



Sun Cluster 3.0 Hardware Guide

Sun Microsystems, Inc.
901 San Antonio Road
Palo Alto, CA 94303-4900
U.S.A. 650-960-1300

Part Number 806-1420
November 2000, Revision A

Copyright Copyright 2000 Sun Microsystems, Inc. 901 San Antonio Road, Palo Alto, California 94303-4900 U.S.A. All rights reserved.

This product or document is distributed under licenses restricting its use, copying, distribution, and decompilation. No part of this product or document may be reproduced in any form by any means without prior written authorization of Sun and its licensors, if any. Third-party software, including font technology, is copyrighted and licensed from Sun suppliers.

Parts of the product may be derived from Berkeley BSD systems, licensed from the University of California. UNIX is a registered trademark in the U.S. and other countries, exclusively licensed through X/Open Company, Ltd.

Sun, Sun Microsystems, the Sun logo, AnswerBook2, docs.sun.com, Solaris, Solstice DiskSuite, and OpenBoot are trademarks, registered trademarks, or service marks of Sun Microsystems, Inc. in the U.S. and other countries. All SPARC trademarks are used under license and are trademarks or registered trademarks of SPARC International, Inc. in the U.S. and other countries. Products bearing SPARC trademarks are based upon an architecture developed by Sun Microsystems, Inc.

The OPEN LOOK and Sun[™] Graphical User Interface was developed by Sun Microsystems, Inc. for its users and licensees. Sun acknowledges the pioneering efforts of Xerox in researching and developing the concept of visual or graphical user interfaces for the computer industry. Sun holds a non-exclusive license from Xerox to the Xerox Graphical User Interface, which license also covers Sun's licensees who implement OPEN LOOK GUIs and otherwise comply with Sun's written license agreements.

Federal Acquisitions: Commercial Software—Government Users Subject to Standard License Terms and Conditions.

DOCUMENTATION IS PROVIDED "AS IS" AND ALL EXPRESS OR IMPLIED CONDITIONS, REPRESENTATIONS AND WARRANTIES, INCLUDING ANY IMPLIED WARRANTY OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE OR NON-INFRINGEMENT, ARE DISCLAIMED, EXCEPT TO THE EXTENT THAT SUCH DISCLAIMERS ARE HELD TO BE LEGALLY INVALID.

Copyright 2000 Sun Microsystems, Inc., 901 San Antonio Road, Palo Alto, CA 94303-4900 Etats-Unis. Tous droits réservés.

Ce produit ou document est distribué avec des licences qui en restreignent l'utilisation, la copie, la distribution, et la décompilation. Aucune partie de ce produit ou document ne peut être reproduite sous aucune forme, par quelque moyen que ce soit, sans l'autorisation préalable et écrite de Sun et de ses bailleurs de licence, s'il y en a. Le logiciel détenu par des tiers, et qui comprend la technologie relative aux polices de caractères, est protégé par un copyright et licencié par des fournisseurs de Sun.

Des parties de ce produit pourront être dérivées des systèmes Berkeley BSD licenciés par l'Université de Californie. UNIX est une marque déposée aux Etats-Unis et dans d'autres pays et licenciée exclusivement par X/Open Company, Ltd.

Sun, Sun Microsystems, le logo Sun, AnswerBook2, docs.sun.com, Solaris, Solstice DiskSuite, et OpenBoot sont des marques de fabrique ou des marques déposées, ou marques de service, de Sun Microsystems, Inc. aux Etats-Unis et dans d'autres pays. Toutes les marques SPARC sont utilisées sous licence et sont des marques de fabrique ou des marques déposées de SPARC International, Inc. aux Etats-Unis et dans d'autres pays. Les produits portant les marques SPARC sont basés sur une architecture développée par Sun Microsystems, Inc.

L'interface d'utilisation graphique OPEN LOOK et Sun[™] a été développée par Sun Microsystems, Inc. pour ses utilisateurs et licenciés. Sun reconnaît les efforts de pionniers de Xerox pour la recherche et le développement du concept des interfaces d'utilisation visuelle ou graphique pour l'industrie de l'informatique. Sun détient une licence non exclusive de Xerox sur l'interface d'utilisation graphique Xerox, cette licence couvrant également les licenciés de Sun qui mettent en place l'interface d'utilisation graphique OPEN LOOK et qui en outre se conforment aux licences écrites de Sun.

LA DOCUMENTATION EST FOURNIE "EN L'ETAT" ET TOUTES AUTRES CONDITIONS, DECLARATIONS ET GARANTIES EXPRESSES OU TACITES SONT FORMELLEMENT EXCLUES, DANS LA MESURE AUTORISEE PAR LA LOI APPLICABLE, Y COMPRIS NOTAMMENT TOUTE GARANTIE IMPLICITE RELATIVE A LA QUALITE MARCHANDE, A L'APTITUDE A UNE UTILISATION PARTICULIERE OU A L'ABSENCE DE CONTREFAÇON.



Contents

Preface 7

1. Introduction to Sun Cluster Hardware 13

Installing Cluster Hardware 13

Maintaining Sun Cluster Hardware 15

Powering On and Powering Off Cluster Hardware 15

Local and Multihost Disks 16

Removable Media 16

2. Installing and Configuring the Terminal Concentrator 17

Installing and Configuring the Terminal Concentrator 17

Installing the Terminal Concentrator 18

▼ How to Install the Terminal Concentrator in a Cabinet 19

▼ How to Cable the Terminal Concentrator 23

Configuring the Terminal Concentrator 24

▼ How to Configure the Terminal Concentrator 25

▼ How to Set Terminal Concentrator Port Parameters 27

▼ How to Correct a Port Configuration Access Error 29

▼ How to Establish a Default Route for the Terminal Concentrator 31

Using the Terminal Concentrator 33

▼ How to Connect to a Node's Console Through the Terminal Concentrator 34

▼ How to Reset a Terminal Concentrator Port	35
3. Installing and Maintaining Cluster Interconnect and Public Network Hardware	39
Installing and Maintaining Cluster Interconnect and Public Network Hardware	40
Installing Cluster Transport Adapters	40
▼ How to Install Cluster Transport Adapters	41
Replacing Cluster Transport Adapters	41
▼ How to Replace Cluster Transport Adapters	42
Removing Cluster Transport Adapters	43
▼ How to Remove Cluster Transport Adapters	43
Installing Cluster Transport Junctions	44
Installing Cluster Transport Cables	45
▼ How to Install Cluster Transport Cables	45
Adding Public Network Connections	46
▼ How to Add Public Network Connections	46
Removing Public Network Connections	47
▼ How to Remove Public Network Connections	47
4. Installing and Maintaining the Sun StorEdge MultiPack Enclosure	49
Installing a StorEdge MultiPack	49
▼ How to Install a StorEdge MultiPack	50
Maintaining a StorEdge MultiPack	54
▼ How to Add a StorEdge MultiPack Disk Drive	55
▼ How to Replace a StorEdge MultiPack Disk Drive	57
▼ How to Remove a StorEdge MultiPack Disk Drive	61
▼ How to Add a StorEdge MultiPack Enclosure	62
▼ How to Replace a StorEdge MultiPack Enclosure	68
▼ How to Remove a StorEdge MultiPack Enclosure	69
5. Installing and Maintaining a Sun StorEdge D1000 Disk Array	73

