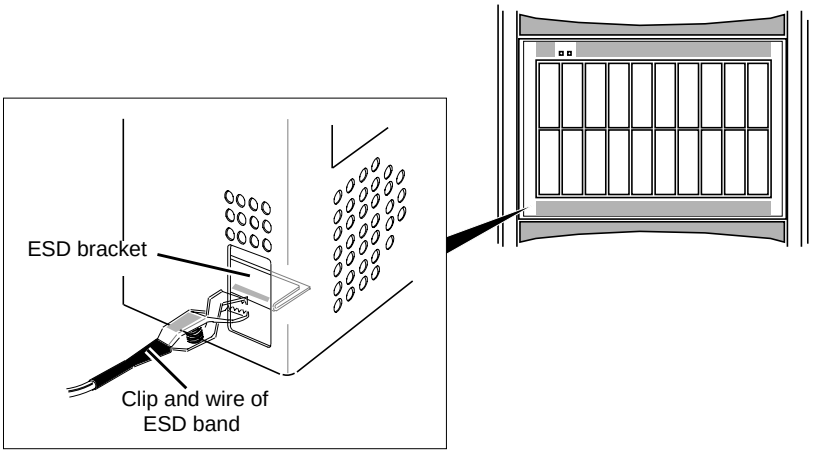

Figure 7-13 Attaching the ESD Clip to the ESD Bracket on a Rack Storage System

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

7. Grasp the disk module by its handle and pull it partway out of the cabinet, as shown in Figure 7-14.

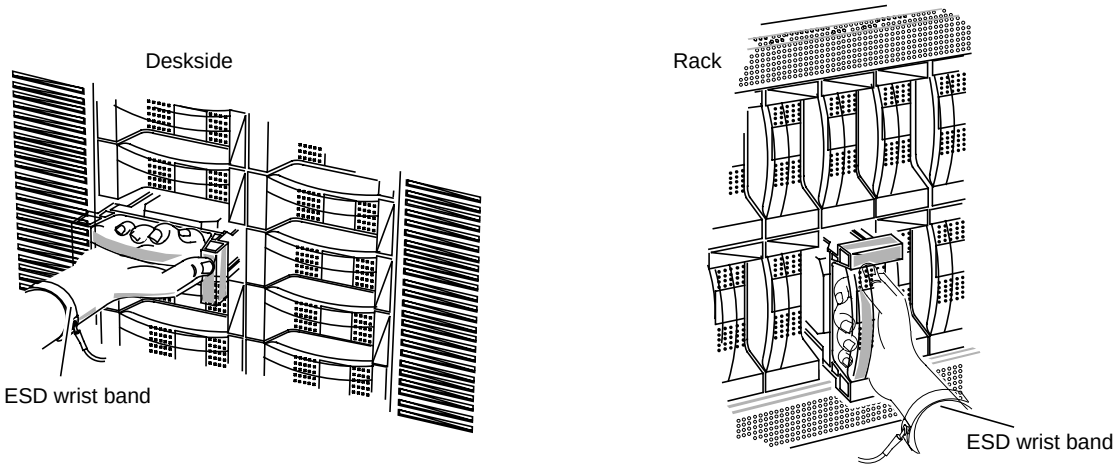

Figure 7-14 Pulling Out a Disk Module

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.**

8. Supporting the disk module with your free hand, pull it all the way out of the cabinet, as shown in Figure 7-15.

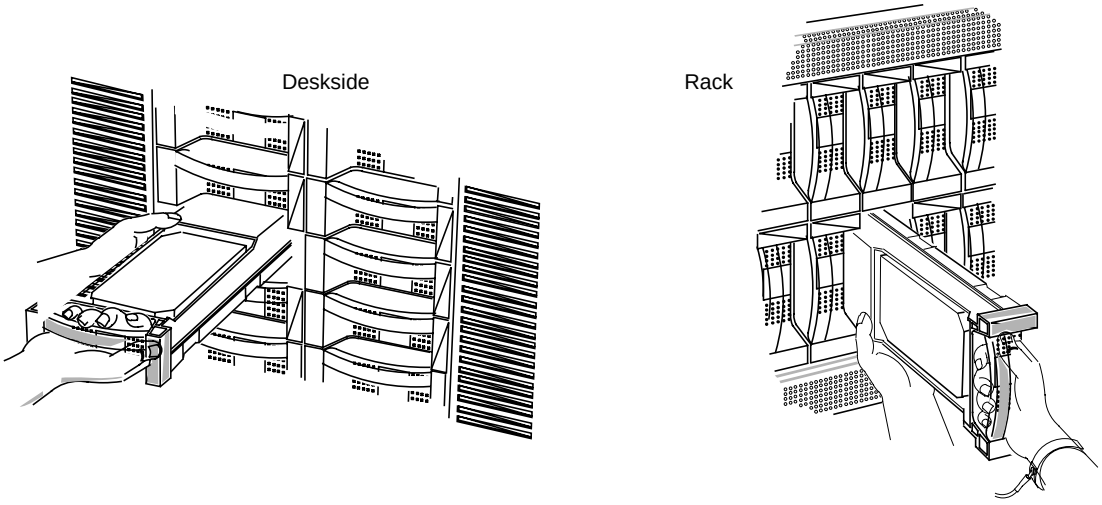

Figure 7-15 Removing a Disk Module

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

9. If necessary, on the label on the side of the disk module, write the ID number for the compartment from which you removed the drive; for example, A3.

   For the compartment ID numbers, refer to Figure 7-8.

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

7.6.2.3 Installing a Replacement Disk Module

To install the replacement disk module, follow these steps:

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

1. Before installing a replacement module, wait at least 15 seconds after removing the failed module to allow the processor time to recognize that the module has been removed. If you insert the replacement module too soon, the processor may report the replacement module as defective.
2. Position the new disk module in its antistatic packaging within reach of the storage system.
3. Locate the slot where you will install the disk module; see Figure 7-8.
4. Touch the new disk module's antistatic packaging to discharge it and the drive module. Remove the new disk module from its packaging.
5. Engage the disk module's rail in the chassis rail slot, as shown in Figure 7-16.

