

Oracle Parallel Server™

Installation and Maintenance Instructions

Document Number 108-0136-003

Contributors

Written by Carolyn Curtis

Illustrated by Cheri Brown, Dan Young, and Carolyn Curtis

Production by Linda Rae Sande

Engineering contributions by Paddy Sreenivasan, Raghu Mallena, Chander Kant, Neville Nandkeshwar, Ewan McKissock, Ravi Chhabria, and Dave Thompson

© 1995-1997, Silicon Graphics, Inc.— All Rights Reserved

This document contains proprietary and confidential information of Silicon Graphics, Inc. The contents of this document may not be disclosed to third parties, copied, or duplicated in any form, in whole or in part, without the prior written permission of Silicon Graphics, Inc.

Restricted Rights Legend

Use, duplication, or disclosure of the technical data contained in this document by the Government is subject to restrictions as set forth in subdivision (c) (1) (ii) of the Rights in Technical Data and Computer Software clause at DFARS 52.227-7013 and/or in similar or successor clauses in the FAR, or in the DOD or NASA FAR Supplement. Unpublished rights reserved under the Copyright Laws of the United States. Contractor/manufacturer is Silicon Graphics, Inc., 2011 N. Shoreline Blvd., Mountain View, CA 94043-1389.

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.

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 dépassant 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、第一種情報装置（商工業地域において使用されるべき情報装置）で商工業地域での電波障害防止を目的とした情報処理装置等電波障害自主規制協議会（VCCI）基準に適合しております。

従って、住宅地域またはその隣接した地域で使用すると、ラジオ、テレビジョン受信機等に受信障害を与えることがあります。

取扱説明書に従って正しい取り扱いをして下さい。

この装置は、第一種情報装置（商工業地域において使用されるべき情報装置）で商工業地域での電波障害防止を目的とした情報処理装置等電波障害自主規制協議会（VCCI）基準に適合しております。

従って、住宅地域またはその隣接した地域で使用すると、ラジオ、テレビジョン受信機等に受信障害を与えることがあります。

取扱説明書に従って正しい取り扱いをして下さい。

Silicon Graphics, CHALLENGE, IRIS, Indy, and the Silicon Graphics logo are registered trademarks and FDDIXPress, Indigo², IRISconsole, IRIS FailSafe, IRIX, IRIX Networker, IRIXPro, Onyx2, Origin, Origin200, Origin2000, O2, Performance Co-Pilot, Ultra SCSI XIO, and XFS are trademarks of Silicon Graphics, Inc. Cray is a registered trademark and CrayLink is a trademark of Cray Research, Inc. ORACLE is a registered trademark and OPS and Oracle Parallel Server are trademarks of Oracle Corporation. PostScript is a trademark of Adobe Systems Incorporated.

Oracle Parallel Server™ Installation and Maintenance Instructions Document Number 108-0136-003

Silicon Graphics, Inc.
Mountain View, California

Contents

About This Guide.....	xiii
Structure of This Document	xiv
Other Required Documentation	xiv
Conventions	xv
1. Oracle Parallel Server Components	1-1
2. Setting Up and Cabling Installations With Origin2000 Servers	2-1
2.1 Setting Up the Component Systems.....	2-2
2.2 Installing SCSI Option Boards.....	2-3
2.3 Planning the Cabling	2-3
2.4 Cabling the Vaults to the OPS Hosts.....	2-5
2.5 Setting the OPS Host SCSI IDs	2-9
2.6 Installing the Network Boards in the OPS Hosts	2-10
2.7 Cabling the Challenge RAID Storage System to the Servers.....	2-11
2.7.1 Selecting a SCSI ID Switch Setting for the Challenge RAID Storage System	2-11
2.7.2 Cabling the Challenge RAID Storage System and the Servers	2-12
2.8 Cabling the IRISconsole Multiplexer to the Workstation	2-17
2.9 Cabling the IRISconsole Multiplexer and the OPS Hosts	2-19
2.10 Powering On the Systems.....	2-26
2.11 Testing the Installed OPS Hardware.....	2-28
3. Setting Up and Cabling Installations With Challenge Servers.....	3-1
3.1 Setting Up the Component Systems.....	3-2
3.2 Upgrading the Remote System Control Firmware	3-4
3.3 Upgrading the IO4 Board Mezzanine Cards	3-4
3.4 Planning the Cabling	3-4
3.5 Cabling the Vaults to the OPS Hosts.....	3-8
3.6 Setting the OPS Host SCSI IDs.....	3-11
3.7 Installing the Network Boards in the OPS Hosts	3-12

3.8	Cabling the Challenge RAID Storage System to the Challenge Servers.....	3-13
3.9	Cabling the IRISconsole Multiplexer to the Workstation	3-18
3.10	Cabling the IRISconsole Multiplexer and the Challenge Servers	3-20
3.11	Powering On the Systems.....	3-25
3.12	Testing the Installed OPS Hardware.....	3-26
4.	Installing the Oracle Parallel Server Software.....	4-1
4.1	Checking the OPS Host SCSI IDs	4-1
4.2	Installing System Software	4-2
4.2.1	Installing IRIX and the Patches on the OPS Hosts	4-2
4.2.2	Converting to the XFS Filesystem.....	4-2
4.2.3	Installing the opscm Software on the OPS Hosts	4-3
4.2.4	Installing the opsnc Software and Configuration File on the Workstation	4-5
4.2.5	Testing the Serial Connection to the OPS Hosts.....	4-7
4.3	Installing the IRISconsole Software and Configuring the Site	4-8
4.4	Installing and Configuring the Challenge RAID Storage System Software for OPS	4-9
4.4.1	Creating the Challenge RAID Config File for Use With OPS	4-10
4.4.2	Restarting the Agent and Checking the Config File	4-10
4.4.3	Creating SCSI Device Nodes (IRIX 6.2 Only).....	4-11
4.4.4	Configuring LUNs	4-12
4.5	Enabling Command-Tagged Queuing on the LUNs.....	4-13
4.6	Partitioning Disks and Creating the XLV Volumes.....	4-14
4.7	Installing the Oracle Database Software.....	4-15
4.8	Installing the OPS Software.....	4-15
4.9	Starting OPS.....	4-17
5.	Maintaining and Upgrading the Oracle Parallel Server.....	5-1
5.1	Isolating a Node	5-1
5.1.1	Breaking the SCSI Chain With the Shared RAID	5-2
5.1.2	Breaking the SCSI Chain With the Log Disks	5-3
5.2	Replacing Cables	5-5
5.2.1	Replacing a Cable on the SCSI Chain With the RAID	5-5
5.2.2	Replacing Network Cables or Serial Lines	5-5
5.2.3	Replacing a Cable on the SCSI Chain With the Log Disks.....	5-6
5.3	Replacing a Log Disk.....	5-7
5.4	Checking OPS Server and Challenge RAID SCSI IDs	5-8
5.5	Checking Disk Partitions	5-9

5.6	Exchanging Challenge RAID Disk Modules.....	5-10
5.6.1	Opening and Closing the Fan Module.....	5-12
5.6.2	Replacing or Adding a Disk Module	5-13
5.6.3	Installing an Add-On Disk Module Array	5-20
5.7	Restoring LUN Ownership on a Challenge RAID System	5-21
5.7.1	Checking XLV Volumes.....	5-21
5.7.2	Rebalancing LUNs Across Two SPs.....	5-22
Index		Index-1

Figures

Figure 2-1	SCSI Bus Selectors on Front of Vault.....	2-3
Figure 2-2	Example OPS Cabling Scheme.....	2-4
Figure 2-3	Connecting SCSI Buses A and B From the First Vault to the First OPS Host (Origin)	2-6
Figure 2-4	Cabling the First Vault to Both Servers	2-7
Figure 2-5	OPS Cabling for Origin Servers and Vaults.....	2-8
Figure 2-6	SCSI-2 Bus Connectors on Back of Challenge RAID Chassis.....	2-13
Figure 2-7	Connecting a SCSI Bus Cable to a Challenge RAID SCSI Port.....	2-13
Figure 2-8	Connecting a Challenge RAID SCSI Bus Cable to a SCSI Port.....	2-14
Figure 2-9	OPS Cabling for Origin Servers and Challenge RAID	2-15
Figure 2-10	Challenge RAID OPS Connection Schemes	2-16
Figure 2-11	Connecting the Multiplexer's SCSI Cable to an O2 Workstation	2-17
Figure 2-12	Setting the Multiplexer SCSI ID.....	2-18
Figure 2-13	Origin2000 Deskside Serial Connector	2-19
Figure 2-14	Origin Rack MMSC ALTERNATE CONSOLE Port	2-20
Figure 2-15	Origin2000 Server tty_1 Port (Console Port).....	2-21
Figure 2-16	Multiplexer Cabling, First OPS Host (Origin2000 Deskside)	2-22
Figure 2-17	Multiplexer Cabling, First OPS Host (Origin2000 Rackmount)	2-23
Figure 2-18	Multiplexer Cabling, Both OPS Hosts (Origin2000 Deskside)	2-24
Figure 2-19	Multiplexer Cabling, Both OPS Hosts (Origin2000 Rackmount)	2-25
Figure 2-20	Connecting the Power Cord (Back of Challenge RAID Storage System)	2-26
Figure 2-21	Turning On the Storage System Power	2-26
Figure 2-22	Unlocking the Fan Module	2-27
Figure 2-23	Enabling a Storage-Control Processor's Power	2-27
Figure 3-1	SCSI Bus Selectors on Front of Vault.....	3-3
Figure 3-2	SCSI Controller Numbering on IO4 and Mezzanine Boards.....	3-5
Figure 3-3	Example OPS Cabling Scheme.....	3-6
Figure 3-4	OPS Cabling for Servers and Vaults.....	3-7
Figure 3-5	Connecting the First Challenge Vault and Server.....	3-8
Figure 3-6	Connecting SCSI Buses A and B from the First Vault to the First Server	3-9

Figure 3-7	Cabling the First Vault to Both Servers	3-10
Figure 3-8	SCSI-2 Bus Connectors on Back of Challenge RAID Chassis	3-14
Figure 3-9	Connecting a SCSI Bus Cable to a Challenge RAID SCSI Port.....	3-14
Figure 3-10	Connecting a Challenge RAID SCSI Bus Cable to a Challenge SCSI Port.....	3-15
Figure 3-11	Challenge RAID OPS Connection Schemes	3-16
Figure 3-12	Connecting SCSI Cables to a Challenge RAID and Both OPS Hosts.....	3-17
Figure 3-13	Cabling the Multiplexer to an Indy Workstation	3-18
Figure 3-14	Connecting the Multiplexer's SCSI Cable to an O2 Workstation	3-19
Figure 3-15	Setting the Multiplexer SCSI ID	3-20
Figure 3-16	Cabling the OPS Host Remote System Control Connector.....	3-21
Figure 3-17	Cabling the OPS Host System Console Connector	3-22
Figure 3-18	Multiplexer Cabling for Challenge XL Server	3-23
Figure 3-19	IRISconsole Cabling for OPS (Challenge L Servers)	3-24
Figure 3-20	Connecting the Power Cord (Back of Challenge RAID Storage System)	3-25
Figure 3-21	Turning On the Storage System Power	3-25
Figure 3-22	Unlocking the Fan Module	3-26
Figure 3-23	Enabling a Storage-Control Processor's Power	3-26
Figure 3-24	Example hinv Output for Installed OPS System (Challenge)	3-27
Figure 4-1	Change Site Config Dialog Box.....	4-9
Figure 5-1	Data Flow Between OPS Challenge Systems and Challenge RAID	5-2
Figure 5-2	Halting Data Flow from SP A to the OPS Systems	5-2
Figure 5-3	Terminating SP A.....	5-2
Figure 5-4	Data Flow with SP A On and Terminated.....	5-3
Figure 5-5	Terminating SP B.....	5-3
Figure 5-6	Disconnecting Cables from Challenge Vault to Host 2	5-4
Figure 5-7	Location of Disks (Front of Challenge RAID).....	5-11
Figure 5-8	Unlocking the Fan Module	5-12
Figure 5-9	Opening the Fan Module	5-12
Figure 5-10	Disk Module Status Lights	5-13
Figure 5-11	Attaching the ESD Clip to the ESD Bracket on the Storage System	5-15
Figure 5-12	Attaching the ESD Clip to the ESD Bracket on a Rack Storage System	5-15
Figure 5-13	Pulling Out a Disk Module.....	5-16
Figure 5-14	Removing a Disk Module	5-16
Figure 5-15	Engaging the Disk Module Rail.....	5-17
Figure 5-16	Engaging the Disk Module Guide	5-18
Figure 5-17	Inserting the Replacement Disk Module	5-18

Tables

Table 1-1	OPS Hardware Bundle Components, Origin2000-Only Configurations.....	1-1
Table 1-2	OPS Hardware Bundle Components, Challenge-Only Configurations.....	1-2
Table 1-3	IRISconsole Administration Software Kit Contents	1-2
Table 2-1	Challenge RAID SCSI ID Switch Settings.....	2-11
Table 3-1	Challenge RAID SCSI ID Switch Settings.....	3-13
Table 4-1	Example Shared Volume Names	4-15
Table 5-1	Challenge RAID SCSI ID Switch Settings.....	5-9
Table 5-2	Ordering Add-On Disk Module Sets	5-20

About This Guide

The Oracle Parallel Server™ (OPS) represents a significant development in RDBMS capability for ORACLE® users. The ability to cluster two computer systems and coordinate their access to a single shared ORACLE database provides for increased throughput with higher levels of application and database availability. Using a CHALLENGE® RAID storage system can provide high availability.

The Silicon Graphics® OPS consists of the following hardware components:

- two servers: Origin™ or CHALLENGE, running IRIX™ 6.2 or 6.4

Note: For ease in reading, CHALLENGE is written as Challenge in the balance of this guide. For brevity, the term IRISconsole™ is used regardless of how many multiplexers are used to connect to systems at a site.

- optionally, for each server, an external peripheral enclosure for SCSI storage devices, such as a Challenge Vault
- one IRISconsole™: O2™ workstation or 32-bit uniprocessor Indy® workstation
- one IRISconsole serial port multiplexer, cables, and software
- one Challenge RAID deskside storage system with two storage-control processors (SPs) and at least five disk modules
- required hardware upgrades and cables

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

OPS software runs on systems running IRIX 6.2 or 6.4. The software making up the OPS configuration includes IRIX patches for OPS, Silicon Graphics Oracle Parallel Server software, and software for the component systems, such as IRISconsole and Challenge RAID storage system software. Optional software includes Performance Co-Pilot™ (PCP), IRIXPro™, and Database Accelerator (DBA). The customer also obtains ORACLE RDBMS and OPS software from Oracle Corp.

The Oracle Parallel Server software version required is 7.3.2 or later, 64-bit version.

Structure of This Document

This guide contains the following chapters and appendices:

- Chapter 1, “Oracle Parallel Server Components,” lists the systems, cables, and other components included in OPS.
- Chapter 2, “Setting Up and Cabling Installations With Origin2000 Servers,” describes how to install and set up servers and vaults, the Challenge RAID storage system, and the IRISconsole on Origin family-only configurations and how to cable them for OPS.
- Chapter 3, “Setting Up and Cabling Installations With Challenge Servers,” describes how to install and set up servers and vaults, the Challenge RAID storage system, and the IRISconsole on Challenge-only configurations and how to cable them for OPS.
- Chapter 4, “Installing the Oracle Parallel Server Software,” explains how to install the various software components of OPS.
- Chapter 5, “Maintaining and Upgrading the Oracle Parallel Server,” explains procedures for swapping log disks, swapping disks on the Challenge RAID, and other tasks.

An index completes this guide.

Other Required Documentation

The following Silicon Graphics documents are required for installing the OPS. 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™ version from *guest@comrade.engr:/usr/people/guest/docdist*.