Installing a StorEdge D1000	73
▼ How to Install a StorEdge D1000	74
Maintaining a StorEdge D1000 Disk Array	79
▼ How to Add a StorEdge D1000 Disk Drive	80
▼ How to Replace a StorEdge D1000 Disk Drive	82
▼ How to Remove a StorEdge D1000 Disk Drive	86
▼ How to Add a StorEdge D1000 Disk Array	87
▼ How to Replace a StorEdge D1000 Disk Array	93
▼ How to Remove a StorEdge D1000 Disk Array	95
6. Installing and Maintaining a Sun StorEdge A5x00 Disk Array	99
Installing a StorEdge A5x00	99
▼ How to Install a StorEdge A5x00	100
Maintaining a StorEdge A5x00 Disk Array	101
▼ How to Add a StorEdge A5x00 Disk Drive	102
▼ How to Replace a StorEdge A5x00 Disk Drive	104
▼ How to Remove a StorEdge A5x00 Disk Drive	108
▼ How to Add a StorEdge A5x00 Disk Array to an Existing Cluster	109
▼ How to Replace a StorEdge A5x00 Disk Array	112
▼ How to Remove a StorEdge A5x00 Disk Array	114
7. Installing, Configuring, and Maintaining a Sun StorEdge A3500 Disk Array	117
Installing a StorEdge A3500	117
▼ How to Install a StorEdge A3500	118
Configuring StorEdge A3500 Disk Drives After Sun Cluster Install	123
▼ How to Create a LUN	125
▼ How to Remove a LUN	126
▼ How to Reset StorEdge A3500 Configuration	126
Maintaining a StorEdge A3500 Disk Array	127

▼	How to Add a StorEdge A3500	130
▼	How to Remove a StorEdge A3500	136
▼	How to Replace a Failed StorEdge A3500 Controller or Restore an Offline StorEdge A3500 Controller	137
	Replacing a Host Adapter	138
▼	How to Replace a Host Adapter	139
A.	Verifying Sun Cluster Hardware Redundancy	141
	Testing Node Redundancy	142
▼	How to Test Nodes Using a Power-off Method	142
	Testing Cluster Interconnect and Network Adapter Failover Group Redundancy	143
▼	How to Test Cluster Interconnects	143
▼	How to Test Network Adapter Failover Groups	144
B.	NVRAMRC Editor and NVEDIT Keystroke Commands	147

Preface

Sun Cluster 3.0 Hardware Guide provides the procedures for installing and servicing Sun™ Cluster hardware.

This document is intended for experienced system administrators with extensive knowledge of Sun software and hardware. This document is not to be used as a planning or presales guide. Determine your system requirements and purchase the appropriate equipment and software before reading this document.

All the procedures in this document require root-level permission.

Using UNIX Commands

This document may not contain information on basic UNIX® commands and procedures such as shutting down the system, booting the system, and configuring devices.

See one or more of the following for this information:

- AnswerBook2™ online documentation for the Solaris™ software environment
- Other software documentation that you received with your system
- Solaris operating environment man pages

Typographic Conventions

Typeface or Symbol	Meaning	Examples
AaBbCc123	The names of commands, files, and directories; on-screen computer output	Edit your <code>.login</code> file. Use <code>ls -a</code> to list all files. % You have mail.
AaBbCc123	What you type, when contrasted with on-screen computer output	% su Password:
<i>AaBbCc123</i>	Book titles, new words or terms, words to be emphasized	Read Chapter 6 in the <i>User's Guide</i> . These are called <i>class</i> options. You <i>must</i> be superuser to do this.
	Command-line variable; replace with a real name or value	To delete a file, type <code>rm filename</code> .

Shell Prompts

Shell	Prompt
C shell	<i>machine_name%</i>
C shell superuser	<i>machine_name#</i>
Bourne shell and Korn shell	\$
Bourne shell and Korn shell superuser	#

Related Documentation

Application	Title	Part Number
Concepts	<i>Sun Cluster 3.0 Concepts</i>	806-1424
Software Installation	<i>Sun Cluster 3.0 Installation Guide</i>	806-1419
Data Services	<i>Sun Cluster 3.0 Data Services Installation and Configuration Guide</i>	805-1421
API Development	<i>Sun Cluster 3.0 Data Services Developer's Guide</i>	805-1422
System Administration	<i>Sun Cluster 3.0 System Administration Guide</i>	806-1423
Sun Cluster Release Notes	<i>Sun Cluster 3.0 Release Notes</i>	805-1428
Sun StorEdge MultiPack Installation	<i>Sun StorEdge MultiPack Installation Guide</i>	805-3953
Sun StorEdge MultiPack Usage	<i>Sun StorEdge MultiPack User's Guide</i>	805-3954
Sun StorEdge MultiPack Hot-Plugging	<i>Sun StorEdge MultiPack Storage Guide</i>	805-3955
Sun StorEdge D1000 Storage	<i>Sun StorEdge D1000 Storage Guide</i>	805-4013
Sun StorEdge D1000 Installation	<i>Sun StorEdge A1000 and D1000 Installation, Operations, and Service Manual</i>	805-2624
Sun StorEdge D1000 Product Note	<i>Sun StorEdge A1000 and D1000 Product Note</i>	805-4866
Sun StorEdge D1000 Rackmount Installation	<i>Sun StorEdge A1000 and D1000 Rackmount Installation Manual</i>	805-2626

Application	Title	Part Number
Sun StorEdge A5x00 Installation and Documentation	<i>Sun StorEdge A5x00 Installation and Documentation Guide</i>	805-1903
Sun StorEdge A5x00 Installation and Service	<i>Sun StorEdge A5x00 Installation and Service Manual</i>	802-7573
Sun StorEdge A5x00 Hardware Configuration	<i>Sun StorEdge A5x00 Hardware Configuration Guide</i>	805-0264
Sun StorEdge RAID Manager Installation	<i>Sun StorEdge RAID Manager Installation and Support Guide</i>	805-7756
Sun StorEdge RAID Manager Release Notes	<i>Sun StorEdge RAID Manager Release Notes</i>	805-7758
Sun StorEdge RAID Manager Usage	<i>Sun StorEdge RAID Manager User's Guide</i>	806-0478
Sun StorEdge RAID Manager Hardware Configuration	<i>Sun StorEdge A3500/A3500FC Hardware Configuration Guide</i>	805-4981
Sun StorEdge A3500 Controller Module	<i>Sun StorEdge A3500/A3500FC Controller Module Guide</i>	805-4980
NVEDIT Editor and Keystroke Commands	<i>OpenBoot 3.x Command Reference Manual</i>	802-3242

Ordering Sun Documentation

Fatbrain.com, an Internet professional bookstore, stocks select product documentation from Sun Microsystems, Inc.