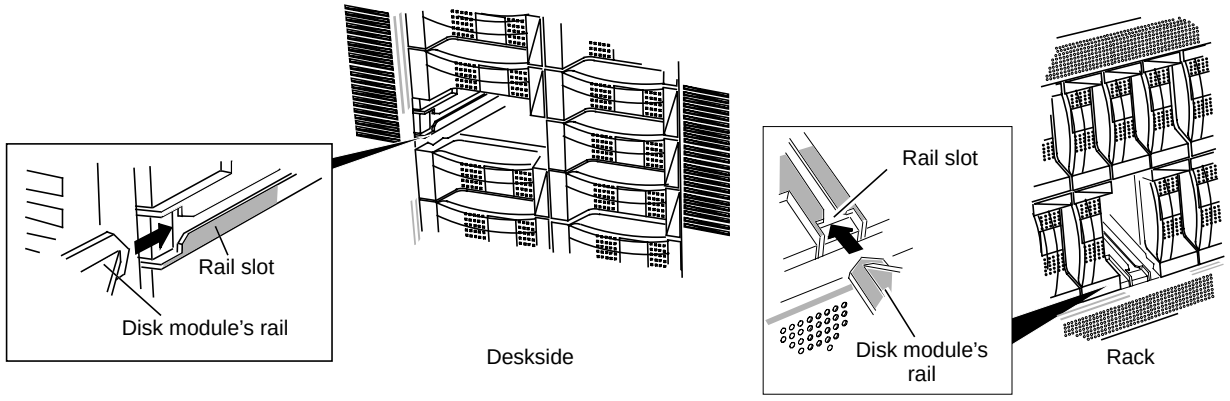

Figure 7-16 Engaging the Disk Module Rail

6. Engage the disk module's guide in the chassis guide slot, as shown in Figure 7-17.

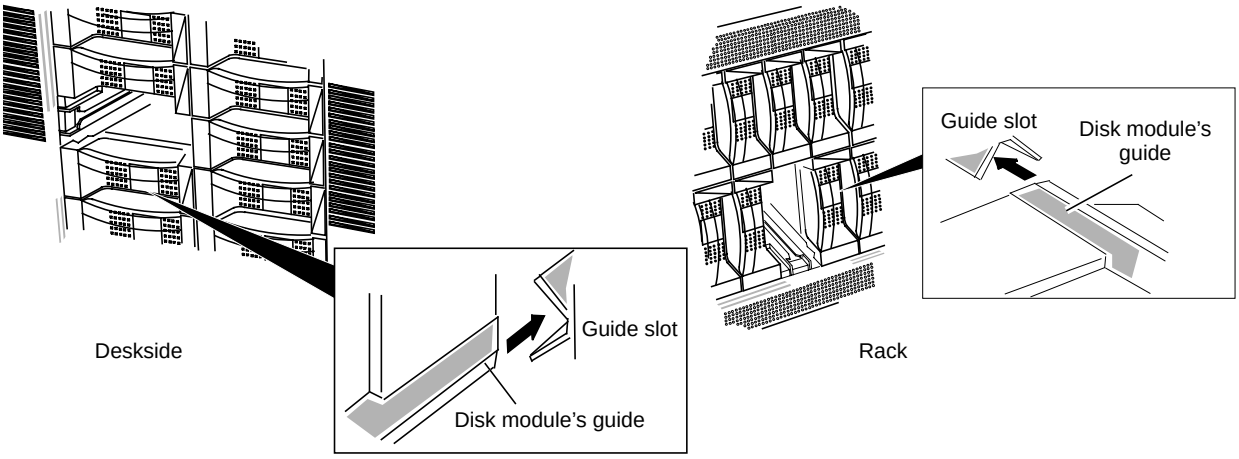

Figure 7-17 Engaging the Disk Module Guide

7. Insert the disk module, as shown in Figure 7-18. Make sure it is completely seated in the slot.