- *Site Preparation for Origin Family and Onyx2* (007-3452-00x)
- *Origin2000 and Onyx2 Deskside and Rackmount Installation Instructions* (108-0155-002 or later)
- *Origin200 and Origin Vault Installation Instructions* (108-0153-001)
- *Origin2000 Rackmount Owner's Guide* (007-3456-00x)
- *Origin2000 Deskside Server Owner's Guide* (007-3453-00x)
- *Origin200 Owner's Guide* (007-3415-00x)
- *Ultra SCSI XIO Board Installation Instructions* (108-0157-001)
- *O2 Workstation Hardware Reference Manual* (007-3275-xxx)
- *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)
- *FDDI PCI Option Installation Instructions* (007-3534-001)
- *100-BASE-T VME Installation Instructions* (108-0148-001)

- *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)
- *FDDIXpress Mezzanine Board for Challenge and Onyx Installation Instructions* (108-0116-002)
- *FDDIXpress Administration Guide* (007-0813-040)
- *FDDIXpress Release Notes*, included in FDDI board shipment
- *CHALLENGE RAID Installation and Maintenance Instructions* (108-0128-006)
- *System Controller PROM Upgrade for Oracle Parallel Server* (108-0140-001)
- *CHALLENGE Vault Rack and SCSIBox 2 Installation Instructions* (108-7044-040)
- *CHALLENGE Vault L Installation Instructions* (108-0124-001)
- *IRIX Admin: Disks and Filesystems* volume, which covers XLV volumes for IRIX 6.2 (007-2825-00x)
- *IRISconsole Administrator's Guide*, included with the IRISconsole CD-ROM (007-2872-00x)
- *IRISconsole Multiplexer Installation Guide* (007-2839-003)
- *The IRIXpro Administrator's Guide* (007-2446-00x)

In addition to the manuals listed above, the latest *CHALLENGE RAID Owner's Guide* (007-2532-007 or later) is useful for maintaining or troubleshooting OPS; bring them to a site requiring maintenance.

The customer manual shipped with the Oracle Parallel Server is called the *Oracle Parallel Server Administrator's Guide for Silicon Graphics Systems* (007-2763-003).

Conventions

The following 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

Oracle Parallel Server Components

This chapter lists OPS components.

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

The OPS option includes:

- Oracle Parallel Server hardware bundle
- Oracle Parallel Server software bundle
- IRISconsole Multiplexer (hardware and software)
- if necessary, XIO or PCI Fast Ethernet boards for the servers
- cables

Table 1-1 lists the contents of the OPS hardware bundle for Origin-only configurations (OPS-Origin2000™-1.2).

Table 1-1 OPS Hardware Bundle Components, Origin2000-Only Configurations

Component	Quantity	Part Number or Marketing Code
IRISconsole kit (see contents in Table 1-3)	1	C0-IRISCONS-1.2
Ethernet cable assembly RJ45-RJ45 (null modem), 40-foot	1	9290132
Cable package: two SCSI cables for connecting to storage, 20-foot	1	X-2SCSI-CABLES
Cable package: two serial cables for connecting deskside Origin/Onyx2™ to ST-1616 multiplexer, 25-foot	1	XVIC25-O2000-DS
Cable package: two serial cables for connecting rackmount Origin/Onyx2 to ST-1616 multiplexer, 25-foot	1	XVIC25-O2000-RM
CD pack, OPS 1.2 for Challenge L and Origin2000	1	SC4-OPS-1.2

Table 1-2 lists the contents of the OPS hardware bundle for Challenge-only configurations (OPS-Challenge-1.2).

Table 1-2 OPS Hardware Bundle Components, Challenge-Only Configurations

Component	Quantity	Part Number or Marketing Code
IRISconsole kit (see contents in Table 1-3)	1	C0-IRISCONS-1.2
Ethernet cable assembly RJ45-RJ45 (null modem), 40-foot	1	9290132
Cable package: two SCSI 20-foot for connecting to storage	1	X-2SCSI-CABLES
Cable package: two serial 25-foot for connecting to ST-1616 multiplexer	1	XVIC25
CD pack, OPS 1.2 for Challenge L and Origin2000	1	SC4-OPS-1.2

Table 1-3 lists components of the IRISconsole kit.

Table 1-3 IRISconsole Administration Software Kit Contents

Component	Quantity	Part Number
16-port Serial Port Multiplexer	1	ST-1616
Cable assembly, 68P/50P high-density SCSI, 0.5-meter	1	018-0689-001
CD pack: IRISconsole software CD 1.2	1	813-0571-002
<i>IRISconsole Multiplexer Installation Guide</i> (customer documentation, hard copy)	1	007-2839-003

Optional hardware includes

- server (one or two)
- Challenge Vault storage
- Challenge RAID storage system(s)
- tape drive(s)
- O2 workstation for use with IRISconsole

Optional software includes latest versions of

- Performance Co-Pilot
- IRIX NetWorker™
- IRIXPro
- Database Accelerator

Note: From Oracle Corporation, the customer also gets ORACLE CD-ROM with software for OPS (including the Distributed Lock Manager software) and for Oracle RDBMS, and documentation.

Chapter 2

Setting Up and Cabling Installations With Origin2000 Servers

Note: Before installing an Oracle Parallel Server (OPS), make sure that the installation site meets the operating limits and AC power requirements as explained in *Site Preparation for Origin Family and Onyx2* (007-3452-00x).

This chapter explains how to cable the two Origin2000 deskside or rackmount servers (OPS hosts), the two Challenge vaults, the Challenge RAID storage system, and the IRISconsole for use as the hardware for the Oracle Parallel Server.

Note: A multirack Origin server should be treated as one OPS server; that is, the CrayLink™ interconnect cannot be used for communication between the two OPS servers.

The process consists of

1. setting up the component systems
2. planning the cabling
3. cabling the Challenge vaults to the OPS hosts
4. setting the OPS host SCSI IDs
5. cabling the Challenge RAID storage system to the servers
6. cabling the IRISconsole multiplexer to the workstation
7. cabling the IRISconsole multiplexer and the servers
8. powering on the systems
9. testing the installed OPS hardware

The following equipment is required for installation:

- manuals:
 - *Site Preparation for Origin Family and Onyx2*
 - installation guides for the component systems, including storage (see the introduction to this guide for part numbers)
 - *IRIS 4-Port Fast Ethernet Adapter with Asynchronous Serial XIO Board Installation Instructions* (108-0151-001) or *Fast Ethernet PCI Option Installation Instructions* (007-3535-001)

- *IRIX Admin: Disks and Filesystems* volume (007-2825-00X)
- laptop or ASCII terminal
- Phillips-head and small flat-blade screwdrivers

2.1 Setting Up the Component Systems

Read through the site preparation guide(s) before unpacking equipment at the site, paying particular attention to the following:

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

Plan to situate the OPS component systems fairly close together. Differential SCSI cables, including cabling inside the chassis, can be no longer than 25 meters (82 feet). Use only the cables included in the shipment.

4. Make sure the site meets safety and operating considerations, as explained in the site preparation guide(s). Note that the fully loaded Challenge RAID rack weighs 899.4 lb (407.96 kg); make sure flooring is adequate.
5. Prepare the site for the systems' electrical requirements.
6. Connect your laptop or ASCII terminal and keyboard to serial port on the back of the first server, as explained in the server installation instructions.
7. If the system includes Challenge vaults, reserve any label sheets that might be included; you need them for labeling SCSI cables for OPS.
8. If the system includes Challenge vaults, set the SCSI bus selectors on each vault to 8 through 15. Figure 2-1 shows the location of the settings.




Figure 2-1 SCSI Bus Selectors on Front of Vault

9. Connect the power cord for the first OPS host (OPS1). 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.
10. Power up each OPS host. When power-on diagnostics are completed, a message appears.
11. Run *hinv* and check for shared disk entries in the output.
12. Set up each system on the network for the site.

2.2 Installing SCSI Option Boards

To install optional Ultra SCSI XIO™ boards (Origin2000 servers), follow instructions in the appropriate guide.

2.3 Planning the Cabling

For Origin2000 servers, you must use the SCSI ports on the optional Ultra SCSI XIO boards, which are differential. Each board has four ports; labels should be attached to each end of each Y cable giving the module, XIO slot, and SCSI channel numbers (for example, **MOD 1 IO6 SCSI0**).

Note: For OPS and other dual hosting, you can use only Ultra SCSI ports 1, 2, and 3; you cannot use port 0.

In an OPS configuration, each SCSI bus runs from a controller on one OPS host through a storage system (vault or RAID) to the second OPS host. For clarity, you connect the controllers with the same numbers on a bus; for example, connect cable SCSI controller 2

on the first OPS host (OPS1) to SCSI controller 2 on the second OPS host (OPS2). Figure 2-2 diagrams a typical OPS configuration.

Note: Make sure that controllers 2 and 5 and bus A on both optional vaults are differential.

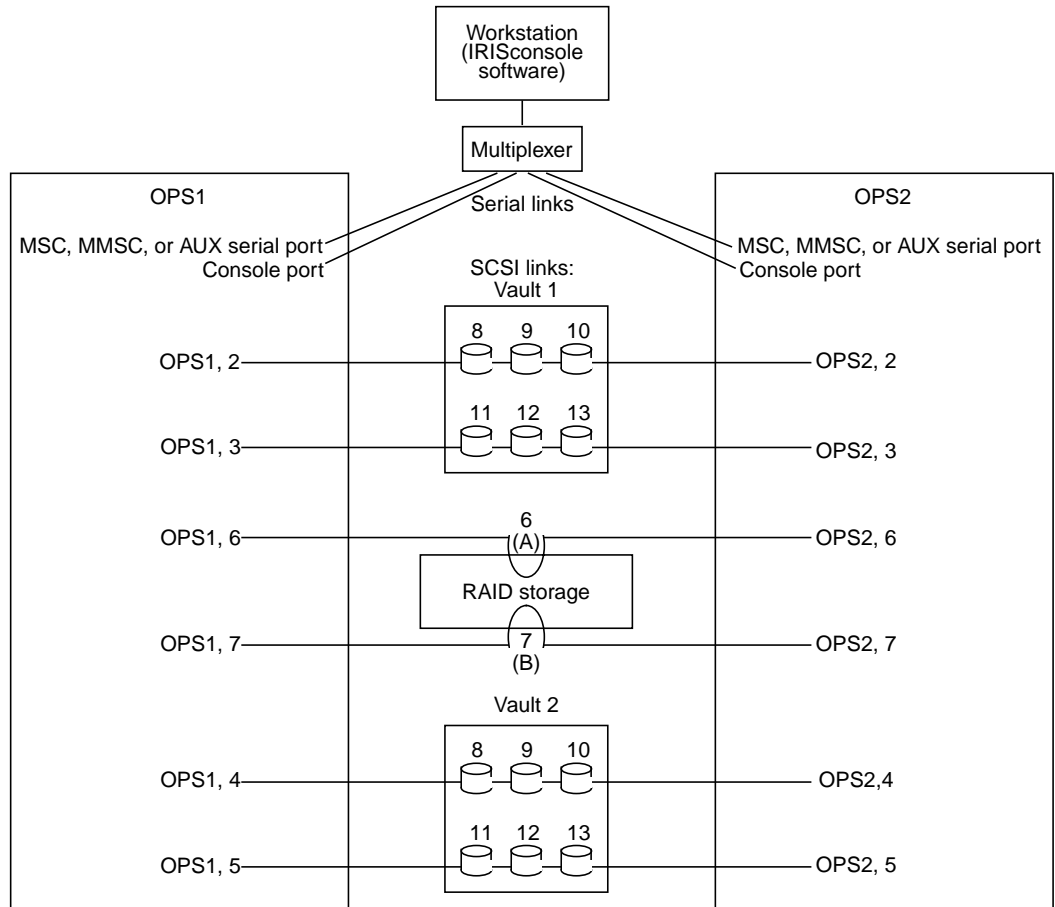


Figure 2-2 Example OPS Cabling Scheme

2.4 Cabling the Vaults to the OPS Hosts

This section explains how to cable each SCSI bus on each vault to each OPS host. Note the following:

- The SCSI buses in the vaults 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's 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 2-3. 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 vault, nor do the SCSI cables shipped with the OPS 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* (007-3499-001) or *Ultra SCSI XIO Board Installation Instructions* (108-0157-001).

Follow these steps:

1. Have ready eight of the OPS host-vault SCSI cables (9290111 is shipped with the OPS option). Use only these cables.
2. Power off both servers and vaults.
3. If necessary, insert an Ultra SCSI Y cable into the connector for ports 2 and 3 on each OPS host. 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 blank labels, create labels **Vault 1** and **Vault 2** and attach them to the appropriate vaults. If necessary, create labels for the hosts.
5. On the back of vault 1, insert a cable into the top SCSI channel A socket, as shown in Figure 2-3.




Figure 2-3 Connecting SCSI Buses A and B From the First Vault to the First OPS Host (Origin)

6. Label the cable end or connector; follow the scheme in Figure 2-4.
7. Connect the other end of this cable to the Ultra SCSI cable coming from port 2 (such as **MOD 1 IO6 SCSI2**), as shown in Figure 2-3. Label the connector or cable end.
8. On the back of vault 1, insert a cable into the top SCSI channel B socket. Label the cable end.
9. Connect the other end of this cable to the Ultra SCSI cable coming from port3 (such as **MOD 1 IO6 SCSI3**) on OPS1, as shown in Figure 2-3. Label the cable end or connector.
10. Insert a third cable into the lower SCSI channel A port on the back of vault 1. Label the cable end.
11. Connect the other end of this cable to Ultra SCSI port 2 on OPS2 (such as **MOD 2 IO6 SCSI2**) that you used on OPS1, as shown in Figure 2-4. Label the cable end.




Figure 2-4 Cabling the First Vault to Both Servers

12. Insert a fourth cable into the lower SCSI channel B port on the back of vault 1 and label it. Connect its other end to same odd-numbered SCSI bus port on OPS2 (such as **MOD 2IO6 SCSI3**) that you used on OPS1. Label the cable end. Figure 2-4 shows all cables for the first vault installed.
13. Connect the four remaining cables to the second vault; follow the pattern in steps 3 through 12. Figure 2-5 shows Origin servers and vaults cabled for OPS.




Figure 2-5 OPS Cabling for Origin Servers and Vaults

14. If necessary, reset the SCSI bus selector settings on the front of each Challenge vault if they have been unset accidentally.
15. If one or more optional digital tape drives were ordered, install them in the server(s).
16. Power on the vaults. Power on the servers.
17. Once the disks are set up, do a dummy I/O to test each link. For example, in the numbering scheme used so far, to test SCSI bus 2, send data from OPS1 to disk 8 on SCSI bus 2 (dks2d8) and see if the disk light for disk 8 on the first vault lights up. Send data from OPS2 and see if the same LED is activated. Repeat this process for each bus.

2.5 Setting the OPS Host SCSI IDs

SCSI host address conflicts have caused the most problems in getting OPS up and running.

Caution: The two servers must have different SCSI IDs. The recommended addresses for OPS are 0 for the first OPS host (OPS1) and 2 for the second OPS host (OPS2). No SCSI device on any bus on any server should have SCSI ID 2.

Setting the host SCSI IDs consists of these steps:

1. For the first OPS host, enter the command `nvrnm`. The last line of this command's output should contain the line

```
scsihostid=
```

This output means that the SCSI host ID for this OPS host is 0, the default.

2. Repeat step 2 in the second window for the second OPS host. The last line of the output should be `scsihostid=2`, meaning that the SCSI host ID for the second OPS host is 2.

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

1. Reboot and start 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
```

Caution: No SCSI device on any bus on any 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`; see Section 2.1, "Setting Up the Component Systems," for a sample output.
6. Enter the command `hinv` for the second server. Compare the output for each server.

2.6 Installing the Network Boards in the OPS Hosts

Follow these steps:

1. Following instructions in their manuals, install XIO or PCI Fast Ethernet or FDDI boards in the hosts.

Caution: This equipment is 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.

2. Edit `/etc/config/netif.options` to include network addresses, for example:

```

# To override the name and/or address of the first gateway interface,
# change the value part and remove the leading : character.