For a list of documents and how to order them, visit the Sun Documentation Center on Fatbrain.com at:

<http://www1.fatbrain.com/documentation/sun>

Accessing Sun Documentation Online

The docs.sun.comSM web site enables you to access Sun technical documentation on the Web. You can browse the docs.sun.com archive or search for a specific book title or subject at:

<http://docs.sun.com>

Getting Help

If you have problems installing or using Sun Cluster, contact your service provider and provide the following information:

- Your name and email address (if available)
- Your company name, address, and phone number
- The model and serial numbers of your systems
- The release number of the operating environment (for example, Solaris 8)
- The release number of Sun Cluster (for example, Sun Cluster 3.0)

Use the following commands to gather information on your system for your service provider:

Command	Function
<code>prtconf -v</code>	Displays the size of the system memory and reports information about peripheral devices
<code>psrinfo -v</code>	Displays information about processors
<code>showrev --p</code>	Reports which patches are installed
<code>prtdiag -v</code>	Displays system diagnostic information
<code>scinstall -pv</code>	Displays Sun Cluster release and package version information

Also have available the contents of the `/var/adm/messages` file.

Introduction to Sun Cluster Hardware

This chapter provides overview information on cluster hardware, including the terminal concentrator, storage devices, and network components.

This chapter contains the following conceptual information:

- “Installing Cluster Hardware” on page 13
- “Maintaining Sun Cluster Hardware” on page 15
- “Powering On and Powering Off Cluster Hardware” on page 15
- “Local and Multihost Disks” on page 16

Installing Cluster Hardware

The following table lists the tasks for building a cluster.

TABLE 1-1 Task Map: Installing Cluster Hardware

Task	For Instructions, Go To...
Plan for cluster hardware capacity, space, and power requirements	The site planning documentation that shipped with your nodes and other hardware
Install the nodes	The documentation that shipped with your nodes

TABLE 1-1 Task Map: Installing Cluster Hardware *(continued)*

Task	For Instructions, Go To...
Install and configure the terminal concentrator	"Installing and Configuring the Terminal Concentrator" on page 17
Install the cluster transport adapters	"How to Install Cluster Transport Adapters" on page 41
Install the cluster transport junction	"Installing Cluster Hardware" on page 13
Install the cluster transport cables	"How to Install Cluster Transport Cables" on page 45
Install public network hardware	The documentation that shipped with your network adapters and nodes
Install and configure the storage	"Installing a StorEdge MultiPack" on page 49
	"Installing a StorEdge D1000" on page 73
	"Installing a StorEdge A5x00" on page 99
	"Installing a StorEdge A3500" on page 117
Install the Solaris operating environment and Sun Cluster software	<i>Sun Cluster 3.0 Installation Guide</i>
Configure the cluster interconnects	<i>Sun Cluster 3.0 System Administration Guide</i>

Maintaining Sun Cluster Hardware

Sun Cluster 3.0 Hardware Guide augments existing documentation that ships with your hardware components by providing information on maintaining this hardware in a cluster environment. The following table describes some of the differences between maintaining cluster hardware as compared to standalone hardware.

TABLE 1-2 Sample Differences Between Servicing Standalone and Cluster Hardware

Task	Standalone Hardware	Cluster Hardware
Shutting down a node	Use the shutdown(1M) command.	To perform an orderly node shutdown, first use the scswitch(1M) command to switch device groups and resource groups to another node. Then shut down the node by running the shutdown(1M) command.
Adding a disk	Run boot -r or devfsadm(1M) to assign a logical device name to the disk. You also need to run volume manager commands to configure the new disk if the disks are under volume management control	Use the devfsadm(1M), scgdevs(1M), and sccdadm(1M) commands. You also need to run volume manager commands to configure the new disk if the disks are under volume management control.
Adding a public network connection	To install the network adapter, perform an orderly node shutdown. After you install the network adapter, update the /etc/hostname.adapter and /etc/inet/hosts files.	To install the network adapter, perform an orderly node shutdown. After you install the network adapter, update the /etc/hostname.adapter and /etc/inet/hosts files. Finally, configure the adapter as part of a NAFO group.

Powering On and Powering Off Cluster Hardware

Consider the following when powering on and powering off cluster hardware:

- Use the power-on and power-off procedures in the manuals that shipped with the hardware *only* for systems that are newly installed or in the process of being installed.
- Use power-on and power-off procedures in *Sun Cluster 3.0 System Administration Guide* for nodes in an existing cluster.



Caution - After the cluster is online and a user application is accessing data on the cluster, do not use the power-on and power-off procedures listed in the manuals that came with the hardware.

Local and Multihost Disks

There are two sets of storage devices within a cluster: local disks and multihost disks. Local disks are directly connected to a single node and hold the Solaris operating environment for each node. Multihost disks are connected to more than one node and hold client application data and other files that need to be accessed from multiple nodes.

For more conceptual information on multihost disks, local disks, and global devices, see *Sun Cluster 3.0 Concepts*.

Removable Media

Removable media includes tape and CD-ROM drives. Tape and CD-ROM drives are local devices. This guide does not contain procedures on adding, removing, or replacing removable media as highly-available storage devices. Although tape and CD-ROM drives are global devices, tape and CD-ROM drives do not have more than one port and do not have multi-initiator firmware support. Dual-ported disks and multi-initiator firmware support enable devices to be highly available. Thus, this guide focuses on disk drives as global devices.

Although tape and CD-ROM drives cannot be highly available at this time, in a cluster environment, you can access tape and CD-ROM drives that are not local to your system. All the various density extensions (such as h, b, l, n, and u) are mapped so that the tape drive can be accessed from any node in the cluster.

Install, add, remove, replace, and use tape and CD-ROM drives as you would in a non-cluster environment. For procedures on adding, removing, and replacing tape and CD-ROM drives, see the documentation that shipped with your hardware.