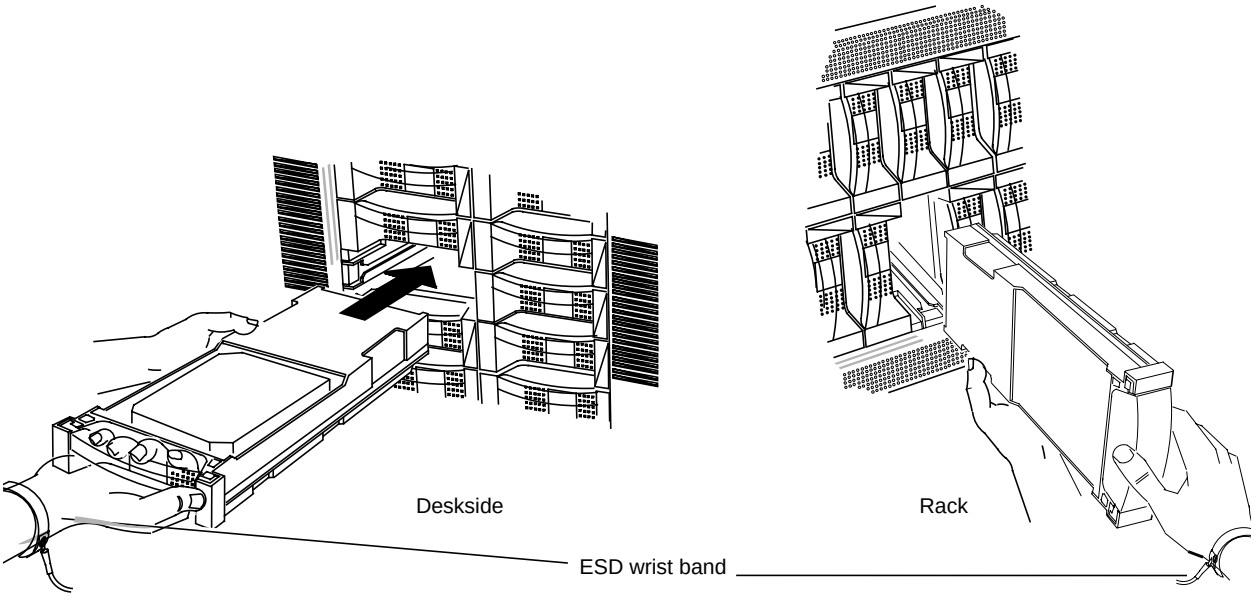

Figure 7-18 Inserting the Replacement Disk Module

8. Remove and store the ESD wrist band, if you are using one.
9. The storage-control processor 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 four 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 of the *CHALLENGE RAID Installation and Maintenance Instructions*.