: if2name=ef0
: if2addr=dlm-$HOSTNAME

# If this host has more than 2 interfaces, you must define values for
# if3name (and if4name if appropriate). Change if3addr (and if4addr) to
# the appropriate names in /etc/hosts if your site has different naming
# conventions.

if3name=
if3addr=gate2-$HOSTNAME

if4name=
if4addr=gate3-$HOSTNAME

```

3. Edit */etc/hosts* to include IP numbers for OPS1 and OPS2. Use test numbers that are not in use on the customer's network. For example:

```

127.0.0.1      localhost
#
# This network number, '192.0.2' is the officially blessed 'test' network.
#   It should be changed immediately to your own official network number.

# start of well known machines--do not remove this line
# main ethernet network
192.26.79.38   ops1
192.26.79.39   ops2
# for DLM FDDI connection only
192.0.2.200    dlm-ops1
192.0.2.201    dlm-ops2
#

```

4. Check the network connections between servers by running ping(1M). Use netstat(1M) to see if the interfaces are displayed.

2.7 Cabling the Challenge RAID Storage System to the Servers

This section explains

- selecting a SCSI ID switch setting for the Challenge RAID system
- cabling the Challenge RAID storage system and the servers

2.7.1 Selecting a SCSI ID Switch Setting for the Challenge RAID Storage System

The SCSI ports on the Challenge RAID storage system are cabled to both OPS hosts; consult the diagram in Figure 2-2.

Table 2-1 charts SCSI ID switch settings for the Challenge RAID storage system. Note that the settings do not conform to frequently used numbering schemes.

Table 2-1 Challenge RAID SCSI ID Switch Settings

SCSI ID Number	Switch Number				Comment
	ID 0	ID 1	ID 2	ID 3	
1	Off	On	On	On	
2	On	Off	On	On	Do not use
3	Off	Off	On	On	
4	On	On	Off	On	
5	Off	On	Off	On	
6	On	Off	Off	On	Shown in Figure 2-9 as A
7	Off	Off	Off	On	Shown in Figure 2-9 as B
8	On	On	On	Off	
9	Off	On	On	Off	
10	On	Off	On	Off	
11	Off	Off	On	Off	
12	On	On	Off	Off	
13	Off	On	Off	Off	
14	On	Off	Off	Off	
15	Off	Off	Off	Off	

2.7.2 Cabling the Challenge RAID Storage System and the Servers

Note the following:

- The SCSI buses in the Challenge RAID system 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's 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 2-8. 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.

- 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 OPS option connect directly to the Ultra SCSI board; you must connect the SCSI cable to the Y cable to make the connection.

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

1. Have ready two SCSI cables shipped with Challenge RAID and the two 20-foot SCSI cables included in the Challenge RAID shipment (p/n 9290111).
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 each OPS host. 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. Using blank labels, create two labels that say **DF SCSI n** and two labels that say **DF SCSI N**, where **n** and **N** are the numbers of two remaining SCSI controllers on the servers (in the example in Figure 2-2, these are 6 and 7).
5. Attach the two even-numbered labels below the **SCSI A** in symbol and **SCSI A out** symbol silk-screens on the Challenge RAID storage system. Attach the two odd-numbered labels below the **SCSI B** in symbol and **SCSI B out** symbol on the Challenge RAID storage system. Figure 2-6 shows the Challenge RAID SCSI ports.

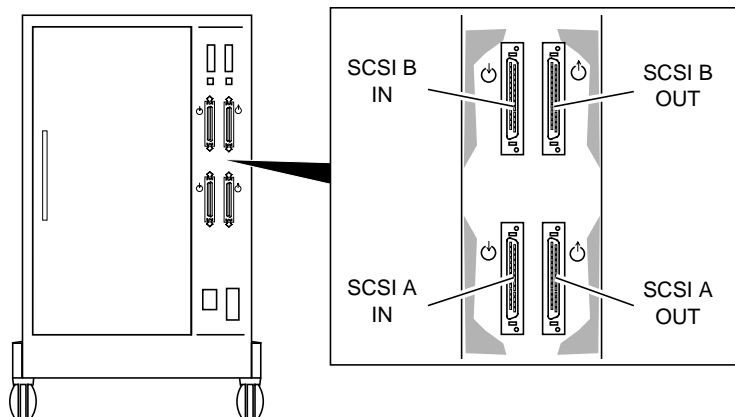


Figure 2-6 SCSI-2 Bus Connectors on Back of Challenge RAID Chassis

6. On the back of the Challenge RAID storage system, attach a SCSI cable to the Challenge RAID **SCSI A** in port. Connect the SCSI cable to the **SCSI A** in connector as shown in Figure 2-7, making sure the connector is inserted securely.




Figure 2-7 Connecting a SCSI Bus Cable to a Challenge RAID SCSI Port

7. Connect the other end of this cable to Ultra SCSI port 2 on OPS1 (for example, **MOD 1 IO6 SCSI2**).

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.

8. Attach another SCSI cable to the Challenge RAID **SCSI B** in port. Connect the other end to Ultra SCSI port 3 on OPS1 (for example, **MOD 1 IO6 SCSI3**), as shown in Figure 2-8.

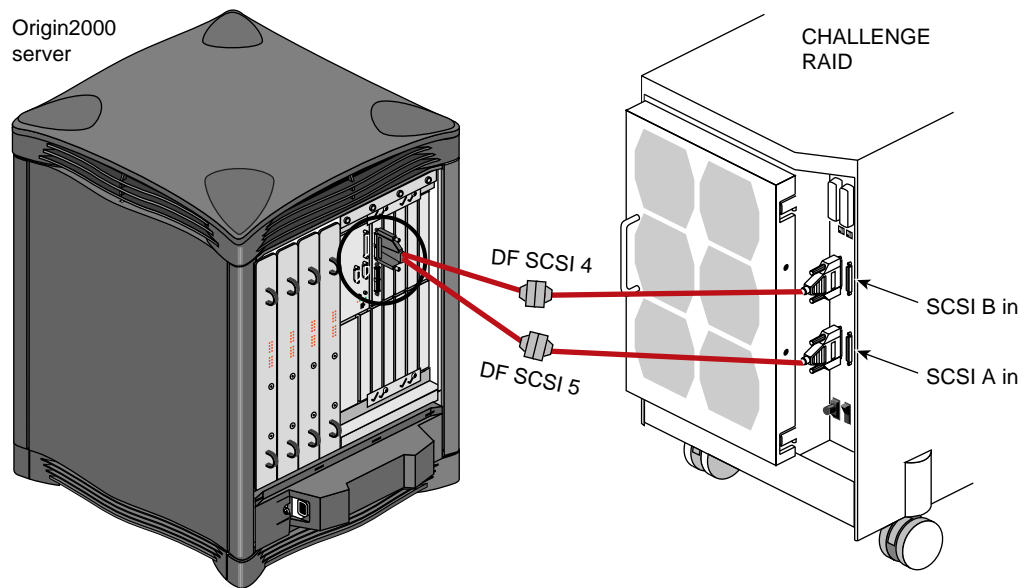


Figure 2-8 Connecting a Challenge RAID SCSI Bus Cable to a SCSI Port

9. Cable the Challenge RAID **SCSI B** in port to the second server's Ultra SCSI port 2. Figure 2-9 shows Origin servers and one Challenge RAID storage system cabled for OPS.




Figure 2-9 OPS Cabling for Origin Servers and Challenge RAID

For installations with more than one RAID system, consult Figure 2-10.

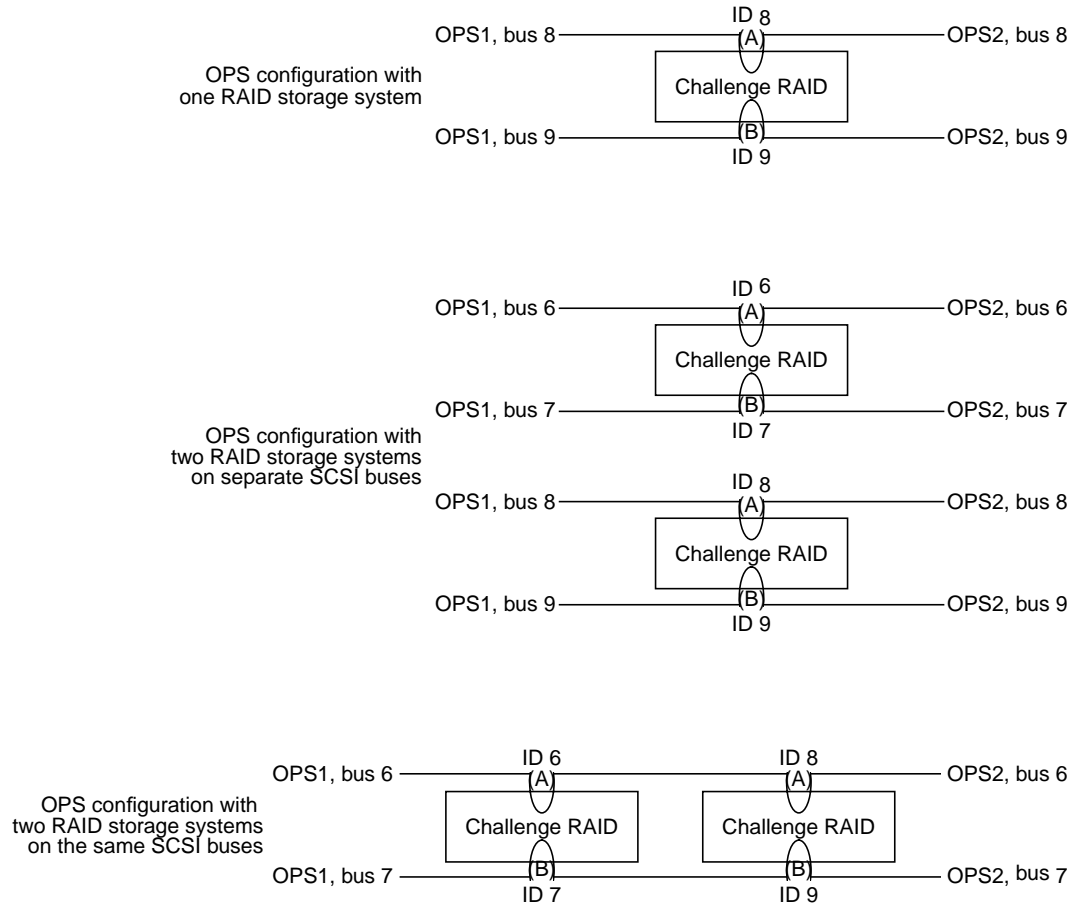


Figure 2-10 Challenge RAID OPS Connection Schemes

2.8 Cabling the IRISconsole Multiplexer to the Workstation

To cable the O2 workstation to the servers, follow these steps:

1. Have ready the contents of the IRISconsole multiplexer kit and the four 2-ft DB25-DB9 serial cables included in the cable kit.
2. Connect one end of the SCSI cable included in the multiplexer kit to the workstation's SCSI port, as shown in Figure 2-11.




Figure 2-11 Connecting the Multiplexer's SCSI Cable to an O2 Workstation

3. Attach the other end of the cable to the leftmost socket on the back of the multiplexer, as shown in Figure 2-11 (for example).
4. If the customer has purchased two multiplexers, attach the cable of the second multiplexer to the middle socket on the first multiplexer.
5. Cable the power supply to the rightmost port on the multiplexer, as shown in Figure 2-12.
6. Select an appropriate SCSI ID number by inserting a small screwdriver or other tool and rotating the arrow, as shown in Figure 2-12. Using SCSI ID 7 on early models of the multiplexer causes the power indication LED to blink. Select another ID and the LED glows steadily with no blink.

Note: The system disk in the O2 workstation is always SCSI target ID 1. Never select SCSI ID 1 for a multiplexer.




Figure 2-12 Setting the Multiplexer SCSI ID

7. If the multiplexer is the last SCSI device on the workstation's standard external SCSI port, terminate it properly by attaching the external SCSI terminator included in the multiplexer shipment. See Figure 2-12.
8. Reboot the IRISconsole workstation.

2.9 Cabling the IRISconsole Multiplexer and the OPS Hosts

The IRISconsole multiplexer is connected to the system controller port on the Origin server.

- Origin2000 deskside: the 9-pin sub-D serial port on the rear of the system (the same electrically as the 8-pin mini-DIN serial port on the rear of the Origin2000 server), as shown in Figure 2-13.




Figure 2-13 Origin2000 Deskside Serial Connector

- Origin2000 rackmount: **ALTERNATE CONSOLE** port on the Origin Rack MMSC; see Figure 2-14.




Figure 2-14 Origin Rack MMSC ALTERNATE CONSOLE Port

To cable the multiplexer to the Origin hosts, follow these steps:

1. Attach a 25-pin end of a serial cable to serial port 1 (or another port, as the customer wishes) on the top of the multiplexer.
2. Connect the 9-pin connector from serial port 1 on the multiplexer to the appropriate serial port of OPS1:
 - Origin2000 deskside: 9-pin sub-D serial port on the 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.

 - Origin2000 rackmount: **ALTERNATE CONSOLE** port on the Origin Rack MMSC
3. Attach a 25-pin end of a serial cable to serial port 2 (or another port, as the customer wishes) on the top of the multiplexer. Connect the other end of this cable to the tty_1 serial port of OPS1: the right-hand tty port (console), as shown in Figure 2-15.




Figure 2-15 Origin2000 Server tty_1 Port (Console Port)

Figure 2-16 shows this connection for an Origin2000 deskside server.




Figure 2-16 Multiplexer Cabling, First OPS Host (Origin2000 Deskside)

Figure 2-17 shows this connection for an Origin2000 rackmount server.




Figure 2-17 Multiplexer Cabling, First OPS Host (Origin2000 Rackmount)

4. Repeat steps 1 through 3 for the second OPS host: attach serial cables from ports 3 and 4 on the multiplexer to the OPS2's system serial port and console ports.

Figure 2-18 diagrams IRISconsole multiplexer cabling for Origin2000 deskside servers.




Figure 2-18 Multiplexer Cabling, Both OPS Hosts (Origin2000 Deskside)

Figure 2-19 diagrams IRISconsole multiplexer cabling for Origin2000 rackmount servers.




Figure 2-19 Multiplexer Cabling, Both OPS Hosts (Origin2000 Rackmount)

2.10 Powering On the Systems

To power on the component systems, follow these steps:

1. Make sure that the Challenge RAID power switch is off.
2. Connect the power cord to the storage system. Plug the other end of the power cord into an AC power outlet. See Figure 2-20.




Figure 2-20 Connecting the Power Cord (Back of Challenge RAID Storage System)

3. Turn on the storage system's power. The green power light on the front of the Challenge RAID storage system comes on, and the fans rotate. See Figure 2-21.




Figure 2-21 Turning On the Storage System Power

4. If the busy light on *none* of the drive modules lights up, make sure that the power for each storage-control processor is enabled. Move the fan module's latch to the **UNLOCK** position, as indicated in Figure 2-22.




Figure 2-22 Unlocking the Fan Module

5. Swing open the fan module and move the power switch for each disabled storage-control processor to the enable position, as shown in Figure 2-23.




Figure 2-23 Enabling a Storage-Control Processor's Power

6. Swing the fan module closed, and then move the latch to the **LOCK** position.

Caution: Do not leave the fan module open for more than two minutes at any time. The Challenge RAID storage system shuts down if the fan module is disabled for more than two minutes.
7. If necessary, connect the power cables and power on the OPS hosts, the Challenge RAID, the workstation, and the multiplexer.

2.11 Testing the Installed OPS Hardware

To test the OPS installation, run *hinv* on a server. Following is a sample output:

```
FPU: MIPS R10010 Floating Point Chip Revision: 0.0
CPU: MIPS R10000 Processor Chip Revision: 2.6
4 180 MHZ IP27 Processors
```