Installing and Configuring the Terminal Concentrator

This chapter provides the hardware and software procedures for installing and configuring a terminal concentrator in a Sun Cluster environment.

This chapter contains the following procedures:

- “How to Install the Terminal Concentrator in a Cabinet” on page 19
- “How to Cable the Terminal Concentrator” on page 23
- “How to Configure the Terminal Concentrator” on page 25
- “How to Set Terminal Concentrator Port Parameters” on page 27
- “How to Correct a Port Configuration Access Error” on page 29
- “How to Establish a Default Route for the Terminal Concentrator” on page 31
- “How to Connect to a Node’s Console Through the Terminal Concentrator” on page 34
- “How to Reset a Terminal Concentrator Port” on page 35

For conceptual information on terminal concentrators, see *Sun Cluster 3.0 Concepts*.

Installing and Configuring the Terminal Concentrator

The following table lists the procedures for installing and configuring the terminal concentrator.

TABLE 2-1 Task Map: Installing and Configuring the Terminal Concentrator

Task	For Instructions, Go To...
Install the terminal concentrator	"How to Install the Terminal Concentrator in a Cabinet" on page 19
Cable the terminal concentrator	"How to Cable the Terminal Concentrator" on page 23
Configure the terminal concentrator	"How to Configure the Terminal Concentrator" on page 25
Set terminal concentrator port parameters	"How to Set Terminal Concentrator Port Parameters" on page 27
Correct a port configuration access error	"How to Correct a Port Configuration Access Error" on page 29
Establish a default route	"How to Establish a Default Route for the Terminal Concentrator" on page 31

Installing the Terminal Concentrator

This section covers the procedure for installing the terminal concentrator hardware and connecting cables from the terminal concentrator to the administrative console and the cluster nodes.

▼ How to Install the Terminal Concentrator in a Cabinet

This procedure provides step-by-step instructions for rack-mounting the terminal concentrator. For convenience, you can rack-mount the terminal concentrator even if your cluster does not contain rack-mounted nodes.

If you do not want to rack-mount your terminal concentrator, place the terminal concentrator in its standalone location, and connect the unit into a utility outlet.

1. Install the terminal concentrator bracket hinge in the primary cabinet.

a. Locate the hinge portion of the terminal concentrator bracket assembly.

b. Loosely install two locator screws in the right-side rail in the rear of the cabinet.

Thread the screws into holes 8 and 29, as shown in Figure 2–1. The locator screws will accept the slotted holes in the hinge piece.

c. Place the slotted holes of the hinge over the locator screws, and let the hinge drop into place.

d. Install the screws in holes 7 and 28.

Tighten these screws, as well as those in holes 8 and 29, as shown in Figure 2–1.




Figure 2-1 Installing the Terminal Concentrator Hinge

2. **Install the terminal concentrator in the bracket, and assemble the bracket.**
 - a. **Place the side pieces of the bracket against the terminal concentrator, as shown in Figure 2-2.**
 - b. **Lower the terminal concentrator (with side pieces) onto the bottom plate. Align the holes in the side pieces with those in the bottom plate.**




Figure 2-2 Assembling the Terminal Concentrator Bracket

- c. **Install and tighten three nuts on threaded studs penetrating through each side plate.**

This step completes the bracket assembly.

- 3. **Install the terminal concentrator bracket onto the hinge in the chassis.**
 - a. **Turn the terminal concentrator bracket on its side so the hinge holes and cable connectors face right.**
 - b. **Align the bracket holes with the boss pins in the hinge assembly, and drop the bracket onto the hinge, as shown in Figure 2-3.**

4. Install the keeper screw in the shorter boss pin to ensure the assembly cannot be accidentally knocked off the hinge.




Figure 2-3 Terminal Concentrator Installed on the Hinge

5. Connect the power cord to the terminal concentrator, as shown in Figure 2-4. Connect the other end to the power distribution unit.




Figure 2-4 Terminal Concentrator Cable Locations

Where to Go From Here

To cable the terminal concentrator, go to “How to Cable the Terminal Concentrator” on page 23.

▼ How to Cable the Terminal Concentrator

1. **Connect the serial cable (530-2152-01 or 530-2151-01) from serial port A on the administrative console to serial port 1 on the terminal concentrator, as shown in Figure 2-5.**

The cable connection from the administrative console enables you to configure the terminal concentrator. The connection serves no other purpose. You can remove this connection after you set up the terminal concentrator.




Figure 2-5 Connecting the Administrative Console

2. **Connect the nodes to the terminal concentrator. Use serial cables with an RJ-45 connector on one end and a DB-25 on the other.**

The cable connections from the concentrator to the nodes enable you to access the ok prompt or OpenBoot™ PROM (OBP) mode by using the Cluster Console (Console Mode) Windows from the Cluster Control Panel (CCP). For more information on using the CCP, see *Sun Cluster 3.0 System Administration Guide*.

3. **Connect the public network Ethernet cable to the appropriate connector on the terminal concentrator.**

Note - The terminal concentrator requires a 10-Mbit/sec Ethernet connection.

4. **Close the terminal concentrator bracket, and install screws in holes 8 and 29 on the left-side rail.**

Where to Go From Here

Go to “Configuring the Terminal Concentrator” on page 24.

Configuring the Terminal Concentrator

This section provides the procedure for configuring the terminal concentrator's network addresses and ports.

▼ How to Configure the Terminal Concentrator

1. From the administrative console, add the following entry to the `/etc/remote` file.

```
tc:\
:dv=/dev/term/a:br#9600:
```

2. Verify that the server and the terminal concentrator are powered on and that the cabinet keyswitch (if applicable) is in the ON position.
3. Establish a connection to the terminal concentrator's serial port.

```
# tip tc
```

- a. If your terminal concentrator is powered off, power on the terminal concentrator and, within five seconds of it being powered on, hold down the Test button for one second, then release the Test button.
- b. If your terminal concentrator is powered on, hold down the Test button until the power LED flashes (about three seconds), then release the Test button. For one second, hold down the Test button again, then release.

Figure 2-6 shows the location of the Test button.

The terminal concentrator performs a self-test, which lasts about 30 seconds. Messages display on the administrative console. If the network connection is not found, press the Q key to stop the message.




Figure 2-6 Terminal Concentrator Test Button

4. If the front panel LEDs light up as shown in Table 2–2, and the administrative console displays a `monitor:: prompt`, proceed to Step 5 on page 26. Otherwise, use Table 2–3 and the documentation that shipped with your terminal concentrator to troubleshoot.