7.6.2.4 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 separate subdirectories in */usr/raid5/flare*:

- */usr/raid5/flare/sauna8/flarecode.bin*: use this for AMD based SPs with FLARE code rev. level 8.00 through 8.49
- */usr/raid5/flare/phoenix8/flarecode.bin*: use this for PowerPC based SPs with FLARE code rev. level 8.50 through 8.99
- */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 *raidcli getagent* or *getsp* to return the new firmware version.

7.6.3 Installing an Add-On Disk Module Array

When you install an add-on disk module, use only Challenge RAID disk modules; only they contain the correct device firmware. Other disk modules, even those from other Silicon Graphics equipment, do 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) for use in other disk module positions.

Table 7-1 gives marketing codes for add-on disk module arrays. Add-on disk modules are available only in arrays of five.

Table 7-1 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

Fill the slots in this order (consult Figure 7-8 if necessary):

- 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

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

1. Position the new disk modules in their antistatic packaging within reach of the storage system.
2. If you are using a wrist band, attach its clip to the ESD bracket on the bottom of the storage system, as shown in Figure 7-12. Put the wrist band around your wrist with the metal button against your skin.
3. Locate the slots where you will install the add-on disk modules; see Figure 7-8.

4. 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. See Figure 7-19.

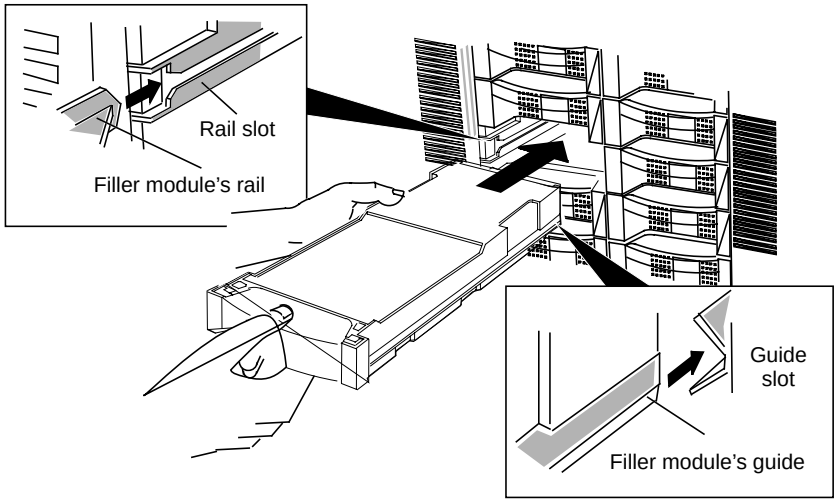

Figure 7-19 Removing a Disk Filler Module

5. Touch the new disk module's antistatic packaging to discharge it and the drive module. Remove the new disk module from its packaging.
6. Engage the disk module's rail in the chassis rail slot, as shown in Figure 7-20.