```

Main memory size: 256 Mbytes
Instruction cache size: 32 Kbytes
Data cache size: 32 Kbytes
Secondary unified instruction/data cache size: 1 Mbyte
Integral SCSI controller 0: Version QL1040B
  Disk drive: unit 1 on SCSI controller 0
  CDROM: unit 6 on SCSI controller 0
Integral SCSI controller 1: Version QL1040B
Integral SCSI controller 2: Version QL1040B
Integral SCSI controller 3: Version QL1040B
Integral SCSI controller 4: Version QL1040B          -----
  RAID controller: unit 1 on SCSI controller 4
  RAID lun: unit 1, lun 0 on SCSI controller 4
  RAID lun: unit 1, lun 1 on SCSI controller 4          RAID SP B
  RAID lun: unit 1, lun 2 on SCSI controller 4
  RAID lun: unit 1, lun 3 on SCSI controller 4          -----
Integral SCSI controller 5: Version QL1040B
  RAID controller: unit 1 on SCSI controller 5          -----
  RAID lun: unit 1, lun 0 on SCSI controller 5
  RAID lun: unit 1, lun 1 on SCSI controller 5          RAID SP A
  RAID lun: unit 1, lun 2 on SCSI controller 5
  RAID lun: unit 1, lun 3 on SCSI controller 5          -----
IOC3 serial port: tty1
IOC3 serial port: tty2
Integral Fast Ethernet: ef0, version 1
IOC3 external interrupts: 1

```


Setting Up and Cabling Installations With Challenge Servers

Note: Before installing an Oracle Parallel Server (OPS), 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*.

This chapter explains how to cable the two Challenge servers (OPS hosts), the two Challenge vaults, the Challenge RAID storage system, and the IRISconsole for use as the hardware for the Oracle Parallel Server. The process consists of

1. setting up the component systems
2. upgrading the Remote System Control firmware
3. upgrading the IO4 board mezzanine cards
4. planning the cabling
5. cabling the vaults to the OPS hosts
6. setting the OPS host SCSI IDs
7. installing the FDDI boards in the OPS hosts
8. cabling the Challenge RAID storage system to the Challenge servers
9. cabling the IRISconsole multiplexer to the workstation
10. cabling the IRISconsole multiplexer and the Challenge servers
11. powering on the systems
12. testing the installed OPS hardware

The following equipment is required for installation:

- installation guides for the component systems (see the introduction to this guide for part numbers)
 - *CHALLENGE/Onyx Site Preparation Guide* (108-7040-040)
 - *CHALLENGE/Onyx XL Rackmount Installation Instructions* (108-7042-020)
 - *CHALLENGE/Onyx Deskside Installation Instructions* (108-7039-020)
 - installation guides for network cards
 - installation guides for storage systems
 - *IRIX Admin: Disks and Filesystems* volume (007-2825-001)
 - *Getting Started with XFS Filesystems* (007-2549-001)

- laptop or ASCII terminal
- Phillips-head and small flat-blade screwdrivers

Note: Make sure that each Challenge L server used for OPS 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 *System Controller PROM Upgrade for Oracle Parallel Server* (108-0140-001), included in this shipment.

3.1 Setting Up the Component Systems

Read through the *CHALLENGE/Onyx Site Preparation Guide* before unpacking equipment at the site, paying particular attention to the following:

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*.

Plan to situate the OPS component systems fairly close together. Differential SCSI cables, including cabling inside the chassis, can be no longer than 25 meters (82 feet). Use only the cables included in the shipment.

4. Make sure the site meets safety and operating considerations, as explained in Chapter 3 of the *CHALLENGE/Onyx Site Preparation Guide*. Note that the fully loaded Challenge RAID rack weighs 899.4 lb (407.96 kg); make sure flooring is adequate.
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. Reserve any label sheets that might be included with the Challenge vault systems; you need them for labeling SCSI cables for OPS.
8. Set the SCSI bus selectors on each optional vault to 8 through 15. Figure 3-1 shows the location of the settings.




Figure 3-1 SCSI Bus Selectors on Front of Vault

9. Connect the power cord for the first OPS host (OPS1). 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.
10. On the OPS host, turn the System Controller key switch to the On position. When power-on diagnostics are completed, a message appears.
11. Type `hinv` on each OPS host; the screen should display output like the following (some lines omitted for clarity):

```

6 150 MHZ IP19 Processors
...
Integral Ethernet controller: et0, Ebus slot 5
FDDI ... : fxp0
EPC external interrupts
Integral SCSI controller 1: Version WD33C95A, differential, revision 0
Disk drive: unit 1 on SCSI controller 1
Integral SCSI controller 0: Version WD33C95A, single ended, revision 0
Tape drive: unit 6 on SCSI controller 0: DAT
Integral SCSI controller 7: Version SCIP/WD33C95A
Integral SCSI controller 6: Version SCIP/WD33C95A
Integral SCSI controller 5: Version SCIP/WD33C95A
Integral SCSI controller 4: Version SCIP/WD33C95A
Integral SCSI controller 3: Version SCIP/WD33C95A
Integral SCSI controller 2: Version SCIP/WD33C95A
Integral EPC parallel port: Ebus slot 5
VME bus: adapter 0 mapped to adapter 21
VME bus: adapter 21
  
```

12. Set up each system on the network for the site.

3.2 Upgrading the Remote System Control Firmware

If necessary, install the Remote System Control port upgrade in one or both Challenge L systems, following instructions in the *System Controller PROM Upgrade for Oracle Parallel Server* (108-0140-001) included in this shipment.

A new system controller firmware PROM is included in the OPS shipment. Install it as explained in the *System Controller PROM Upgrade for Oracle Parallel Server*.

3.3 Upgrading the IO4 Board Mezzanine Cards

Two SCIP mezzanine cards are included in the OPS option. Install them following instructions in the *CHALLENGE/Onyx XL Rackmount Installation Instructions* or *CHALLENGE/Onyx L Deskside Installation Instructions*.

3.4 Planning the Cabling

The numbering of SCSI controllers (ports) on the OPS host IO4 board and mezzanine card(s) (daughter boards), if present, is standardized. Figure 3-2 diagrams the numbering scheme. Note that the numbering is constant regardless of which daughter board is installed.




Figure 3-2 SCSI Controller Numbering on IO4 and Mezzanine Boards

In an OPS configuration, each SCSI bus runs from a controller on one Challenge server through a vault or Challenge RAID to the second OPS host. For clarity, connect the controllers with the same numbers on a bus; for example, connect cable SCSI controller 2 on the first OPS host (OPS1) to SCSI controller 2 on the second OPS host (OPS2). Figure 3-3 diagrams a typical OPS configuration.

Note: Make sure that controllers 2 and 5 and bus A on both optional vaults are differential.

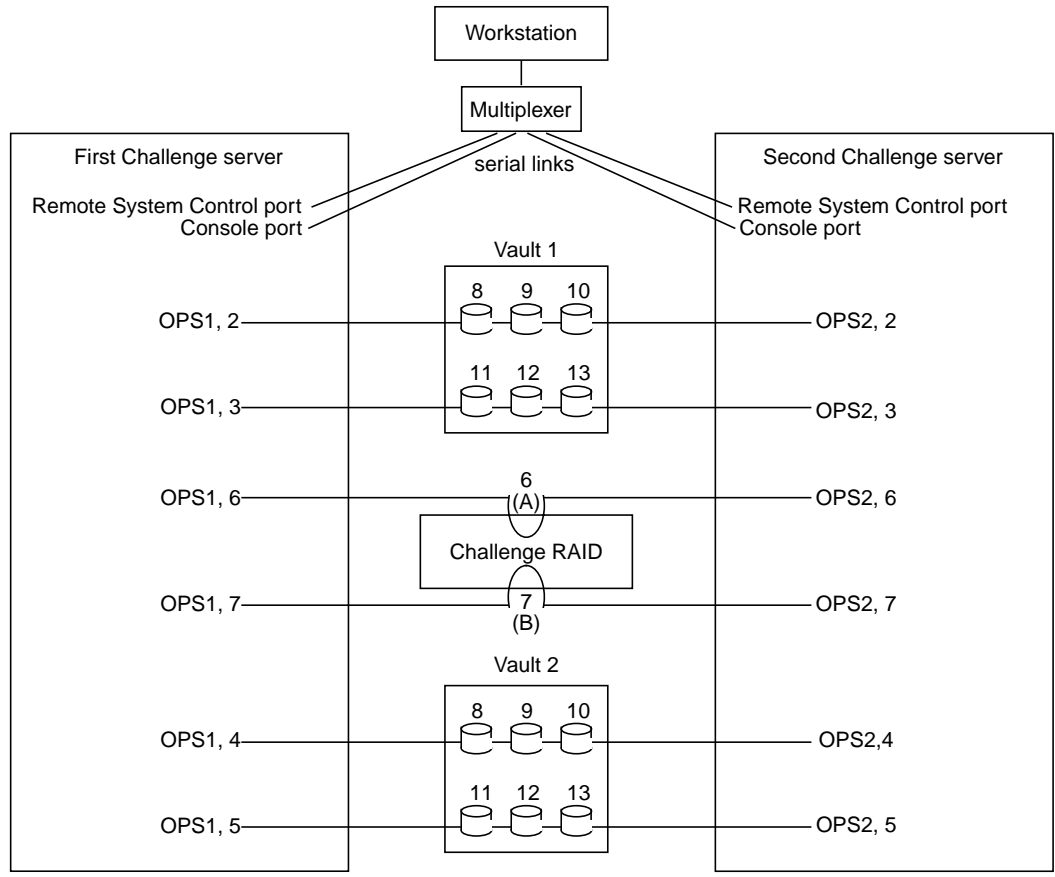


Figure 3-3 Example OPS Cabling Scheme

Figure 3-4 shows servers and vaults cabled for OPS.




Figure 3-4 OPS Cabling for Servers and Vaults

Note: The IRISconsole is also linked to the Challenge RAID storage system; this link is not shown in Figure 3-4.

The SCSI buses in the vaults and the Challenge RAID storage system are named A and B. The recommended procedure is to connect even-numbered SCSI buses to A and odd-numbered buses to B.

3.5 Cabling the Vaults to the OPS Hosts

This section explains how to cable each SCSI bus on each vault to each OPS host. Follow these steps:

1. Have ready eight of the OPS host-vault cables. Use only these cables for this connection.
2. On blank labels, create labels **OPS1**, **OPS2**, **Vault 1**, and **Vault 2**. Attach the labels to the appropriate OPS hosts and vaults.
3. Power off both servers and vaults.
4. On the back of vault 1, insert one of the cables into the top SCSI channel A socket, as shown in Figure 3-5.
5. Label the cable end or connector; follow the scheme in Figure 3-3.
6. Connect the other end of this cable to an even-numbered SCSI bus port on the first OPS host (OPS1), as shown in Figure 3-5. Label the connector or cable end.




Figure 3-5 Connecting the First Challenge Vault and Server

7. On the back of vault 1, insert a cable into the top SCSI channel B socket. Label the cable end.
8. Connect the other end of this cable to an odd-numbered SCSI bus port on the first OPS host (OPS1), as shown in Figure 3-6. Label the cable end or connector.




Figure 3-6 Connecting SCSI Buses A and B from the First Vault to the First Server

9. Insert a third cable into the lower SCSI channel A port on the back of vault 1. Label the cable end.
10. Connect the other end of this cable to the same even-numbered SCSI bus port on the second OPS host (OPS2) that you used on OPS1, as shown in Figure 3-7. Label the cable end.
11. Insert a fourth cable into the lower SCSI channel B port on the back of vault 1 and label it. Connect its other end to same odd-numbered SCSI bus port on the second OPS host (OPS2) that you used on OPS1. Label the cable end. Figure 3-7 shows all cables for the first vault installed.




Figure 3-7 Cabling the First Vault to Both Servers

12. Connect the four remaining cables to the second vault; follow the pattern in steps 4 through 11. Figure 3-4 earlier in this chapter shows this cabling.
13. If necessary, reset the SCSI bus selector settings on the front of each Challenge vault if they have been unset accidentally.
14. If one or more optional Digital Linear Tape drives (DLT) were ordered, install them in the server(s).
15. Power on the vaults. Power on the servers.

16. Do a dummy I/O to test each link. For example, in the numbering scheme used so far, to test SCSI bus 2, send data from OPS1 to disk 8 on SCSI bus 2 (dks2d8) and see if the disk light for disk 8 on the first vault lights up. Send data from OPS2 and see if the same LED is activated. Repeat this process for each bus.
17. Halt and power off each system.

3.6 Setting the OPS Host SCSI IDs

SCSI host address conflicts have caused the most problems in getting OPS up and running.

Caution: The two Challenge servers must have different SCSI IDs. The recommended addresses for OPS are 0 for the first OPS host (OPS1) and 2 for the second OPS host (OPS2). No SCSI device on any bus on any server should have SCSI ID 2.

Setting the host SCSI IDs consists of these steps:

1. For the first OPS host, enter the command *nvrnm*. The last line of this command's output should contain the line

```
scsihostid=
```

This output means that the SCSI host ID for this OPS host is 0, the default.
2. Repeat step 2 in the second window for the second OPS host. The last line of the output should be `scsihostid=2`, meaning that the SCSI host ID for the second OPS host is 2.

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

1. Reboot and start 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
```

Caution: No SCSI device on any bus on any 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*; see Section 3.1, "Setting Up the Component Systems," for a sample output.
6. Enter the command *hinv* for the second server. Compare the output for each server.

3.7 Installing the Network Boards in the OPS Hosts

Follow these steps:

1. Install the cards following instructions in the installation guides for the network cards (FDDI or VME 100-BASE-T).

Caution: This equipment is 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.

2. Edit */etc/config/netif.options* to include network addresses, for example:

```
# To override the name and/or address of the first gateway interface,  
# change the value part and remove the leading : character.
```

```
: if2name=ipg0  
: if2addr=d1m-$HOSTNAME
```

```
# If this host has more than 2 interfaces, you must define values for  
# if3name (and if4name if appropriate). Change if3addr (and if4addr) to  
# the appropriate names in /etc/hosts if your site has different naming  
# conventions.
```

```
if3name=  
if3addr=gate2-$HOSTNAME
```

```
if4name=  
if4addr=gate3-$HOSTNAME
```

Edit */etc/hosts* to include IP numbers for OPS1 and OPS2. Use test numbers that are not in use on the customer's network. For example:

```
127.0.0.1      localhost  
#  
# This network number, '192.0.2' is the officially blessed 'test' network.  
# It should be changed immediately to your own official network number.  
  
# start of well known machines--do not remove this line  
# main ethernet network  
192.26.79.38  ops1  
192.26.79.39  ops2  
# for DLM FDDI connection only  
192.0.2.200   d1m-ops1  
192.0.2.201   d1m-ops2  
#
```

3.8 Cabling the Challenge RAID Storage System to the Challenge Servers

The SCSI ports on the Challenge RAID storage system are cabled to both Challenge servers; consult the diagram in Figure 3-3.

Table 3-1 charts SCSI ID switch settings for the Challenge RAID storage system. Note that the settings do not conform to frequently used numbering schemes.

Table 3-1 Challenge RAID SCSI ID Switch Settings

SCSI ID Number	Switch Number				Comment
	ID 0	ID 1	ID 2	ID 3	
1	Off	On	On	On	
2	On	Off	On	On	Do not use
3	Off	Off	On	On	
4	On	On	Off	On	
5	Off	On	Off	On	
6	On	Off	Off	On	Shown in Figure 3-10 as A
7	Off	Off	Off	On	Shown in Figure 3-10 as B
8	On	On	On	Off	
9	Off	On	On	Off	
10	On	Off	On	Off	
11	Off	Off	On	Off	
12	On	On	Off	Off	
13	Off	On	Off	Off	
14	On	Off	Off	Off	
15	Off	Off	Off	Off	

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

1. Have ready the two remaining SCSI cables and the two 20-foot SCSI cables included in the Challenge RAID shipment (p/n 9290111).
2. Using blank labels, create two labels that say **DF SCSI n** and two labels that say **DF SCSI N**, where **n** and **N** are the numbers of two remaining SCSI controllers on the servers (in the example in Figure 3-3, these are 6 and 7).
3. Attach the two even-numbered labels below the **SCSI A** in symbol and **SCSI A** out symbol silk-screens on the Challenge RAID storage system. Attach the two odd-numbered labels below the **SCSI B** in symbol and **SCSI B** out symbol on the Challenge RAID storage system. Figure 3-8 shows the Challenge RAID SCSI ports.