TABLE 2-2 Front Panel LEDs: Indicating a Successful Boot or Monitor Mode Reset

Power (Green)	Unit (Green)	Net (Green)	Attn (Amber)	Load (Green)	Active (Green)	Test (Orange)
ON	ON	ON	ON	OFF	Intermittent blinking	ON

TABLE 2-3 Front Panel LEDs: Indicating a Failed Boot

Mode	Power (Green)	Unit (Green)	Net (Green)	Attn (Amber)	Load (Green)	Active (Green)
Hardware failure	ON	Blinking	OFF	Blinking	OFF	OFF
Network test failure	ON	ON	Blinking	OFF	OFF	Intermittent blinking
Network test aborted, or net command failed	ON	ON	OFF	Blinking	OFF	Intermittent blinking
Booted wrong image	ON	ON	ON	Blinking	OFF	OFF
Other failure	One or more Status LEDs (1-8) are ON					

5. Use the `addr` command to assign an IP address, subnet mask, and network address to the terminal concentrator.

In the following example (Class B network, Class C subnet), the broadcast address is the terminal concentrator's address with the host portion set to 255 (all binary 1's).

```
monitor:: addr
Enter Internet address [<uninitialized>]: 172.25.80.6
Internet address: 172.25.80.6
Enter Subnet mask [255.255.0.0]: 255.255.255.0
Subnet mask: 255.255.255.0
Enter Preferred load host Internet address [<any host>]: 172.25.80.6
*** Warning: Load host and Internet address are the same ***
Preferred load host address: 172.25.80.6
Enter Broadcast address [0.0.0.0]: 172.25.80.255
Broadcast address: 172.25.80.255
Enter Preferred dump address [0.0.0.0]: 172.25.80.6
Preferred dump address: 172.25.80.6
Select type of IP packet encapsulation (ieee802/ethernet) [<ethernet>]:
Type of IP packet encapsulation: <ethernet>
Load Broadcast Y/N [Y]: n
Load Broadcast: N
```

6. After the addr session, power-cycle the terminal concentrator.

The Load and Active lights should briefly blink, and then the Load light turns off.

7. Use the ping(1M) command to confirm that the network connection works.

8. Exit the tip utility by entering <Return> and typing a tilde followed by a period.

```
<Return>~.
~
[EOT]
#
```

Where to Go From Here

Go to “How to Set Terminal Concentrator Port Parameters” on page 27.

▼ How to Set Terminal Concentrator Port Parameters

This procedure explains how to determine if the port type variable must be set and how to set the variable.

The port type parameter must be set to `dial_in`. If the parameter is set to `hardwired`, the cluster console might be unable to detect when a port is already in use. The port type parameter must be set to `Y`.

1. Find the Sun label on the top panel of the terminal concentrator.

Figure 2-7 shows the Sun label.

2. Examine the serial number to see if it is in the lower serial number range. The serial number consists of 7 digits, followed by a dash and 10 more digits.

- If the numbers after the dash start with 9520 or higher, the port type variable is set correctly. Skip to Step 4 on page 29.
- If the numbers after the dash start with 9519 or lower, you must change the port type variable. Proceed to Step 3 on page 28.




Figure 2-7 Determining the Terminal Concentrator Version

3. Using an administrative console, change the port type variable by setting the port parameters, then reboot the terminal concentrator as shown in the following example.

The boot command causes the changes to take effect. The terminal concentrator will be unavailable for approximately one minute.

```
admin-ws# telnet tc_name
Trying terminal concentrator IP address
Connected to tc_name
Escape character is '^]'.
Rotaries Defined:
  cli
Enter Annex port name or number: cli
Annex Command Line Interpreter * Copyright 1991 Xylogics, Inc.
annex: su
```

(continued)

```

Password: password (default password is the terminal concentrator IP address)
annex# admin
Annex administration MICRO-XL-UX R7.0.1, 8 ports
admin : set port=1-8 type dial_in imask_7bits Y
      You may need to reset the appropriate port, Annex subsystem or
        reboot the Annex for changes to take effect.
admin : set port=1-8 mode slave
admin : quit
annex# boot
bootfile: <return>
warning:   <return>

```

Note - Ensure that the terminal concentrator is powered on and has completed booting.

4. Verify that you can log in from the administrative console to the consoles of each node.

For information on how to connect to the nodes' consoles, see "How to Connect to a Node's Console Through the Terminal Concentrator" on page 34.

▼ How to Correct a Port Configuration Access Error

A misconfigured port not accepting network connections may return a `Connect : Connection refused` message when you use `telnet(1)`. Use the following procedure to correct the port configuration.

1. Connect to the terminal concentrator without specifying a port.

```
# telnet tc_name
```

tc_name Specifies the hostname of the terminal concentrator

2. Press an extra Return after making the connection, and specify the port number.

```
Trying ip_address ..
Connected to 192.9.200.1
Escape character is "^]".
...
[RETURN]
Rotaries Defined:
    cli
Enter Annex port name or number: 2
```

- If you see a Port(s) busy, do you wish to wait? (y/n) message, answer **N** and go to “How to Reset a Terminal Concentrator Port” on page 35.
- If you see a Error: Permission denied message, the port mode is configured incorrectly to the command-line interface and must be set to slave. Go to Step 3 on page 30.

3. Select the terminal concentrator’s command-line interface.

```
...
Enter Annex port name or number: cli
annex:
```

4. Type the su command and password.

The default password is the terminal concentrator’s IP address.

```
annex: su
Password:
```

5. Reset the port.

```
annex# admin
Annex administration MICRO-XL-UX R7.0.1, 8 ports
admin: port 2
admin: set port mode slave
    You may need to reset the appropriate port, Annex subsystem or
    reboot the Annex for changes to take effect.
admin: reset 2
```

Example—Correcting a Terminal Concentrator Port Configuration Access Error

The following example shows how to correct an access error on the terminal concentrator port 2.

```
admin-ws# telnet tc1
Trying 192.9.200.1 ...
Connected to 192.9.200.1.
Escape character is '^]'.
[Return]
Enter Annex port name or number: cli
...
annex: su
Password: root_password
annex# admin
Annex administration MICRO-XL-UX R7.0.1, 8 ports
admin: port 2
admin: set port mode slave
You may need to reset the appropriate port, Annex subsystem or
reboot the Annex for changes to take effect.
admin: reset 2
```

▼ How to Establish a Default Route for the Terminal Concentrator

A routing table overflow in the terminal concentrator can cause network connections to be intermittent or lost altogether. Symptoms include connection timeouts and routes that are reestablished then disappear, even though the terminal concentrator itself has not rebooted.

The following procedure fixes this problem by establishing a default route within the terminal concentrator. To preserve the default route within the terminal concentrator, you must also disable the routed feature.