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

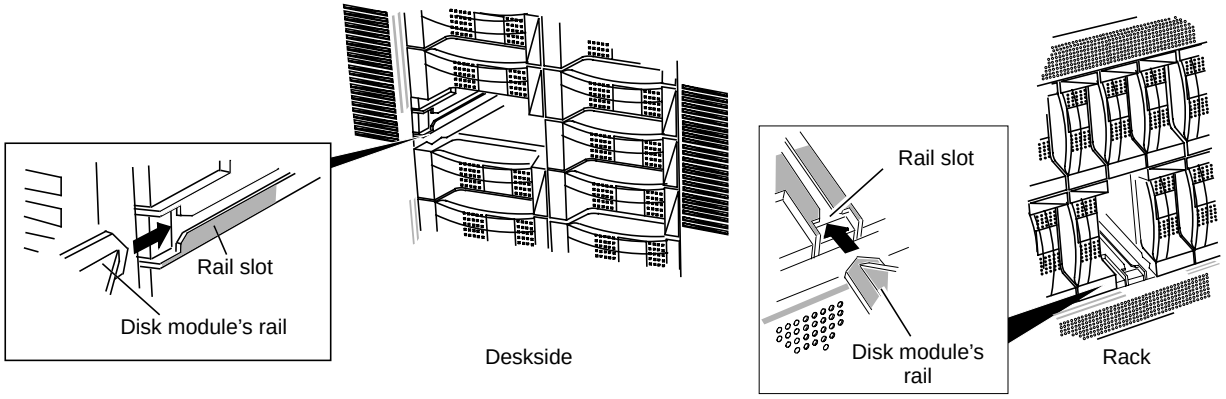

Figure 7-20 Engaging the Disk Module Rail

7. Engage the disk module's guide in the chassis guide slot, as shown in Figure 7-21.

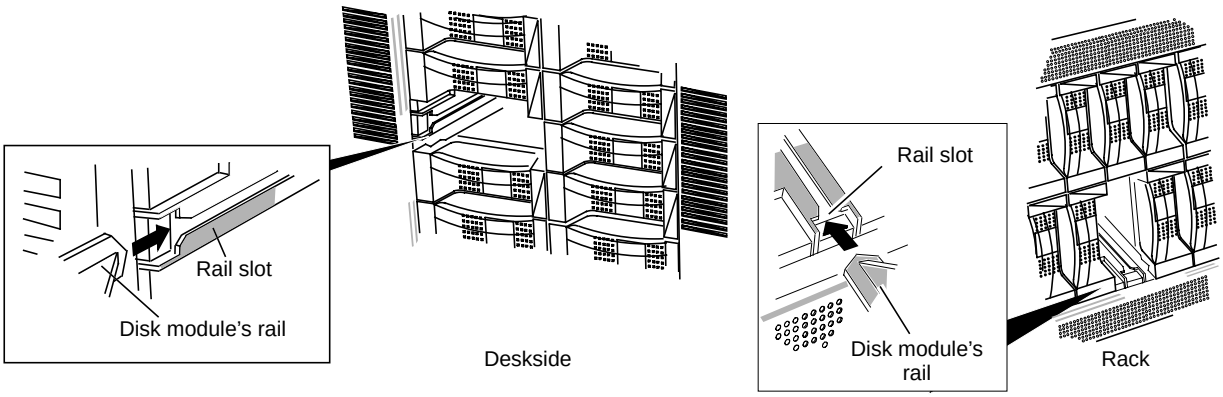

Figure 7-21 Engaging the Disk Module Guide

8. Insert the disk module, as shown in Figure 7-21. Make sure it is completely seated in the slot.

7.6.3.1 Updating the Disk Module Firmware

If you have replaced 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 separate directories in */usr/raid5/flare*:

- */usr/raid5/flare/sauna8/flarecode.bin*: use this for AMD based SPs with FLARE code rev. level 8.00 through 8.49
- */usr/raid5/flare/phoenix8/flarecode.bin*: use this for PowerPC based SPs with FLARE code rev. level 8.50 through 8.99
- */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 filename in the command contains microcode that runs on the storage-control processor and possibly also a microcode image for the storage-control processor PROM, which runs the power-on diagnostics. This command has no output.

7.7 Replacing Fibre Channel Disk Modules

For this procedure, have handy the *Origin FibreVault and Fibre Channel RAID Installation Instructions* (108-0154-002 or later).

Note the following:

- There is no need to power off the array enclosure while replacing disks in an array.
- When adding or replacing a fibre channel disk module, use only the proper Silicon Graphics FC-AL disk modules.

  The two types of Fibre Channel disk modules, RAID and non-RAID, look identical, but have different capacities; they are not interchangeable. Always confirm the type of fibre drive (RAID or non-RAID) before you install it. Each drive module has an identifying sticker with its part number:

  – 9470138 for (8.8 GB) RAID drives (520-byte sectors)

  – 9470140 for (9.1 GB) non-RAID drives (512-byte sectors)

  The two types of drives are not compatible and must not be mixed in the same enclosure.
- If you are replacing RAID disk modules 0, 1, or 2 (which contain internal code), or all the disk drive modules in a RAID LUN, you must follow instructions in "Updating Licensed Internal (flare) Code" in Appendix C of the *Origin FibreVault and Fibre Channel RAID Installation Instructions.*

  In this case, the SPs must be rebooted and the agent restarted after the new disk is in place. Rebooting restarts the SPs in the array, which terminates all outstanding I/O to the array. If you plan to reboot, you must unmount any filesystems or partitions on the array and quiesce the bus. Plan to work with the customer for this procedure.

Caution: The disk modules have no ESD shielding; do not stack them on top of each other. Use standard ESD precautions as listed in the instructions and always leave disks in an antistatic container until ready for installation.