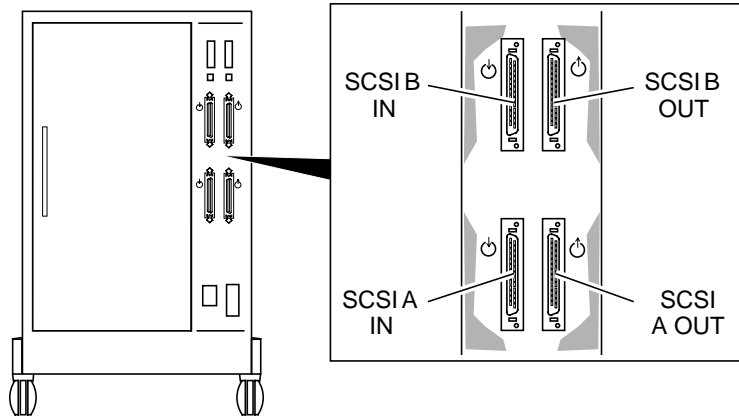


Figure 3-8 SCSI-2 Bus Connectors on Back of Challenge RAID Chassis

4. On the back of the Challenge RAID storage system, attach a SCSI cable to the Challenge RAID **SCSI A** in port. Connect the SCSI cable to the **SCSI A** in connector as shown in Figure 3-9, making sure the connector is inserted securely.




Figure 3-9 Connecting a SCSI Bus Cable to a Challenge RAID SCSI Port

5. Connect the other end of this cable to the even-numbered SCSI bus on OPS1, as shown in Figure 3-10.




Figure 3-10 Connecting a Challenge RAID SCSI Bus Cable to a Challenge SCSI Port

6. Continue cabling as diagrammed in Figure 3-11.




Figure 3-11 Challenge RAID OPS Connection Schemes

Figure 3-12 shows a typical hookup for an OPS configuration with one Challenge RAID storage system.




Figure 3-12 Connecting SCSI Cables to a Challenge RAID and Both OPS Hosts

3.9 Cabling the IRISconsole Multiplexer to the Workstation

To cable the Indy or O2 workstation to the Challenge servers, follow these steps:

1. Have ready the contents of the IRISconsole multiplexer kit and the pair of cables included in the cable kit, so that you have four 2-ft DB25-DB9 serial cables (p/n 018-0527-001).
2. Connect one end of the SCSI cable (018-0526-001) included in the multiplexer kit to the workstation's SCSI port, as shown in Figure 3-13 and Figure 3-14.




Figure 3-13 Cabling the Multiplexer to an Indy Workstation




Figure 3-14 Connecting the Multiplexer's SCSI Cable to an O2 Workstation

3. Attach the other end of the cable to the leftmost socket on the back of the multiplexer, as shown in Figure 3-13.
4. If the customer has purchased two multiplexers, attach the cable of the second multiplexer to the middle socket on the first multiplexer.
5. Cable the power supply to the rightmost port on the multiplexer, as shown in Figure 3-15.
6. Select an appropriate SCSI ID number by inserting a small screwdriver or other tool and rotating the arrow, as shown in Figure 3-15. Using SCSI ID 7 on early models of the multiplexer causes the power indication LED to blink. Select another ID and the LED glows steadily with no blink.




Figure 3-15 Setting the Multiplexer SCSI ID

7. If the multiplexer is the last SCSI device on the workstation's standard external SCSI port, terminate it properly by attaching the external SCSI terminator included in the multiplexer shipment. See Figure 3-15.
8. Reboot the workstation.

3.10 Cabling the IRISconsole Multiplexer and the Challenge Servers

Follow these steps:

1. Attach a 25-pin end of a serial cable to serial port 1 (or another port, as the customer wishes) on the top of the multiplexer.
2. Connect the 9-pin connector from serial port 1 on the multiplexer to the connector labeled **System Console tty_1** on the I/O panel of the first OPS host (OPS1), as shown in Figure 3-16.




Figure 3-16 Cabling the OPS Host Remote System Control Connector

3. Attach a 25-pin end of a serial cable to serial port 2 (or another port, as the customer wishes) on the top of the multiplexer. Connect the other end of this cable to the **Remote Sys Control SSE Use Only** connector, which is to the right of the I/O panel on the first OPS host, as shown in Figure 3-17.




Figure 3-17 Cabling the OPS Host System Console Connector

Figure 3-18 shows cabling for a Challenge XL.




Figure 3-18 Multiplexer Cabling for Challenge XL Server

- Repeat steps 1 through 3 for the second OPS host: attach serial cables from ports 3 and 4 on the multiplexer to the OPS2's **System Console tty_1** and **Remote Sys Control SSE Use Only** ports, respectively.

Figure 3-19 diagrams IRISconsole multiplexer cables.




Figure 3-19 IRISconsole Cabling for OPS (Challenge L Servers)

Note: Do not use the 8-pin powered serial connector in systems that have the IRISconsole multiplexer installed.

3.11 Powering On the Systems

To power on the component systems, follow these steps:

1. Make sure that the Challenge RAID power switch is off.
2. Connect the power cord to the storage system. Plug the other end of the power cord into an AC power outlet. See Figure 3-20.




Figure 3-20 Connecting the Power Cord (Back of Challenge RAID Storage System)

3. Turn on the storage system's power. The green power light on the front of the Challenge RAID storage system comes on, and the fans rotate. See Figure 3-21.




Figure 3-21 Turning On the Storage System Power

4. If the busy light on *none* of the drive modules lights up, make sure that the power for each storage-control processor is enabled. Move the fan module's latch to the **UNLOCK** position, as indicated in Figure 3-22.




Figure 3-22 Unlocking the Fan Module

5. Swing open the fan module and move the power switch for each disabled storage-control processor to the enable position, as shown in Figure 3-23.




Figure 3-23 Enabling a Storage-Control Processor's Power

6. Swing the fan module closed, and then move the latch to the **LOCK** position.

Caution: Do not leave the fan module open for more than two minutes at any time. The Challenge RAID storage system shuts down if the fan module is disabled for more than two minutes.
7. If necessary, connect the power cables and power on the OPS hosts, the Challenge RAID, the workstation, and the multiplexer.

3.12 Testing the Installed OPS Hardware

To test the OPS installation, run *hinv* on a server. Figure 3-24 shows a sample output.

```

6 150 MHZ IP19 Processors
CPU: MIPS R4400 Processor Chip Revision: 5.0
FPU: MIPS R4010 Floating Point Chip Revision: 0.0
Data cache size: 16 Kbytes
Instruction cache size: 16 Kbytes
Secondary unified instruction/data cache size: 1 Mbyte
Main memory size: 256 Mbytes, 2-way interleaved
I/O board, Ebus slot 5: IO4 revision 1
Integral EPC serial ports: 4
Integral Ethernet controller: et0, Ebus slot 5
Efast FXP controller: fxp0
EPC external interrupts
Integral SCSI controller 1: Version WD33C95A,
differential, revision 0
Disk drive: unit 1 on SCSI controller 1
Integral SCSI controller 0: Version WD33C95A, single
ended, revision 0
Tape drive: unit 6 on SCSI controller 0: DAT
Integral SCSI controller 7: Version SCIP/WD33C95A
Integral SCSI controller 6: Version SCIP/WD33C95A
Integral SCSI controller 5: Version SCIP/WD33C95A
Disk drive: unit 5, lun 3 on SCSI controller 5
Disk drive: unit 5, lun 2 on SCSI controller 5
Disk drive: unit 5, lun 1 on SCSI controller 5
Disk drive: unit5 on SCSI controller 5
Integral SCSI controller 4: Version SCIP/WD33C95A
Disk drive: unit 15 on SCSI controller 4
Disk drive: unit 14 on SCSI controller 4
Disk drive: unit 13 on SCSI controller 4
Disk drive: unit 12 on SCSI controller 4
Disk drive: unit 11 on SCSI controller 4
Disk drive: unit 10 on SCSI controller 4
Disk drive: unit 9 on SCSI controller 4
Disk drive: unit 8 on SCSI controller 4
Integral SCSI controller 3: Version SCIP/WD33C95A
Disk drive: unit 3, lun 3 on SCSI controller 3
Disk drive: unit 3, lun 2 on SCSI controller 3
Disk drive: unit 3, lun 1 on SCSI controller 3
Disk drive: unit 3 on SCSI controller 3
Integral SCSI controller 2: Version SCIP/WD33C95A
Disk drive: unit 15 on SCSI controller 2
Disk drive: unit 14 on SCSI controller 2
Disk drive: unit 13 on SCSI controller 2
Disk drive: unit 12 on SCSI controller 2
Disk drive: unit 11 on SCSI controller 2
Disk drive: unit 10 on SCSI controller 2
Disk drive: unit 9 on SCSI controller 2
Disk drive: unit 8 on SCSI controller 2
Integral EPC parallel port: Ebus slot 5
VME bus: adapter 0 mapped to adapter 21
VME bus: adapter 21

```

RAID SP B

Vault 1

RAID SP A

Vault 2

```

4 160 MHS IP19 Processors
CPU: MIPS R4400 Processor Chip Revision: 5.0
FPU: MIPS R4010 Floating Point Chip Revision: 0.0
Data cache size: 16 Kbytes
Instruction cache size: 16 Kbytes
Secondary unified instruction/data cache size: 1 Mbyte
Main memory size: 256 Mbytes, 2-way interleaved
I/O board, Ibus slot 5: IO4 revision 1
Integral IPC serial ports: 4
Integral Ethernet controller: et0, Ibus slot 5
IFast FXP controller: fxp0
IPC external interrupts
Integral SCSI controller 1: Version WD33C95A, differential, revision 0
Disk drive: unit 1 on SCSI controller 1
Integral SCSI controller 0: Version WD33C95A, single ended, revision 0
Tape drive: unit 6 on SCSI controller 0: DAT
Integral SCSI controller 7: Version SCIF/WD33C95A
Integral SCSI controller 6: Version SCIF/WD33C95A
Integral SCSI controller 5: Version SCIF/WD33C95A
Disk drive: unit 5, lun 3 on SCSI controller 5
Disk drive: unit 5, lun 2 on SCSI controller 5
Disk drive: unit 5, lun 1 on SCSI controller 5
Disk drive: unit 5 on SCSI controller 5
Integral SCSI controller 4: Version SCIF/WD33C95A
Disk drive: unit 15 on SCSI controller 4
Disk drive: unit 14 on SCSI controller 4
Disk drive: unit 13 on SCSI controller 4
Disk drive: unit 12 on SCSI controller 4
Disk drive: unit 11 on SCSI controller 4
Disk drive: unit 10 on SCSI controller 4
Disk drive: unit 9 on SCSI controller 4
Disk drive: unit 8 on SCSI controller 4
Integral SCSI controller 3: Version SCIF/WD33C95A
Disk drive: unit 3, lun 3 on SCSI controller 3
Disk drive: unit 3, lun 2 on SCSI controller 3
Disk drive: unit 3, lun 1 on SCSI controller 3
Disk drive: unit 3 on SCSI controller 3
Integral SCSI controller 2: Version SCIF/WD33C95A
Disk drive: unit 16 on SCSI controller 2
Disk drive: unit 14 on SCSI controller 2
Disk drive: unit 13 on SCSI controller 2
Disk drive: unit 12 on SCSI controller 2
Disk drive: unit 11 on SCSI controller 2
Disk drive: unit 10 on SCSI controller 2
Disk drive: unit 9 on SCSI controller 2
Disk drive: unit 8 on SCSI controller 2
Integral IPC parallel port: Ibus slot 5
VME bus: adapter 0 mapped to adapter 21
VME bus: adapter 21

```

RAID SP B

Vault 1

RAID SP A

Vault 2

Figure 3-24 Example hinv Output for Installed OPS System (Challenge)

Installing the Oracle Parallel Server Software

After the hardware is set up, cabled, and tested, install the required software, following these procedures:

1. checking the Challenge server (OPS host) SCSI IDs
2. installing system software
3. installing the IRISconsole software and configuring the site
4. installing and configuring the Challenge RAID storage system software for OPS
5. enabling command-tagged queuing on the LUNs
6. partitioning log files and creating the XLV volumes
7. installing the Oracle database software
8. installing the OPS software
9. starting OPS

Each step is explained in a separate section in this chapter.

4.1 Checking the OPS Host SCSI IDs

This section explains how to ensure that the two OPS hosts (Origin or Challenge servers or other SCSI devices on the same bus) do not have the same SCSI ID number. If you have just completed the hardware installation, skip to the next section; this information is provided for cases in which software installation does not immediately follow hardware installation.

Follow these steps:

1. Using your laptop or a customer ASCII terminal, log on to each OPS host.
2. For the first OPS host, enter the command *nvrnm*. The output should contain the line

```
scsihostid=
```

This output means that the SCSI host ID for this OPS host is 0.

3. For the second OPS host, repeat step 2 in the second window. The output should contain the line

```
scsihostid=2
```

This output means that the SCSI host ID for the second OPS host is 2.

If the host SCSI IDs are not 0 and 2, respectively, follow the instructions in Chapter 2 and Chapter 3 to set them.

4.2 Installing System Software

The system software consists of patches for IRIX and the OPS software package. For efficiency, you can install all the software described in this section on the first OPS host (OPS1) and copy the disk to OPS2 using the *sash cp* command.

The section explains

- installing IRIX and the patches on the OPS hosts
- converting to the XFS™ filesystem
- installing the *opscm* software on the OPS hosts
- installing the *opsnc* software and configuration file on the workstation
- testing the serial connection to the OPS hosts

4.2.1 Installing IRIX and the Patches on the OPS Hosts

To install system software, follow these steps:

1. Check to see if IRIX 6.2 or IRIX 6.4 is installed on the first OPS host.
2. Load *oe2.sw.xlvplex*.
3. Install all recommended IRIX patches.
4. If necessary, change the root filesystem to an XFS filesystem.
5. Repeat steps 1 through 3 for the second OPS host and the workstation.

For complete information on installing an XFS filesystem, see *IRIX Admin: Disks and Filesystems* volume, which covers XLV volumes for IRIX 6.2, as well as IRIX 6.4 patches (007-2825-002).

4.2.2 Converting to the XFS Filesystem

To convert the filesystems on both hosts and the workstation to XFS, see *IRIX Admin: Disks and Filesystems*. When you use the *mkfs* command to create the new XFS filesystem, use

```
mkfs -d name=<partition_name> -l internal,size=1000b
```

4.2.3 Installing the opscm Software on the OPS Hosts

Installing the *opscm* software on the first OPS host consists of these steps:

- installing the *opscm* software
- editing the configuration file
- editing *sidconf*

4.2.3.1 Installing the opscm Software

To install the *opscm* software on the hosts, follow these steps:

1. Enter the following commands:

```
ops1 1# su - root
ops1 2# inst -f /CDROM/opscm_eoe
```

The Inst Main Menu appears.

2. Select item 2, "list."
3. At the *Inst* prompt, type *go*.
4. After the program completes the install process, type *quit* at the prompt.