Note - The routing table overflow is not a problem for connections made from a host that resides on the same network as the terminal concentrator.

1. Connect to the terminal concentrator.

```
# telnet tc_name
```

tc_name Specifies the name of the terminal concentrator

2. Press an extra Return after making the connection, and select the command-line interface to connect to the terminal concentrator.

```
...
Enter Annex port name or number: cli
annex:
```

3. Type the su command and password.

The default password is the terminal concentrator's IP address.

```
annex: su
Password:
```

4. Start the editor to change the config.annex file.

```
annex# edit config.annex
```

The keyboard commands for this editor are Control-W:save and exit, Control-X:exit, Control-F:page down, and Control-B:page up.

The config.annex file, which is created in the terminal concentrator's EEPROM file system, defines the default route. The config.annex file can also define rotaries that enable a symbolic name to be used instead of a port number.

5. Add the following lines to the file.

Substitute the appropriate IP address for your default router.

```
%gateway
net default gateway 192.9.200.2 metric 1 active ^W
```

6. Disable the local routed feature.

```
annex# admin set annex routed n
```

7. Reboot the terminal concentrator.

```
annex# boot
bootfile: <reboot>
warning: <return>
```

While the terminal concentrator is rebooting, the node consoles are not accessible.

Example—Establishing a Default Route for the Terminal Concentrator

The following example shows how to establish a default route for the terminal concentrator.

```
admin-ws# telnet tc1
Trying 192.9.200.1 ...
Connected to 192.9.200.1.
Escape character is '^]'.
[Return]
Enter Annex port name or number: cli
...
annex: su
Password: root_password
annex: edit config.annex
(Editor starts)
Ctrl-W:save and exit Ctrl-X:exit Ctrl-F:page down Ctrl-B:page up
%gateway
net default gateway 192.9.200.2 metric 1 active ^W
annex# admin set annex routed n
You may need to reset the appropriate port, Annex subsystem or
reboot the Annex for changes to take effect.
annex# boot
```

Using the Terminal Concentrator

This section provides the procedures for using the terminal concentrator in a cluster.

TABLE 2-4 Task Map: Using the Terminal Concentrator

Task	For Instructions, Go To...
Connect to a node's console through the terminal concentrator	"How to Connect to a Node's Console Through the Terminal Concentrator" on page 34
Reset a terminal concentrator port	"How to Reset a Terminal Concentrator Port" on page 35

▼ How to Connect to a Node's Console Through the Terminal Concentrator

The following procedure enables remote connections from the administrative console to a cluster node's console by first connecting to the terminal concentrator.

1. Connect to a node by starting a session with the appropriate terminal concentrator port.

You need to know which terminal concentrator port the node is physically cabled to. See Step 2 on page 24 in "How to Cable the Terminal Concentrator" on page 23 for details on determining which terminal concentrator port the node is physically cabled to.

```
# telnet tc_name tc_port_number
```

tc_name Specifies the name of the terminal concentrator.

tc_port_number Specifies the port number on the terminal concentrator. Port numbers are configuration dependent. Typically, ports 2 and 3 (5002 and 5003) are used for the first cluster installed at a site.

Note - If you set up node security, you are prompted for the port password.

2. Log in to the node's console.

After establishing the connection, the system prompts for the login name and password.

3. Set the terminal type based on the type of window used in Step 1 on page 34.

```
# TERM=xterm
# export TERM
```

Example—Connecting to a Node’s Console Through the Terminal Concentrator

The following example shows how to connect to a cluster node in a configuration that uses a terminal concentrator. A Shell tool has already been started by using an xterm type window.

```
admin-ws# telnet tc1 5002
Trying 192.9.200.1 ...
Connected to 192.9.200.1.
Escape character is '^]'.
[Return]
pys-palindrome-1 console login: root
password: root_password
(for sh or ksh)
phys-palindrome-1# TERM=xterm; export TERM
(for csh)
phys-palindrome-1# set term=xterm
```

▼ How to Reset a Terminal Concentrator Port

When a port on the terminal concentrator is busy, that is, in use by a user, you can reset the port to disconnect that user. This procedure is useful if you need to perform an administrative task on the busy port.

A busy port returns the following message when you try to connect to the terminal concentrator.

```
telnet: Unable to connect to remote host: Connection refused
```

If you use the port selector, you might see a port busy message. See “How to Correct a Port Configuration Access Error” on page 29 for details on the port busy message.

1. Connect to the terminal concentrator port.

```
# telnet tc_name
```

tc_name Specifies the name of the terminal concentrator

2. Press an extra Return after making the connection and select the command-line interface to connect to the terminal concentrator.

```
...  
Enter Annex port name or number: cli  
annex:
```

3. Type the su command and password.

The default password is the terminal concentrator's IP address.

```
annex: su  
Password:
```

4. Determine which port to reset.

The who command shows ports that are in use.

```
annex# who
```

5. Reset the port that is in use.

```
annex# admin reset port_number
```

6. Disconnect from the terminal concentrator.

```
annex# hangup
```

You can now connect to the port.

Example—Resetting a Terminal Concentrator Connection

The following example shows how to reset the terminal concentrator connection on port 2.

```

admin-ws# telnet tc1
Trying 192.9.200.1 ...
Connected to 192.9.200.1.
Escape character is '^]'.
[Return]
...
Enter Annex port name or number: cli
...
annex: su
Password: root_password
annex: who
Port   What   User   Location   When   Idle   Address
2   PSVR   ---   ---   ---   1:27   192.9.75.12
v1   CLI   ---   ---   ---   192.9.76.10
annex# admin reset 2
annex# hangup

```


Installing and Maintaining Cluster Interconnect and Public Network Hardware

This chapter provides the procedures for adding, replacing, and removing cluster interconnect and public network hardware. Cluster interconnect components include transport adapters, transport junctions, and transport cables.

This chapter contains the following procedures:

- “How to Install Cluster Transport Adapters” on page 41
- “How to Replace Cluster Transport Adapters” on page 42
- “How to Remove Cluster Transport Adapters” on page 43
- “How to Install Cluster Transport Cables” on page 45
- “How to Add Public Network Connections” on page 46
- “How to Remove Public Network Connections” on page 47

For conceptual information on cluster interconnects and public network interfaces, see *Sun Cluster 3.0 Concepts*.

Installing and Maintaining Cluster Interconnect and Public Network Hardware

The following table provides procedures for installing and maintaining cluster interconnect and public network hardware.