To replace a Silicon Graphics Fibre Channel storage option disk module, follow these steps:

1. Have ready the proper replacement disk module(s), RAID or non-RAID.
2. Identify the failed disk drive module(s); not e the following:

   A failed fibre channel disk drive module's yellow LED lights, or both of its LEDs are off to show that the host is bypassing it.

   Although you can remove a disk module within a RAID-3 LUN without damaging the data on the LUN, you should do this only when the disk module has actually failed.

   If you are replacing more than one disk drive, use the GUI to determine the LUN(s) the drives are in so that you can avoid killing an operating LUN when you remove the drives.

   Read the event log for the SP that owns the LUN containing the faulty disk module. Also check for any other messages that indicate a related failure, such as a failure of a fibre channel loop connection or a general shutdown of an enclosure. Such a message could mean the disk module itself has not failed.

A message about the disk module contains its module ID. In the Fibre Channel RAID graphics user interface, if the disk module button displays the ID, in the Array Configuration dialog, select View, and then select Show Disk IDs.

3. Use one of the following methods to take the disk offline and prepare for installation:

   - Non-RAID disks: If necessary, back up the information on the disk you are replacing. To tell the host that you are removing a (non-RAID) disk module and taking it offline, use the `fccli REMOVE` option with the **-c** and **-t** arguments, which specify channel designator and target drive. This example takes drive 10 in channel 12 offline:

     ```
     fccli REMOVE -c 12 -t 10
     ```

   - RAID disks:

     Unbind the LUN containing the failed disk module, following instructions in Appendix B (unbind command) in the *Origin FibreVault and Fibre Channel RAID Installation Instructions.* On an Onyx2, or on an Origin2000 equipped with optional visualization console (SI) graphics, you can use the Graphical User Interface (described in Appendix C of the *Origin FibreVault and Fibre Channel RAID Installation Instructions*) to take a RAID disk module offline.

     In a RAID array, disks 0, 1, and 2 contain special internal code; see "Updating Licensed Internal (flare) Code" in Appendix C in the *Origin FibreVault and Fibre Channel RAID Installation Instructions.*

   - Either type of disk: Shut down the host and fibre channel enclosure systems, which terminates all outstanding I/O to the array. If you plan to reboot, which restarts the SPs in the array, work with the customer to unmount any filesystems or partitions on the array and quiesce the bus.

4. Open the FibreVault or Fibre Channel RAID enclosure's front door by inserting the key in the door's latch and turning the key to the unlocked icon position. Remove the key. Push the door's latch, and pull the door down until it is fully open.

   If you pull the door down with enough force that it becomes unhinged, snap it back on the hinges.

5. Attach the clip of the ESD wristband to the FibreVault or Fibre Channel RAID enclosure chassis and put the wristband around your wrist with the metal button against your skin.

   Caution: The enclosures have no mechanism to prevent you from removing a disk module while the drive is spinning. Do not remove a module unless its green active light is off and its amber check LED is on (or the entire enclosure is shut down).

6. Grasp the disk module's handle so that your thumb is on its latch. Push the latch and slowly pull the module from the slot, supporting it with your other hand, as shown in Figure 7-22. The disk module's amber check LED turns off (if power is on).