4.2.3.2 Editing the opscm Configuration File

The configuration file for *opscm*, */usr/opscm/conf/opsconf*, lists information on

- the cluster controller (workstation and IRISconsole software) managing the OPS hosts
- each OPS instance and its peer, plus OPS instance/peer pairs running on other nodes
- time parameters for the *opscm* and lock manager software

To edit the configuration file, follow these steps:

1. In your favorite editor, edit */usr/opscm/conf/opsconf* to match the following.

Note: In this output, IRISconsole refers to the workstation acting as the node controller for OPS (using IRISconsole software).

```
# ident "$Revision: 1.5 $"

#
#CLUSTER id name size console
#
CLUSTER 1 test 2 ops-indy

#id any integer -- not used
#name just a ascii string -- not used
#size better be 2 -- not used
#console hostname of IRISConsole
#cons_port port number on IRISConsole where cluster_daemon listens
```

```

#
#NODE dom inst ndname ndaddress      cmsvc  apsvc  wt
#
NODE 0  0   host1  150.166.42.37  opscm  opsdlm 1
NODE 0  1   host2  150.166.42.38  opscm  opsdlm 1

#dom      domain nr, (LKDOM environment variable)
#inst     instance id (LKINST environment variable)
#ndname   hostname, same as in the /usr/opsnc/conf/nconf file on IRISconsole
#ndaddr   ip address tied to the OPS private network interface
#cmsvc    from services(4):port nr on which the opscm listens for connections
#apsvc    from services(4):port nr on which the OPS dlm listens for connections
#wt       any integer - not used in this release

#
#TIME parameters (all units are seconds, except tick which is microsec)
#TIME init ping ack sync cnct sdwn tick   hb hbto hold pwr rcfg
#
TIME 60  5   50  0   120 30   500000 2 10 20 30 5

#init CM waits init sec on startup to hear from peer
#ping CM sends a ping message to DLM every ping sec
#ack CM exits if it does not receives a ping ack from DLM for ack sec
#sync CM aborts DLM sync op unless it completes in sync sec [0=inf]
#cnct CM exits if DLM does not connect in cnct sec
#sdwn CM exits if DLM does not detach sdwn sec after Stop is sent [0=inf]
#tick CM looks for work every tick MICROsec
#hb CM sends heartbeat message to peer every hb sec
#hbto CM declares peer crashed if no HB received for hbto sec
#hold CM tries to reconnect with peer hold sec after peer crashes
#pwr assume power failed if both HB and remote sysctlr fail within pwr sec
#rcfg DLM specific parameter

```

2. In the CLUSTER line, enter the machine name of the workstation being used for IRISconsole.
3. Define the DLM domain(s) and DLM instances, as explained in the *Oracle Parallel Server Administrator's Guide*. Check that the lines under

```

#NODE dom inst ndname ndaddress      cmsvc  apsvc  wt

```

have accurate information on your OPS servers. For example:

```

NODE 0  0   host1  150.166.42.37  opscm  opsdlm 1
NODE 0  1   host2  150.166.42.38  opscm  opsdlm 1

```

4. Save and close */usr/opsnc/conf/opsconf*.
5. Edit the file */etc/services*; add three services:

```

opscm      newnumber1/tcp
opsdlm     newnumber2/tcp
opsnc      newnumber3/tcp

```

where *newnumbern* is a number not used elsewhere in this file. For example:

```

opscm      7018/tcp
opsdlm     7019/tcp
opsnc      7020/tcp

```

6. Save and exit the file.
7. Set the ownership, permissions, and timeouts in the *config* file according to the customer's wishes.

4.2.3.3 Editing *sidconf*

The *sidconf* file maps each Oracle instance (sid, or Oracle system ID) to a DLM domain-instance pair. For information on DLM domains and instances, see Chapter 1 of the *Oracle Parallel Server Administrator's Guide*.

To edit the *sidconf* configuration file, follow these steps:

1. Determine the DLM domain and DLM instance (sid) of the Oracle database for each instance.
2. In */usr/opscm/conf/sidconf*, enter correct information for the two OPS servers in the lines

```
MAP sid0 domainnumber instancenumber
MAP sid1 domainnumber instancenumber
```

The variables *sid0* and *sid1* must contain no more than four characters. For example:

```
MAP fin1 0 0
MAP fin2 0 1
```

3. Save and close */usr/opscm/conf/sidconf*.
4. If necessary, change permissions on this file so that it can be read by everyone.
5. When you have finished installing the system software on the first Challenge server, copy its disk to the second Challenge server; change the hostname.

4.2.4 Installing the *opsnc* Software and Configuration File on the Workstation

To install the *opsnc* software on the workstation, follow these steps:

1. Enter the following commands:

```
ops1 1# su - root
ops1 2# inst -f /CDROM/opsnc_eoe
```

The Inst Main Menu appears.

2. At the *Inst* prompt, type *go*.
3. After the program completes the install process, type *quit* at the prompt. The following messages appear:

```
Building dynamic ELF inventory file for rqs(1) processing .. 100% Done.
Invoking rqs(1) on necessary dynamic ELF objects .. 100% Done.
```

4. If desired, list the installed software; type *ls -lR /usr/opsnc*. The following should appear:

```
total 0
drwxr-xr-x  2 root  sys      36 Apr 20 13:28 bin/
drwxr-xr-x  2 root  sys      24 Apr 20 13:28 conf/

./bin:
```

```

total 107
-r-xr-x--- 1 root sys 39792 Apr 20 13:28 opsnc*
-r-xr-x--- 1 root sys 14368 Apr 20 13:28 spng*

./conf:
total 1
-rw-rw-r-- 1 root sys 346 Apr 20 13:28 ncconf

```

5. Use your favorite editor to edit the configuration file `/usr/opsnc/conf/ncconf`. Enter information for each OPS server under

```
#nodename      ttyname      systemcontrollertype
```

For example:

```
ops1  /dev/ttyf031  SYSCTLR
ops2  /dev/ttyf033  SYSCTLR
```

The last digit in each `tty` entry should be the multiplexer port into which the Remote System Control for each server is cabled. The numbers 1 through 16 on the multiplexer correspond to 0 through f in the `tty` entries. In the example above, the Remote System Control ports are cabled to ports 2 and 4 on the multiplexer.

The third field is `SYSCTLR` for Challenge servers, `MSC` for deskside Origin2000 servers, and `MMSC` for rackmount Origin2000 servers.

6. To test the connection to the multiplexer, become superuser at the shell prompt to determine if the system recognizes the multiplexer. As the superuser, enter the `/sbin/cdscanbus` command; the output should look like the following:

```
# /sbin/cdscanbus
-----
SCSI ID: 1
Device type: 0x0
Vendor ID: 'SGI '
Product ID: 'SEAGATE ST31200N'
Firmware: '8640'
Vendor Unique:
30 30 32 33 34 31 34 34 00 00 00 00 00 00 00
00234144.....
-----
SCSI ID: 5
Device type: 0x9
Vendor ID: 'CenData '
Product ID: '<multiplexername>'
Firmware: 'V3.9'
Vendor Unique:
20 53 65 6e 64 20 4c 55 4e 20 00 00 00 00 00 00
.Send.LUN.....
```

Available SCSI ID numbers on bus 0: 2 3 4 6 7

7. If the multiplexer does not appear in the `/sbin/cdscanbus` listing, recheck all the following items:

- SCSI cable connections
- SCSI terminator and ID
- multiplexer power connection

- proper software installation

You must have the multiplexer software drivers installed, or the `cdscanbus` command will not work. See the section “Basic Troubleshooting” at the end of Chapter 2 of the *IRISconsole Multiplexer Installation Guide*.

8. If the OPS servers are Origin servers and have MSC or MMSC system controller passwords, you must initialize them on the IRISconsole workstation. Enter

```
/usr/opsnc/bin/spng -w password -n nodename -d systemcontrollertype
```

For example:

```
/usr/opsnc/bin/spng -w abcd -n ops1 -d MSC
```

4.2.5 Testing the Serial Connection to the OPS Hosts

To test the serial connection between the workstation and the OPS hosts, follow these steps:

1. Make sure the OPS hosts are powered on.
2. If OPS hosts are Challenge servers, enter

```
/usr/opsnc/bin/spng -i 10 -f /dev/ttyfn
```

where `/dev/ttyfn` is the name of the device on the workstation connected to the Remote System Control port on the OPS host.

If the OPS hosts are deskside Origin2000 or Origin200™ servers, enter

```
/usr/opsnc/bin/spng -i 10 -dMSC -f /dev/ttyfn
```

If the OPS hosts are rackmount Origin2000 servers, enter

```
/usr/opsnc/bin/spng -i 10 -dMMSC -f /dev/ttyfn
```

The command should return zero error states if the serial connection is correct.

3. If this type of output does not appear, check the following:
 - verify that the OPS host is powered on
 - verify the cable connections from the multiplexer to the Remote System Control port and from the workstation to the multiplexer
 - debug the IRISconsole installation; see “Basic Troubleshooting” at the end of Chapter 2 in the *IRISconsole Multiplexer Installation Guide*

Note: For more information on troubleshooting, see Chapter 5, “Maintaining and Upgrading the Oracle Parallel Server.”

4.3 Installing the IRISconsole Software and Configuring the Site

To set up the workstation for OPS, follow these steps:

1. If the customer has purchased a new workstation for use with OPS, install its software and configure it.
2. From the distribution CDs included in the IRISconsole kit, load the IRISconsole software as explained in Chapter 2 of the *IRISconsole Multiplexer Installation Guide*.
3. Use the IRISconsole GUI to set up the customer's *site*, which is the IRISconsole grouping of the two OPS hosts for administration purposes: as root; type `/usr/sbin/ic`. The IRISconsole main window appears.
4. In the Edit menu, select "Add New Site." The Change IRISconsole Config dialog box appears.
5. In the dialog box, type a name for the site, such as `ops`. Click the *Apply* button; select "Save IRISconsole Config" in the File menu and close the dialog box.
6. To add the systems to the site, double-click on the site icon in the main menu. The site window appears.
7. In the site window, select "Add New System..." in the Edit menu. The Change Site Config dialog box appears, as shown in Figure 4-1.

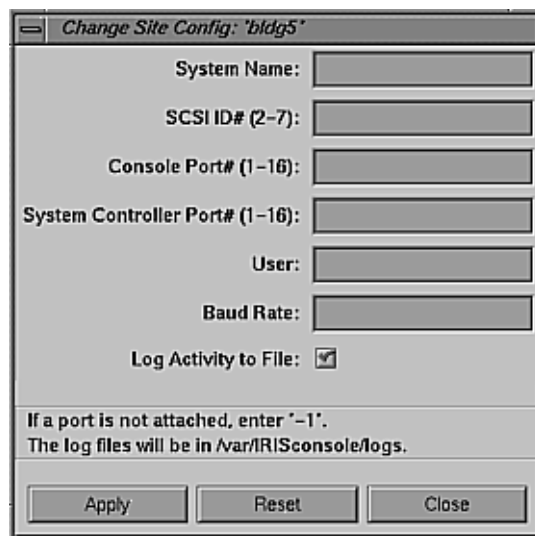


Figure 4-1 Change Site Config Dialog Box

8. Enter information for the first OPS server. If necessary, check the multiplexer and the cables for the ports to enter. Ask the customer for the correct entry for the User: field. When the information for the first OPS server is correct, click the *Apply* button.
9. Click *Reset* to clear the fields and enter the information for the second OPS server. When the information for the second OPS server is correct, click *Apply*. Click *Close* to exit the dialog box.
10. Save the configuration by selecting "Save System Config" in the site window File menu. If the customer desires, save the configuration to a particular name (Save As). Exit the IRISconsole GUI.

Note: For more detail on steps 3 through 10, see Chapter 2 of the *IRISconsole Administrator's Guide*.

4.4 Installing and Configuring the Challenge RAID Storage System Software for OPS

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

Configuring the Challenge RAID storage system for OPS consists of these processes:

- creating the Challenge RAID config file for use with OPS
- restarting the agent and checking the config file
- creating SCSI device nodes (IRIX 6.2 only)
- configuring LUNs

4.4.1 Creating the Challenge RAID Config File for Use With OPS

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

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

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

1. Open */usr/raid5/raid5-agent.config.proto* in your favorite editor. This file is self-documenting and contains samples.
2. Save */usr/raid5/raid5-agent.config.proto* as */usr/raid5/raid5-agent.config*.
3. In the RAID specification section, add RAID devices, using the controller numbers displayed in the *hinv* output (see Section 3.1 earlier in this chapter). For example, to add SPs:

```
device sc6d610 sgi_clariion_SPA "SGI Agent Development System"  
device sc6d510 sgi_clariion_SPB "SGI Agent Development System"
```

Note: In */usr/raid5/raid5-agent.config*, only the SCSI device number for LUN 0 is required.

4. Consulting with the customer, enter the userids of those who can perform write operations in the User Specification section.

5. Consulting with the customer, enter the poll specification in the Poll Specification section.
6. Save and exit. Restart the agent as explained in the next section.
7. Repeat steps 1 through 6 for the second OPS host.

4.4.2 Restarting the Agent and Checking the Config 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/raidcli getagent` command to display information on RAID devices. The output of this command covers all devices; because the disks have not been bound into LUNs, information for only the storage control processors (SPs) is shown; for example

```
Agent Rev: 1.55
Name: PHX-SPA-sc7d2
Desc: TOWER A
Node: sc7d210
Signature: 0x65a91c00
Peer Signature: 0xc3d81f00
Revision 9.3.0, Model: 7305
SCSI Id: 2
Prom Rev: 1.55.0
SP Memory: 32
Serial No: 96-2192-4
```

```
Agent Rev: 1.55
Name: PHX-SPB-sc5d5
Desc: TOWER A
Node: sc5d510
Signature: 0xc3d81f00
Peer Signature: 0x65a91c00
Revision 9.3.0, Model: 7305
SCSI Id: 5
Prom Rev: 1.55.0
SP Memory: 32
Serial No: 96-2192-4
```

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

4.4.3 Creating SCSI Device Nodes (IRIX 6.2 Only)

Note: The instructions in this section are for IRIX 6.2 only. They are not required for IRIX 6.4; `ioconfig(1M)` in that version of IRIX automatically creates device nodes.

To create SCSI device nodes, follow these steps:

1. Become superuser on the first OPS host using your laptop or an ASCII terminal.

Note: If you use the IRISconsole, you can open a window for each OPS host.

2. If you are creating more than the default number of partitions, enter this command in the `/dev` directory:

```
MAKEDEV dks scsi
```

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

- for 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
- for 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 OPS server during the software installation process, no device nodes are created.

3. 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/sc6d610
crw-----  1 root    sys      195,403 Mar 21 15:40 /dev/scsi/sc6d611
crw-----  1 root    sys      195,419 Mar 21 15:40 /dev/scsi/sc6d612
crw-----  1 root    sys      195,435 Mar 21 15:40 /dev/scsi/sc6d613
crw-----  1 root    sys      195,451 Mar 21 15:45 /dev/scsi/sc6d614
crw-----  1 root    sys      195,467 Mar 21 15:45 /dev/scsi/sc6d615
crw-----  1 root    sys      195,483 Mar 21 15:45 /dev/scsi/sc6d616
crw-----  1 root    sys      195,499 Mar 21 15:45 /dev/scsi/sc3d317
```

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 `d` is the number of the RAID disk module
- the digit after `l` (the last character in the address) is the LUN number

4. Reboot the system.
5. Repeat steps 1 through 5 for the second OPS host.

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

Follow these steps:

1. Change to the `/dev` directory:

```
cd /dev
```

2. Enter

```
./MAKE_VLUNS <controller-number> <target-number>
```

This command creates the device nodes for the new disks.

3. Bind the newly installed modules into one or more LUNs, as described in Section 4.4.4, “Configuring LUNs,” later in this chapter.

4.4.4 Configuring LUNs

For OPS, configure RAID LUNs as RAID-5 or RAID-3 (“fast” RAID-3). For instructions, see Chapter 4 of the *CHALLENGE RAID Installation and Maintenance Instructions* (007-0128-006).

4.5 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 `fx`, consult its reference page.

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)-----
[ex]lt          [d]ebug/          [l]abel/          [a]uto
[b]adblock/     [ex]ercise/          [r]epartition/   [f]ormat
```

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

```
fx> label/set/param
```

Output such as the following appears:

```
fx/label/set/parameters: Error correction = (enabled)
fx/label/set/parameters: Data transfer on error = (enabled)
fx/label/set/parameters: Report recovered errors = (enabled)
```

```

fx/label/set/parameters: Delay for error recovery = (enabled)
fx/label/set/parameters: Err retry count = (0)
fx/label/set/parameters: Transfer of bad data blocks = (enabled)
fx/label/set/parameters: Auto bad block reallocation (write) =
(enabled)
fx/label/set/parameters: Auto bad block reallocation (read) =
(enabled)
fx/label/set/parameters: Read ahead caching = (disabled)

```

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

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

4. At the CTQ depth prompt, type **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, type **yes**.

The following output appears:

```

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

```

6. To repartition the disk, type **r**.

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

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

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

4.6 Partitioning Disks and Creating the XLV Volumes

When you specify the number of files in the process of creating the Oracle database, do not exceed the XLV limit of 52 volumes, which is 13 partitions (volumes) per RAID LUN. Partitions are numbered 0-7, 11-15; partitions 8, 9, and 10 cannot be used.