TABLE 3-1 Task Map: Installing and Maintaining Cluster Interconnect and Public Network Hardware

Task	For Instructions, Go To...
Install cluster transport adapters	"How to Install Cluster Transport Adapters" on page 41
Replace cluster transport adapters	"How to Replace Cluster Transport Adapters" on page 42
Remove cluster transport adapters	"How to Remove Cluster Transport Adapters" on page 43
Install cluster transport junctions	"Installing Cluster Transport Junctions" on page 44
Install cluster transport cables	"How to Install Cluster Transport Cables" on page 45

Installing Cluster Transport Adapters

This section provides the procedure for adding cluster transport adapters. Use this procedure as part of the initial cluster hardware installation. Also use this procedure if you want to install cluster transport adapters after a Sun Cluster software installation.

For conceptual information on cluster transport adapters, see *Sun Cluster 3.0 Concepts*.

▼ How to Install Cluster Transport Adapters

1. Shut down the node in which you will install the cluster transport adapter.

```
# scswitch -S -h nodename  
# shutdown -y -g 0
```

For the procedure on shutting down a node, see *Sun Cluster 3.0 System Administration Guide*.

2. Power off the node.

For the procedure on powering off a node, see the documentation that shipped with your node.

3. Install the cluster transport adapter.

This procedure is the same procedure in a non-cluster environment. For the procedure on installing host adapters, see the documentation that shipped with your network adapter and node.

4. Power on and boot the node.

For the procedures on powering on and booting a node, see *Sun Cluster 3.0 System Administration Guide*.

Where to Go From Here

If you are performing an initial installation of cluster hardware, see “Installing Cluster Transport Junctions” on page 44. If you have Sun Cluster software installed and want the cluster transport adapter you added to appear in the Sun Cluster configuration, see *Sun Cluster 3.0 System Administration Guide*.

Replacing Cluster Transport Adapters

This section provides the procedure for replacing cluster transport adapters.

For conceptual information on cluster transport adapters, see *Sun Cluster 3.0 Concepts*.

▼ How to Replace Cluster Transport Adapters



Caution - You must maintain at least one cluster interconnect between the nodes of a cluster. The cluster will not function properly without a functioning cluster interconnect.

1. Shut down the node with the cluster transport adapter you want to replace.

```
# scswitch -S -h nodename  
# shutdown -y -g 0
```

For the procedure on shutting down a node, see *Sun Cluster 3.0 System Administration Guide*.

2. Power off the node.

For the procedure on powering off a node, see the documentation that shipped with your node.

3. Disconnect the cluster transport cable from the cluster transport adapter and other devices.

This procedure is the same in a non-cluster environment. For the procedure on disconnecting cables from cluster transport adapters, see the documentation that shipped with your network adapter and node.

4. Replace the cluster transport adapter.

This procedure is the same in a non-cluster environment. For the procedure on replacing cluster transport adapters, see the documentation that shipped with your network adapter and node.

5. Reconnect the cluster transport cable to the new cluster transport adapter.

This procedure is the same in a non-cluster environment. For the procedure on connecting cables to cluster transport adapters, see the documentation that shipped with your network adapter and node.

6. Power on and boot the node.

For the procedures on powering on and booting a node, see *Sun Cluster 3.0 System Administration Guide*.

Where to Go From Here

If you are performing an initial installation of cluster hardware, see “Installing Cluster Transport Junctions” on page 44. If you have Sun Cluster software installed, the cluster transport adapter you replaced appears in the Sun Cluster software configuration, and you want to reconfigure Sun Cluster with the new cluster transport adapter, see *Sun Cluster 3.0 System Administration Guide*.

Removing Cluster Transport Adapters

This section provides the procedure for removing cluster transport adapters.

For conceptual information on cluster transport adapters, see *Sun Cluster 3.0 Concepts*.

▼ How to Remove Cluster Transport Adapters



Caution - You must maintain at least one cluster interconnect between the nodes of a cluster. The cluster will not function properly without a functioning cluster interconnect.

1. Verify that the cluster transport adapter you want to replace appears in the Sun Cluster software configuration.

- If the cluster transport adapter you want to replace appears in the Sun Cluster software configuration, remove the cluster transport adapter from the Sun Cluster configuration. To remove a cluster interconnect, see the *Sun Cluster 3.0 System Administration Guide*, then proceed to Step 2 on page 43.
- If the cluster transport adapter you want to replace does not appear in the Sun Cluster software configuration, proceed to Step 2 on page 42.

2. Disconnect any cluster interconnect connections cabled to the cluster transport adapter.

This procedure is the same procedure in a non-cluster environment. For the procedure on disconnecting network connections from cluster transport adapters, see the documentation that shipped with your network adapter.

3. Shut down the node with the cluster transport adapter you want to remove.

```
# scswitch -S -h nodename  
# shutdown -y -g 0
```

For the procedure on shutting down a node, see *Sun Cluster 3.0 System Administration Guide*.

4. Power off the node.

For the procedure on powering off a node, see the documentation that shipped with your node.

5. Disconnect the cluster transport cables from the cluster transport adapter and other devices.

This procedure is the same in a non-cluster environment. For the procedure on disconnecting cables from cluster transport adapters, see the documentation that shipped with your network adapter and node.

6. Remove the cluster transport adapter.

This procedure is the same in a non-cluster environment. For the procedure on removing cluster transport adapters, see the documentation that shipped with your network adapter and node.

7. Power on and boot the node.

For the procedures on powering on and booting a node, see *Sun Cluster 3.0 System Administration Guide*.

Installing Cluster Transport Junctions

A cluster with two nodes can use a point-to-point connections, requiring no transport junctions. A cluster with more than two nodes requires two transport junctions, which are Ethernet-based switches (customer-supplied), to provide means of communication among all nodes in the cluster.

Installing Cluster Transport Cables

This section provides the procedure for cabling the cluster transport.

▼ How to Install Cluster Transport Cables

1. Locate the host adapters.

When cabling the cluster interconnect, be aware that the location of the host adapters might differ by platform. For example, some platforms have the host adapters mounted on the motherboards. In others, the cluster transport adapters are mounted on system and I/O boards, respectively.

For information on locating your cluster transport adapters, see the documentation that shipped with your network adapter and node.

2. Cable the nodes, as shown in Figure 3–1 or Figure 3–2.

You need to use a point-to-point (crossover) Ethernet cable if you are using 100BaseT or TPE ports on a node directly to ports on another node (no switch). Gigabit Ethernet uses the standard fiber cable for both point-to-point and switch configurations.




Figure 3–1 Typical Two-Node Sun Cluster Interconnect




Figure 3-2 Typical Four-Node Cluster Interconnect

Where to Go From Here

See “Adding Public Network Connections” on page 46.