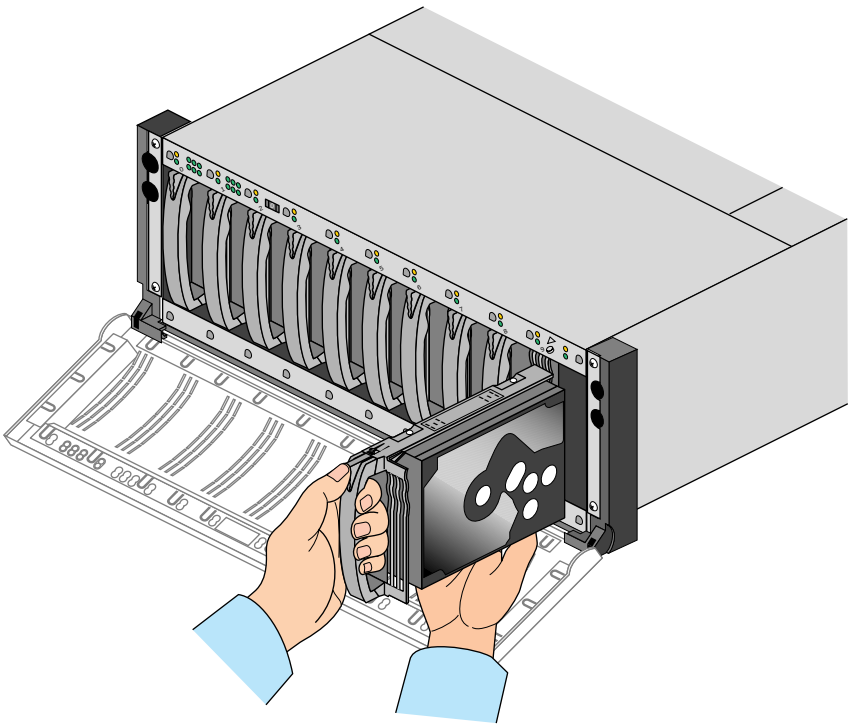

Figure 7-22 Removing an FC Disk Module

Note: Do not remove a disk module or filler unless you have its replacement ready to install. If you must replace all the disk drive modules in a RAID LUN, see Appendix C in the *Origin FibreVault and Fibre Channel RAID Installation Instructions*.

7. Grasp the handle of the new disk module and align the module with the guides in the empty slot. Slowly push the disk module into the slot until it clicks into place. The disk module's active LED flashes and then remains on (if power is on).
8. Remove and store the ESD wristband.
9. Power the host and FC enclosure back on if they were powered off during the replacement procedure.
10. Use one of the following procedures to bring the disk online:
    - Take the disk targeted for replacement online with the graphical user interface (see Appendix C in the *Origin FibreVault and Fibre Channel RAID Installation Instructions*).
    - Use the `fccli INSERT` option in the command-line interface to inform the host (that owns the loop) of its presence; see Appendix B in the *Origin FibreVault and Fibre Channel RAID Installation Instructions*. The fault LED remains lit until the drive is online.

- If the disk you have replaced has new firmware (RAID disks 0, 1, or 2), reboot; see "Updating Licensed Internal (flare) Code" in Appendix C in the *Origin FibreVault and Fibre Channel RAID Installation Instructions.*

11. Push the fibre channel enclosure's front door up until it latches into place.
12. Set the enclosure address, as explained in Section 2.10.4.3, "Setting the Enclosure Address on a Fibre Channel RAID Enclosure," or Section 2.10.5.3, "Setting Enclosure Addresses on the FibreVault," in Chapter 2.

7.8 Restoring LUN Ownership on a Challenge RAID System

If you have replaced a SCSI component (SPs, cables, SCIP mezzanine cards, SCSI boards, I/O boards), you have powered off the RAID system's SPs in the process. At this time, accesses to the LUNs currently owned by the disabled SP are normally no longer possible.

When the SPs are turned back on (at the end of the replacement procedure), all LUNs continue to be owned by the other SP. This section explains how to restore the original status:

- Section 7.8.1, "Checking XLV Volumes"
- Section 7.8.2, "Rebalancing LUNs Across Two SPs"

Note: Although IRIX 6.4 includes the failover feature, IRIX 6.2 requires a patch.

7.8.1 Checking XLV Volumes

If accesses to the LUNs on the RAID system are made through XLV volumes under IRIX 6.2, and an SP fails (or is turned off), XLV code forces the remaining SP to take over ownership of all LUNs and redirects the accesses to the remaining SP.

This failover is transparent to the user (such as the Oracle RDBMS) making the accesses, provided the following points are true:

- The RAID system is dual-SP.
- The access to the RAID is through XLV volumes.
- The volumes have been assembled correctly.
- Both SPs are operational at the time that one fails (or is turned off).
- The default ownership of the LUNs is distributed across the two SPs (this is not a strict requirement, but is a useful simplification).

7.8.1.1 Checking XLV Volumes Under IRIX 6.2

For IRIX 6.2 only, to check that all volumes have been assembled correctly, run *xlv_make*, and at the prompt, enter **show**. If each volume element of each volume shows for each disk partition a second one in curly braces ({}), the volume has been assembled correctly.