Follow these steps:

1. If there are more than ten volumes, you must edit *master.d/xlv* (see “Configuring the System for More Than Ten XLV Logical Volumes” in Chapter 7 of the *IRIX Admin: Disks and Filesystems* volume) and regenerate the kernel using `autoconfig(1M)`.
2. Stop the RAID agent:

```
/etc/init.d/raid5 stop
```
3. Use *fx* to create partitions for the log files on the volumes.
4. Create an *xlv_make* file for the volumes in all partitions.
5. Run *xlv_make* to create the volumes.

6. Edit `/etc/init.d/xlv` to add lines that change the ownership of the Oracle volumes to "oracle." Look for these commented lines:

```
# change ownership of xlv volumes here:
# chown $USER.$GROUP /dev/dsk/xlv/<volname>
# chown $USER.$GROUP /dev/rdisk/xlv/<volname>
```

7. When the volumes are created, change the owner to *oracle*.

Note: On every reboot, *xlv* recreates the volumes; they will be owned by *root.sys*, and the names will be different on each node.

8. Change the node name for all XLV volumes to be used by the OPS software so that the full XLV volume name is the same on both nodes.

When XLV creates volumes, it inserts a node name element in each volume name. For example, if volume *finance1* is created on node *ops1*, XLV names the volume *ops1.finance1*. On node *ops1*, the device name for the volume appears as the matching volume `/dev/dsk/xlv/finance1`, but on node *ops2*, the full XLV volume name appears: `/dev/dsk/xlv/ops1.finance1`. The volume name has this form because XLV on *ops2* sees that the volume *finance1* was built on another server (*ops1*).

However, OPS requires that these shared volumes have exactly the same name. You must rename the volumes manually; do not use a node name. Give the shared volumes the same names on each node. Table 4-1 gives examples.

Table 4-1 Example Shared Volume Names

XLV Shared Volume Name		Renamed as OPS Shared Volume	
<i>ops1</i>	<i>ops2</i>	<i>ops1</i>	<i>ops2</i>
<i>ops1.finance1</i>	<i>ops2.finance1</i>	<i>OPS.finance1</i>	<i>OPS.finance1</i>
<i>ops1.finance2</i>	<i>ops2.finance2</i>	<i>OPS.finance2</i>	<i>OPS.finance2</i>
<i>ops1.finance3</i>	<i>ops2.finance3</i>	<i>OPS.finance3</i>	<i>OPS.finance3</i>

9. Restart the RAID agent:

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

4.7 Installing the Oracle Database Software

Installing the Oracle database software, which requires entering information in a series of menus, takes two to three hours. Follow instructions in the *Oracle Parallel Server for SGI IRIX Installation & Configuration Guide*, an Oracle document (A22258-x). Add the Oracle *bin* directory to *PATH*.

4.8 Installing the OPS Software

To install the OPS software, follow these steps:

1. Log in on each server as the Oracle software owner.
2. Set environment variables:
 - set `ORACLE_HOME` to the name of the directory where the Oracle server is installed on the OPS server
 - set `REMOTE_NODE` to the name of the second OPS server
 - set `REMOTE_ORACLE_HOME` to the name of the directory on `REMOTE_NODE` where you have created an `ORACLE_HOME` directory

3. Use `sysctl -i` to increase the number of semaphores in the kernel:

```
semaem = 16384 (0x4000)
semvmx = 32767 (0x7fff)
semume = 10 (0xa)
semopm = 200 (0xc8)
semmsl = 200 (0xc8)
semnu = 100 (0x64)
semms = 200 (0xc8)
semnmi = 100 (0x64)
```

4. If necessary, adjust the shared memory segment size using `sysctl`.

5. Run the OPS install script as follows:

```
% cd $ORACLE_HOME/ps0
% ./ps0inst.sh
```

6. The `ps0inst.sh` script does three things:

- moves files needed by OPS, such as the Distributed Lock Manager executable file
- installs the OPS and relinks the Oracle server with this option, using the following commands:

```
% cd $ORACLE_HOME rdbms/lib
% make -f oracle.mk parropt
% make -f oracle.mk ioracle
```

(To remove the OPS software, execute these commands substituting `no_parropt` for `parropt` in the second command.)

- copies the locally installed Oracle server to the second OPS server

During installation, the script outputs messages informing you of the installation status. When the installation has finished, the following message appears:

```
Finished installing Oracle Parallel Server
```

7. If the message shown above does not appear, correct any errors and enter

```
% ./copyoh.sh
```

The file `osh` is owned by root and is not copied automatically.

Note: If the message `rcp: osh: Permission denied` appears, ignore it.

8. To put the file *osh* on the second Challenge server, log in as root and copy this file, as follows:

```
# rcp &ORACLE_HOME/bin/osh \  
    $REMOTE_NODE:$REMOTE_ORACLE_HOME/bin  
# rsh &REMOTE_NODE chmod 4711 $REMOTE_ORACLE_HOME/bin/osh
```

If you encounter errors, make sure that the user equivalence is set up for root, or execute the following steps as root:

```
# cd $ORACLE_HOME/bin  
# ftp $REMOTE_NODE
```

Enter the username *root* and the root user password.

```
ftp> cd $REMOTE_ORACLE_HOME/bin  
ftp> bin  
ftp> put osh  
ftp> quit  
# rlogin $REMOTE_NODE -l root  
Password: password  
# chmod 4711 $REMOTE_ORACLE_HOME/bin/osh
```

9. Install any optional software that the customer has purchased.

4.9 Starting OPS

To start the OPS software for normal dual-host operation on both hosts and the workstation node controller, follow these steps:

1. As root on the workstation, start the OPS node controller software by entering

```
/usr/opsnc/bin/opsnc
```
2. Start IRISconsole by entering

```
/usr/sbin/ic
```
3. In the IRISconsole site window, click the *Get Console* button for each OPS host.
4. Check for the presence of the CM lock file in */tmp*. This filename has the format *.nn*, where each digit stands for DLM domain and instance, for example, *.00*. If this file exists, delete it.
5. Check if the CM is already running by entering

```
ps -ef | grep opscm
```

If it is running, type the following to kill it:

```
killall -TERM opscm
```
6. Check if the DLM is already running by entering

```
ps -ef | grep dlm
```

If the DLM is running, enter the following to kill its processes:

```
killall -TERM dlmon  
killall -TERM dlmd
```

7. Run *ipcs* to determine the shared memory and semaphores used on the host. The following is an example output:

```
IPC status from /dev/kmem as of Thu May 18 14:31:22 1995
T      ID      KEY          MODE          OWNER        GROUP
Message Queues:
Shared Memory:
m      0 0x53637444 --rw-r--r--   root         sys
m      301 0x000022bb --rw-rw----   oracle        dba
m      2202 0x0c33b7c9 --rw-r-----  oracle        dba|
Semaphores:
s      2200 0x00000000 --ra-r-----  oracle        dba
```

8. If Oracle or DLM is using any shared memory segments or semaphores, save them to another location if you need them for debugging a DLM or Oracle crash; otherwise, delete them with *ipcrm*. In the example in step 7, you would use

```
ipcrm -m 301 -m 2202 -s 2200
```

9. For each host, set environment variables in the */etc/init.d/opscmgr* script:

```
ORACLE_HOME=/usr/people/oracle
ORACLE_SID=sidname
LKDOM=0
LKINST=0
USER=oracle
GROUP=dba

export ORACLE_HOME ORACLE_SID LKDOM LKINST

/usr/opscm/bin/opscm
$ORACLE_HOME/bin/lkmgrd -u $USER -g $GROUP
```

In each script, make sure that the values for *LKDOM=* and *LKINST=* are accurate for the domain and instance on that host. These values must match those in */usr/opscm/conf/sidconf*, as explained in “Editing *sidconf*,” earlier in this chapter.

10. Run the startup script on each host to bring up OPS.
11. In each host console window, start the Oracle database.
12. To check DLM operation, run *d1mstat*:

```
setenv ORACLE_SID sidname
$ORACLE_HOME/bin/d1mstat -i -q
```

No error messages should appear.

Chapter 5

Maintaining and Upgrading the Oracle Parallel Server

This chapter explains

- isolating a node (preparing to swap an IO4 board)
- replacing cables
- replacing a log disk
- checking OPS server and Challenge RAID SCSI IDs
- checking disk partitions
- exchanging Challenge RAID disk modules
- restoring LUN ownership on a Challenge RAID system

5.1 Isolating a Node

This section explains how to isolate a node (for example, 1) with a failed component from the other node (in this case, 2), which is still running OPS. This technique allows several maintenance procedures to be applied. For example, you can replace the IO4 or the SCIP mez card that affects the SCSI chains joining the two nodes.

The procedures are

- breaking the SCSI chain with the shared RAID
- breaking the SCSI chain with the log disks

5.1.1 Breaking the SCSI Chain With the Shared RAID

Figure 5-1 shows that data between the two OPS systems (OPS1 and OPS2) flows to and from the shared Challenge RAID.




Figure 5-1 Data Flow Between OPS Challenge Systems and Challenge RAID

To break the SCSI chain with the shared Challenge RAID, follow these steps:

1. Power off the host (node) whose IO4 board you want to replace. In this example, it is node 2 (OPS2).
2. Turn off SP A, as diagrammed in Figure 5-2.




Figure 5-2 Halting Data Flow from SP A to the OPS Systems

3. Remove the cable connecting SP A to the host (node) whose I/O or SCSI board you want to replace (in this example, OPS2). Place a differential terminator on A1, as diagrammed in Figure 5-3.




Figure 5-3 Terminating SP A

4. Turn on SP A. Data flow changes, as diagrammed in Figure 5-4.




Figure 5-4 Data Flow with SP A On and Terminated

5. Wait five minutes for SP A to come online.
6. After five minutes have passed, verify that SP A is online by issuing this *raidcli* command:

```
/usr/raid5/raidcli getcrus
```

The output should contain the line

```
SPA State: Present
```

7. Power off SP B. Remove the cable connecting SP B to node 1; place a differential terminator on B1, as diagrammed in Figure 5-5.




Figure 5-5 Terminating SP B

8. Turn on SP B.
9. Follow instructions for replacing the I/O or SCSI board in the server's installation instructions or the SCSI board's installation instructions, respectively.

5.1.2 Breaking the SCSI Chain With the Log Disks

In this section, terms are used as follows:

- log disks: 1X 1Y 1Z are the log disks primarily written by node 1 and mirrored on 1x 1y 1z. Similarly, 2X 2Y 2Z are the log disks primarily written by node 2 and mirrored on 2x 2y 2z.
- vaults: V and v, containing 1X 1Y 1Z 2X 2Y 2Z, and 1x 1y 1z 2x 2y 2z, respectively. 1X 1Y 1Z should be on one SCSI bus within V; 2X 2Y 2Z should be on the other SCSI bus of V. Similarly for the mirror disks.

To break the SCSI chain with the log disks on a vault V, follow these steps:

1. Verify that you can see all the log files by issuing the following SQLDBA command from an instance:

```
connect internal
select * from v$logfile
```

2. Make all disks within vault V inactive by issuing the following SQLDBA commands from any instance:

```
ALTER DATABASE DROP LOGFILE MEMBER '/dev/rdisk/1X';
ALTER DATABASE DROP LOGFILE MEMBER '/dev/rdisk/1Y';
ALTER DATABASE DROP LOGFILE MEMBER '/dev/rdisk/1Z';
ALTER DATABASE DROP LOGFILE MEMBER '/dev/rdisk/2X';
ALTER DATABASE DROP LOGFILE MEMBER '/dev/rdisk/2Y';
ALTER DATABASE DROP LOGFILE MEMBER '/dev/rdisk/2Z';
```

Note: A log file will not be dropped if it is part of the currently active log group or if it is part of the log group that is currently being archived.

3. For the log group that is currently being archived, wait for archiver to finish archiving it and reexecute the DROP command for the file.
4. For the log group that is the current group, switch to the next group and then reexecute the DROP command for the file.

To switch a log group for an instance, the following command must be executed from that instance:

```
ALTER SYSTEM SWITCH LOGFILE;
```

5. Verify that the log file members have actually been dropped by issuing the following SQLDBA command from any instance:

```
connect internal
select * from v$logfile
```

6. Power off vault V.
7. Disconnect the cables connecting V to OPS2; place differential terminators on V where the cables were, as diagrammed in Figure 5-6.




Figure 5-6 Disconnecting Cables from Challenge Vault to Host 2

8. Disconnect FDDI cables and lines, and other network or communication links. OPS2 is now isolated.
9. Add the LOG files back to Oracle by executing the following SQLDBA commands from any instance:

```
ALTER DATABASE ADD LOGFILE MEMBER '/dev/rdsk/1X' REUSE TO GROUP 1;
ALTER DATABASE ADD LOGFILE MEMBER '/dev/rdsk/1Y' REUSE TO GROUP 2;
ALTER DATABASE ADD LOGFILE MEMBER '/dev/rdsk/1Z' REUSE TO GROUP 3;
ALTER DATABASE ADD LOGFILE MEMBER '/dev/rdsk/2X' REUSE TO GROUP 4;
ALTER DATABASE ADD LOGFILE MEMBER '/dev/rdsk/2Y' REUSE TO GROUP 5;
ALTER DATABASE ADD LOGFILE MEMBER '/dev/rdsk/2Z' REUSE TO GROUP 6;
```

10. Verify that the log files actually got added by executing the following SQLDBA commands from any instance:

```
connect internal
select * from v$logfile
```

The newly added log files might show a status of INVALID.

11. Cycle through all the log groups by executing the following from each instance:

```
ALTER SYSTEM SWITCH LOGFILE;
ALTER SYSTEM SWITCH LOGFILE;
ALTER SYSTEM SWITCH LOGFILE;
```

5.2 Replacing Cables

This section contains procedures on replacing cables:

- replacing a cable on the SCSI chain with the RAID
- replacing network cables or serial lines
- replacing a cable on the SCSI chain with the log disks

5.2.1 Replacing a Cable on the SCSI Chain With the RAID

To replace a nonfunctional cable connecting the IO4 board in OPS1 (I1) to SP A, follow these steps:

1. Power down node 1.
2. Turn off SP A.
3. Replace the cable from A1 to I1 (host 1).
4. Turn on SP A.

5.2.2 Replacing Network Cables or Serial Lines

Follow these steps:

1. Back up and shut down Oracle.
2. Stop the CM:


```
killall -TERM opscm
```
3. Replace the defective cable.
4. Bring up the CM and Oracle.

5.2.3 Replacing a Cable on the SCSI Chain With the Log Disks

In this section, terms are used as follows:

- log disks: 1X 1Y 1Z are the log disks primarily written by node 1 and mirrored on 1x 1y 1z. Similarly, 2X 2Y 2Z are the log disks primarily written by node 2 and mirrored on 2x 2y 2z.
- vaults: V and v, containing 1X 1Y 1Z 2X 2Y 2Z and 1x 1y 1z 2x 2y 2z, respectively. 1X 1Y 1Z should be on one SCSI bus within V; 2X 2Y 2Z should be on the other SCSI bus of V. Similarly for the mirror disks.

To replace a nonfunctional cable from the IO4 board in OPS2 (I2) to vault 1 (V), follow these steps:

1. Follow steps 1 through 7 in Section 5.1.2, “Breaking the SCSI Chain With the Log Disks,” earlier in this chapter.
2. Replace the cable between OPS2 and vault V.
3. Power on vault V.
4. Do a dummy I/O on all the disks in the vault from both nodes.

```
dd if=/dev/rdsk/1X of=/dev/null count=1
dd if=/dev/rdsk/1Y of=/dev/null count=1
dd if=/dev/rdsk/1Z of=/dev/null count=1
dd if=/dev/rdsk/2X of=/dev/null count=1
dd if=/dev/rdsk/2Y of=/dev/null count=1
dd if=/dev/rdsk/2Z of=/dev/null count=1
```

It is OK if some of the dd's return an error.

5. Repeat the dd's from both nodes.

```
dd if=/dev/rdsk/1X of=/dev/null count=1
dd if=/dev/rdsk/1Y of=/dev/null count=1
dd if=/dev/rdsk/1Z of=/dev/null count=1
dd if=/dev/rdsk/2X of=/dev/null count=1
dd if=/dev/rdsk/2Y of=/dev/null count=1
dd if=/dev/rdsk/2Z of=/dev/null count=1
```

This time all the dd's should succeed.

6. Verify the current log files by executing the following SQLDBA statement from any node.

```
connect internal
select * from v$logfile
```

7. Add the log files back to Oracle by executing the following SQLDBA commands from any node:

```
ALTER DATABASE ADD LOGFILE MEMBER '/dev/rdsk/1X' REUSE TO GROUP 1;
ALTER DATABASE ADD LOGFILE MEMBER '/dev/rdsk/1Y' REUSE TO GROUP 2;
ALTER DATABASE ADD LOGFILE MEMBER '/dev/rdsk/1Z' REUSE TO GROUP 3;
ALTER DATABASE ADD LOGFILE MEMBER '/dev/rdsk/2X' REUSE TO GROUP 4;
ALTER DATABASE ADD LOGFILE MEMBER '/dev/rdsk/2Y' REUSE TO GROUP 5;
ALTER DATABASE ADD LOGFILE MEMBER '/dev/rdsk/2Z' REUSE TO GROUP 6;
```

8. Verify that the log files actually got added by executing the following SQLDBA commands from any node:

```
connect internal
select * from v$logfile
```

The newly added log files may show a status of INVALID.

9. Cycle through all the log groups by executing the following from each instance:

```
ALTER SYSTEM SWITCH LOGFILE;
ALTER SYSTEM SWITCH LOGFILE;
ALTER SYSTEM SWITCH LOGFILE;
```

5.3 Replacing a Log Disk

To replace a nonfunctional log disk (1X), follow these steps:

1. Follow steps 1 through 6 in Section 5.1.2, “Breaking the SCSI Chain With the Log Disks,” earlier in this chapter.
2. Replace the nonfunctional disk in vault V.
3. Power on vault V.
4. Do a dummy I/O on all the disks in the vault from both nodes.

```
dd if=/dev/rdsk/1X of=/dev/null count=1
dd if=/dev/rdsk/1Y of=/dev/null count=1
dd if=/dev/rdsk/1Z of=/dev/null count=1
dd if=/dev/rdsk/2X of=/dev/null count=1
dd if=/dev/rdsk/2Y of=/dev/null count=1
dd if=/dev/rdsk/2Z of=/dev/null count=1
```

It is OK if some of the dd's return an error.

5. Repeat the dd's from both nodes.

```
dd if=/dev/rdsk/1X of=/dev/null count=1
dd if=/dev/rdsk/1Y of=/dev/null count=1
dd if=/dev/rdsk/1Z of=/dev/null count=1
dd if=/dev/rdsk/2X of=/dev/null count=1
dd if=/dev/rdsk/2Y of=/dev/null count=1
dd if=/dev/rdsk/2Z of=/dev/null count=1
```

This time all the dd's should succeed.

6. Verify the current log files by executing the following SQLDBA statement from any instance:

```
connect internal
select * from v$logfile
```

7. Add the LOG files back to Oracle by executing the following SQLDBA commands from any instance:

```
ALTER DATABASE ADD LOGFILE MEMBER '/dev/rdsk/1X' REUSE TO GROUP 1;
ALTER DATABASE ADD LOGFILE MEMBER '/dev/rdsk/1Y' REUSE TO GROUP 2;
ALTER DATABASE ADD LOGFILE MEMBER '/dev/rdsk/1Z' REUSE TO GROUP 3;
ALTER DATABASE ADD LOGFILE MEMBER '/dev/rdsk/2X' REUSE TO GROUP 4;
ALTER DATABASE ADD LOGFILE MEMBER '/dev/rdsk/2Y' REUSE TO GROUP 5;
ALTER DATABASE ADD LOGFILE MEMBER '/dev/rdsk/2Z' REUSE TO GROUP 6;
```

8. Verify that the log files actually got added by executing the following SQLDBA commands from any instance:

```
connect internal
select * from v$logfile
```

The newly added log files might show a status of INVALID.

9. Cycle through all the log groups by executing the following from each instance:

```
ALTER SYSTEM SWITCH LOGFILE;
ALTER SYSTEM SWITCH LOGFILE;
ALTER SYSTEM SWITCH LOGFILE;
```

5.4 Checking OPS Server and Challenge RAID SCSI IDs

This section explains how to make sure that the two OPS servers (or other SCSI devices on the same bus) do not have the same SCSI ID number. Follow these steps:

1. On the first OPS server, enter the command *nvrnm*. This command's output should contain the line

```
scsihostid=0
```

This output means that the SCSI host ID for this OPS server is 0.

2. Repeat step 1 in the second window, for the second OPS server. The output should contain the line

```
scsihostid=2
```

This output means that the SCSI host ID for the second OPS server is 2.

3. If the host SCSI IDs are not 0 and 2, respectively, reboot and start the System Maintenance menu; choose item 5, the Command Monitor.

4. In the Command Monitor, set the SCSI ID for this system:

```
setenv scsihostid 2
```

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

```
printenv
```

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

6. Exit the System Maintenance menu and restart the system.
7. Enter the command *hinv* on each server and compare the output.

Note: No SCSI device on any bus on the second server (OPS2) should have SCSI ID 2.

Table 5-1 charts SCSI ID switch settings for the Challenge RAID storage system. Note that the settings do not conform to frequently used numbering schemes.

Table 5-1 Challenge RAID SCSI ID Switch Settings

SCSI ID Number	Switch Number				Comment
	ID 0	ID 1	ID 2	ID 3	
1	Off	On	On	On	
2	On	Off	On	On	Do not use
3	Off	Off	On	On	
4	On	On	Off	On	
5	Off	On	Off	On	
6	On	Off	Off	On	Shown in Figure 2-9 and Figure 3-10 as A
7	Off	Off	Off	On	Shown in Figure 2-9 and Figure 3-10 as B
8	On	On	On	Off	
9	Off	On	On	Off	
10	On	Off	On	Off	
11	Off	Off	On	Off	
12	On	On	Off	Off	
13	Off	On	Off	Off	
14	On	Off	Off	Off	
15	Off	Off	Off	Off	

5.5 Checking Disk Partitions

Never create disk partitions 8, 9, and 10 manually. If partition 9 on the drive has been put on the drive by Silicon Graphics or during the installation process, do not reuse it or the sectors that it spans. Thus, 13 usable partitions per LUN remain: 0, 1, 2, 3, 4, 5, 6, 7, 11, 12, 13, 14, and 15.

5.6 Exchanging Challenge RAID Disk Modules

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

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

- 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 5.6.2.4, “Updating the Disk Module Firmware.”

Figure 5-7 diagrams the location of disks, which can be replaced by owners or SSEs.




Figure 5-7 Location of Disks (Front of Challenge RAID)

Note: For information on replacing other field-replaceable units (FRUs) in the Challenge RAID, 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 explains

- opening and closing the fan module
- replacing or adding a disk module
- installing an add-on disk module array

5.6.1 Opening and Closing the Fan Module

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

1. On the back of the storage system, move the fan module's latch to the **UNLOCK** position, as shown in Figure 5-8.




Figure 5-8 Unlocking the Fan Module

2. Swing open the fan module, as shown in Figure 5-9.




Figure 5-9 Opening the Fan Module

Close the fan module by reversing these steps.

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

5.6.2 Replacing or Adding a Disk Module

This section explains

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

5.6.2.1 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, as explained in Chapter 3 in this guide, you can replace the defective module and rebuild your data without powering off the Challenge RAID storage system or interrupting user applications.

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

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

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

1. Look for the module with its amber fault light on. Figure 5-10 shows the fault indicator light and other lights on a disk module.




Figure 5-10 Disk Module Status Lights

2. Determine the failed module's ID; see Figure 5-7.

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

3. If you have not already checked the module status with *raidcli getdisk*, do so now. See Appendix B of the *CHALLENGE RAID Installation and Maintenance Instructions* (007-0128-006) for complete information on *raidcli* parameters.
4. If you have not already checked the unsolicited error log with *raidcli getlog* for a message about the disk module, as explained in Appendix B of the *CHALLENGE RAID Installation and Maintenance Instructions* (007-0128-006), do so now.

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

Note: If you are using storage system caching, the system uses modules A0, B0, C0, D0, and E0 for its cache vault. If one of these modules fails, the storage system dumps its cache image to the remaining modules in the vault; then it writes all dirty (modified) pages to disk and disables caching. The cache status changes, as indicated in the output of the *raidcli getcache* command. Caching remains disabled until you insert a replacement module and the storage system rebuilds the module into the physical disk unit. For information on caching, see Chapter 7 of the *CHALLENGE RAID Owner's Guide* (007-2523-007).

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 the correct disk modules available from Silicon Graphics.

5.6.2.2 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 that the module is being replaced.

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

To remove a disk module, follow these steps:

1. Verify that the suspected module has actually failed.

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

2. When a disk fails, the 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*. See Appendix B for information on these parameters.
3. Locate the disk module that you want to remove; see Figure 5-7 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 5-11. Put the wrist band around your wrist with the metal button against your skin.




Figure 5-11 Attaching the ESD Clip to the ESD Bracket on the Storage System
 Figure 5-12 shows where to attach the clip on a rack storage system.




Figure 5-12 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. Make sure that the disk has stopped spinning and that the heads have unloaded.

8. Grasp the disk module by its handle and pull it partway out of the cabinet, as shown in Figure 5-13.




Figure 5-13 Pulling Out a Disk Module

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



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

9. Supporting the disk module with your free hand, pull it all the way out of the cabinet, as shown in Figure 5-14.




Figure 5-14 Removing a Disk Module

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

10. If 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 5-7.

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

Caution: Before installing a replacement module, wait at least 15 seconds after removing the failed module to allow the storage-control 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.

5.6.2.3 Installing a Replacement Disk Module

To install the replacement disk module, follow these steps:

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

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

2. Before installing a replacement module, wait at least 15 seconds after removing the failed module to allow the 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.
3. Position the disk module in its antistatic packaging within reach of the storage system.
4. Locate the slot where you will install the disk module; see Figure 5-7.
5. Engage the disk module's rail in the chassis rail slot, as shown in Figure 5-15.




Figure 5-15 Engaging the Disk Module Rail

- Engage the disk module's guide in the chassis guide slot, as shown in Figure 5-16.




Figure 5-16 Engaging the Disk Module Guide

- Insert the disk module, as shown in Figure 5-17. Make sure it is completely seated in the slot.




Figure 5-17 Inserting the Replacement Disk Module

- Remove and store the ESD wrist band, if you are using one.
- 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* (007-0128-006).

5.6.2.4 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 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.

This command has no output.

5.6.3 Installing an Add-On Disk Module Array

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

Table 5-2 gives marketing codes for add-on disk module arrays. Add-on disk modules are available only in arrays of five.

Table 5-2 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 5-7 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.

Follow these steps:

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 5-11. 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 5-7.



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

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.
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 5-15.
Caution: Disk modules are extremely sensitive to shock and vibration. Even a slight jar can severely damage them.
7. Engage the disk module's guide in the chassis guide slot, as shown in Figure 5-16.
8. Insert the disk module, as shown in Figure 5-16. Make sure it is completely seated in the slot.

5.7 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:

- checking XLV volumes
- rebalancing LUNs across two SPs

Note: Although IRIX 6.4 includes the failover feature, IRIX 6.2 requires a patch.

5.7.1 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:

1. The RAID system is dual-SP.
2. The access to the RAID is through XLV volumes.
3. The volumes have been assembled correctly.
4. Both SPs are operational at the time that one fails (or is turned off).
5. The default ownership of the LUNs is distributed across the two SPs (this is not a strict requirement, but is a useful simplification).

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
```

For systems running under IRIX 6.2, use `xlv_mgr` to check for dual connectivity to RAID LUNs.

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)
```

5.7.2 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.

5.7.2.1 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) for more information on *raidcli trespass*; */sbin/trespass* has identical options.

5.7.2.2 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 5.7.1, “Checking XLV Volumes”). Use this command to switch to the other path; for example:

```
scsifo -s /hw/scsi_ctlr/7/target/6/lun/0
```


Index

C

cables, replacing, 5-5 through 5-7

CHALLENGE RAID

agent

and config file, 4-10

stopping, restarting, 4-10

breaking SCSI chain with, 5-2 through 5-3

config file, 4-10

exchanging disk modules, 5-10 through 5-21

fan module, unlocking, 2-27, 3-26

hardware installation, 2-11 through 2-16, 3-13 through 3-17

replacing cable on SCSI chain with, 5-5

SP, powering on, 2-27, 3-26

storage-control processor, powering on, 2-27, 3-26

CHALLENGE server

cabling to multiplexer, 3-20 through 3-24

SCSI ID, 3-11

command tag queuing, 4-13 through 4-14

conventions, xv

D

disk module

A0, B0, C0, D0, E0

failed, 5-14

removing, 5-10, 5-20

adding, 5-13

add-on array, 5-20 through 5-21

installing, 5-20 through 5-21

order, 5-20

ordering, 5-20

identifying failed, 5-13 through 5-14

installing replacement, 5-17 through 5-18

removing failed, 5-14 through 5-17

replacing, 5-13 through 5-19

status lights, 5-13

documentation required, xiv through xv

E

electrostatic discharge, 2-10, 3-12

ESD preventive measures, 2-10, 3-12

F

fan module in CHALLENGE RAID

locking, 2-27, 3-26, 5-12

unlocking, 2-27, 3-25, 5-12

FDDI

boards, installing, 3-12

cables, replacing, 5-5

fx, 4-13 through 4-14

H

hardware

optional, 1-2

see also OPS hardware

I

Inst, 4-9

IRISconsole multiplexer

cabling, 2-17 through 2-25, 3-18 through 3-24

kit contents, 1-2

isolating a node, 5-1 through 5-5

K

kit contents, 1-1 through 1-2

L

log disk

- breaking SCSI chain with, 5-3 through 5-5
- replacing, 5-7 through 5-8
- replacing cable on SCSI chain with, 5-6 through 5-7

LUN, restoring ownership, 5-21 through 5-23

M

MAKE_VLUNS, 4-12

N

*nvr*am to determine host SCSI ID, 2-9, 3-11

O

OPS

- components, 1-1 through 1-2
- software, xiii
- hardware
 - installing, 2-1 through 2-27, 3-1 through 3-24
 - cabling
 - CHALLENGE RAID, 2-11 through 2-16, 3-13 through 3-17
 - CHALLENGE server, 3-2 through 3-24
 - multiplexer, 2-17 through 2-25, 3-18 through 3-24
 - Origin server, 2-2 through 2-25
 - planning, 2-3 through 2-5, 3-4 through 3-7
 - vaults, 2-5 through 2-8, 3-8 through 3-11
 - workstation, 2-17 through 2-18, 3-18 through 3-20
 - Ethernet boards, 2-10
 - FDDI boards, 3-12
 - host SCSI IDs, 2-9, 3-11

required equipment, 2-2, 3-1 through 3-2

powering on, 2-26 through 2-27, 3-25 through 3-26

setting up, 2-2 through 2-3, 3-2 through 3-3

testing installation, 2-28, 3-26 through 3-27

host, *see* Origin server, CHALLENGE server

kit contents, 1-1 through 1-2

software

- installing, 4-1 through 4-18

Origin server

cabling to multiplexer, 2-21 through 2-25

SCSI ID, 2-9

R

RAID, *see* CHALLENGE RAID

raid5-agent.config, 4-10

rebuild time, *raid5 bind*, 5-19

Remote System Control port

firmware, installing, 3-4

upgrade

installing, 3-4

part number, 3-2

required documentation, xiv through xv

S

SCIP mezzanine cards, installing, 3-4

SCSI

connectors on CHALLENGE RAID chassis, 2-13, 3-14

ID

CHALLENGE RAID, 2-11, 3-13, 5-9

CHALLENGE server, 3-11

checking, 5-8 through 5-9

multiplexer, 2-18, 3-20

OPS server, 5-8 through 5-9

Origin server, 2-9

scsifo, 5-22, 5-23

software

optional, 1-2

see also OPS software

T

trespass, 5-23

W

warnings, 2-10, 3-12

X

XLV

and restoring LUN ownership, 5-21 through 5-23

creating volumes, 4-14 through 4-15

volume limit, 4-14