The most common cause for incorrectly assembled volumes (barring hardware failures) is that the RAID agent was running at the time *xlv_assemble* was invoked. If both SPs are still operational, this problem can be rectified as follows:

1. Stop the RAID agent:

   /etc/init.d/raid5 stop

2. Enter **xlv_assemble**.
3. View the output carefully to verify that assembly was done correctly.
4. Restart the agent:

   /etc/init.d/raid5 start

7.8.1.2 Checking XLV Volumes Under IRIX 6.4

For systems running under IRIX 6.4, use *scsifo* to check for dual connectivity to the RAID LUNs:

```
scsifo [-d (dump failover structures)]
        [-s <device_pathname> (perform failover switch)]
        [-P <device_pathname> [pathname is currently active]
```

For example:

```
# scsifo -d
Group 0:
  [P] /hw/scsi_ctlr/5/target/10/lun/3 (89)
  [ ] /hw/scsi_ctlr/7/target/6/lun/3 (227)
Group 1:
  [ ] /hw/scsi_ctlr/5/target/10/lun/2 (87)
  [P] /hw/scsi_ctlr/7/target/6/lun/2 (219)
Group 2:
  [P] /hw/scsi_ctlr/5/target/10/lun/1 (79)
  [ ] /hw/scsi_ctlr/7/target/6/lun/1 (217)
Group 3:
  [ ] /hw/scsi_ctlr/5/target/10/lun/0 (77)
  [P] /hw/scsi_ctlr/7/target/6/lun/0 (209)
```

7.8.2 Rebalancing LUNs Across Two SPs

When the original SP is turned back on (at the end of the replacement procedure), all LUNs continue to be owned by the other SP. While this situation is viable, it is not optimal. Performance is significantly improved if accesses to the LUNs are balanced across the two SPs, instead of all being funneled through one SP.

Follow the appropriate procedure below whenever an SP (storage-control processor) of a dual-SP RAID system is turned off, then turned back on again.

7.8.2.1 Rebalancing LUNs Under IRIX 6.2

To restore the original status on systems running IRIX 6.2, use the *raidcli trespass* command (RAID agent 1.55; for earlier versions, use */sbin/trespass*). 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. See Appendix B of the *CHALLENGE RAID Installation and Maintenance Instructions* (007-0128-006 or later) for more information on *raidcli trespass*; */sbin/trespass* has identical options.

7.8.2.2 Rebalancing LUNs Under IRIX 6.4

To restore the original status on systems running IRIX 6.4 (RAID agent 2.3 or later), use *scsifo* to get the status of the dual paths to the RAID LUNs, as shown in Section 7.8.1.2, "Checking XLV Volumes Under IRIX 6.4". Use this command to switch to the other path; for example:

```
scsifo -s /hw/scsi_ctlr/7/target/6/lun/0/disk
```

7.9 Replacing Batteries in the Remote Power Control Unit

The remote power control unit accepts inlet power from a standard 12 VDC-AC wall adapter.

The remote power control unit uses eight AA alkaline batteries as an on-board uninterruptable power supply to power the unit and sustain it for up to five hours during a power outage. The batteries are in a removable battery tray, which is accessed from the front of the unit.

An LED on the battery tray alerts you that the batteries have lost sufficient power to drive the remote power control unit. This LED also lights up immediately after you power on the remote power control unit.

If the batteries fail and are not replaced, the unit will retain all configuration settings in the absence of power for ten years, although it cannot control the power control units.

You can replace batteries in the remote power control unit with the unit powered on and the IRIS FailSafe cluster running (the battery tray is hot-pluggable).

To install fresh batteries, follow these steps:

1. Have ready eight fresh AA alkaline batteries. Do not mix fresh and used batteries.
2. Unlock the battery tray by turning the lock screw from the **LOCK** to the **UNLOCK** position. Pull out the battery tray using its handle.
3. Remove the three screws on the outer part of the battery holder cage; slide the holder cage off. The two inner battery holders are exposed; each holds four batteries.

4. With your fingers, carefully pry the used batteries out of one holder. Replace the batteries, following the correct orientation. (Use the other holder as a guide if necessary.) Repeat the process for the second holder.
5. When you have replaced the batteries, slide the outer battery holder cage back onto the battery tray, and replace the screws.
6. Replace the battery tray in the remote power control unit and turn the lock screw back to the **LOCK** position.

Index