Adding Public Network Connections

Adding public network connections to a node in a cluster is no different than doing so in a non-cluster environment.

For procedures related to administering public network connections, see *Sun Cluster 3.0 System Administration Guide*.

▼ How to Add Public Network Connections

For instructions on adding public network connections, see the hardware documentation that shipped with your node and public network connection devices.

Where to Go From Here

To install your storage, see “Installing a StorEdge MultiPack” on page 49, “Installing a StorEdge D1000” on page 73, “Installing a StorEdge A5x00” on page 99, or “Installing a StorEdge A3500” on page 117.

Removing Public Network Connections



Caution - You must maintain at least one public network connection for each node in the cluster. The cluster is inaccessible without a public network connection.

Removing public network connections from a node in a cluster is no different than doing so in a non-cluster environment. For procedures related to administering public network connections, see *Sun Cluster 3.0 System Administration Guide*.

▼ How to Remove Public Network Connections

For instructions on removing public network connections, see the hardware documentation that shipped with your node and public network connection devices.

Installing and Maintaining the Sun StorEdge MultiPack Enclosure

This chapter provides the procedures for installing and maintaining a Sun StorEdge™ MultiPack enclosure.

This chapter contains the following procedures:

- “How to Install a StorEdge MultiPack” on page 50
- “How to Add a StorEdge MultiPack Disk Drive” on page 55
- “How to Replace a StorEdge MultiPack Disk Drive” on page 57
- “How to Remove a StorEdge MultiPack Disk Drive” on page 61
- “How to Add a StorEdge MultiPack Enclosure” on page 62
- “How to Replace a StorEdge MultiPack Enclosure” on page 68
- “How to Remove a StorEdge MultiPack Enclosure” on page 69

For conceptual information on multihost disks, see *Sun Cluster 3.0 Concepts*.

Installing a StorEdge MultiPack

This section provides the procedure for an initial installation of a StorEdge MultiPack enclosure. The following table lists the steps involved in an initial installation of a StorEdge MultiPack enclosure.

TABLE 4-1 Task Map:Installing a StorEdge MultiPack

Task	For Instructions, Go To...
Install the host adapters	The documentation that shipped with your nodes
Cable the enclosure	“How to Install a StorEdge MultiPack” on page 50
Ensure that each device in the SCSI chain has a unique SCSI address	“How to Install a StorEdge MultiPack” on page 50
Instruct the OpenBoot PROM Monitor to use the new nvramrc script	“How to Install a StorEdge MultiPack” on page 50
Continue with the Solaris operating environment, Sun Cluster software, and volume management software installation tasks	<i>Sun Cluster 3.0 Installation Guide</i>

▼ How to Install a StorEdge MultiPack

Use this procedure for an initial installation of a StorEdge MultiPack enclosure, prior to installing the Solaris operating environment and Sun Cluster software. Perform this procedure in conjunction with the procedures in *Sun Cluster 3.0 Installation Guide* and your server hardware manual.

Multihost storage in clusters uses the multi-initiator capability of the SCSI (Small Computer System Interface) specification. For conceptual information on multi-initiator capability, see *Sun Cluster 3.0 Concepts*.

1. Ensure that each device in the SCSI chain has a unique SCSI address.

The default SCSI address for host adapters is 7. Reserve SCSI address 7 for one host adapter in the SCSI chain. This procedure refers to the host adapter you choose for SCSI address 7 as the host adapter on the second node. To avoid conflicts, in Step 7 on page 52 you will change the `scsi-initiator-id` of the remaining host adapter in the SCSI chain to an available SCSI address. This procedure refers to the host adapter with an available SCSI address as the host adapter on the first node. Depending on the device and configuration settings of the device, either SCSI address 6 or 8 is usually available.



Caution - Even though a slot in the enclosure might not be in use, you should avoid setting the `scsi-initiator-id` for the first node to the SCSI address for that disk slot. This precaution minimizes future complications if you install additional disk drives.

For more information, see the *OpenBoot 3.x Command Reference Manual* and the labels inside the storage device.

2. Install the host adapters in the nodes that will be connected to the enclosure.

For the procedure on installing host adapters, see the documentation that shipped with your host adapters and nodes.

3. Connect the cables to the enclosure, as shown in Figure 4-1.

Make sure that the *entire* SCSI bus length to each enclosure is less than 6 m. This measurement includes the cables to both nodes, as well as the bus length internal to each enclosure, node, and host adapter. Refer to the documentation that shipped with the enclosure for other restrictions regarding SCSI operation.




Figure 4-1 Example of a StorEdge MultiPack enclosure mirrored pair

4. **Connect the AC power cord for each enclosure of the pair to a different power source.**
5. **Without allowing the node to boot, power on the first node. If necessary, abort the system to continue with OpenBoot PROM Monitor tasks.**
6. **Find the paths to the host adapters.**

```
{0} ok show-disks
a) /pci@1f,4000/pci@4/SUNW,isptwo@4/sd
b) /pci@1f,4000/pci@2/SUNW,isptwo@4/sd
```

Identify and record the two controllers that will be connected to the storage devices, and record these paths. Use this information to change the SCSI addresses of these controllers in the `nvrarc` script. Do not include the `/sd` directories in the device paths.

7. **Edit the `nvrarc` script to set the `scsi-initiator-id` for the host adapters on the first node.**

For a list of `nvrarc` editor and `nvedit` keystroke commands, see Appendix B. The following example sets the `scsi-initiator-id` to 6. The OpenBoot PROM Monitor prints the line numbers (0:, 1:, and so on).



Caution - Insert exactly one space after the first double quote and before `scsi-initiator-id`.

```
{0} ok nvedit
0: probe-all
1: cd /pci@1f,4000/pci@4/SUNW,isptwo@4
2: 6 " scsi-initiator-id" integer-property
3: device-end
4: cd /pci@1f,4000/pci@2/SUNW,isptwo@4
5: 6 " scsi-initiator-id" integer-property
6: device-end
7: install-console
8: banner <Control C>
{0} ok
```

8. **Store the changes.**

The changes you make through the `nvedit` command are done on a temporary copy of the `nvrarc` script. You can continue to edit this copy without risk. After you complete your edits, save the changes. If you are not sure about the changes, discard them.

- To store the changes, type:

```
{0} ok nvstore
{0} ok
```

- To discard the changes, type:

```
{0} ok nvquit
{0} ok
```

9. Verify the contents of the `nvrarc` script you created in Step 7 on page 52.

If the contents of the `nvrarc` script are incorrect, use the `nvedit` command to make corrections.

```
{0} ok printenv nvrarc
nvrarc =
        probe-all
        cd /pci@1f,4000/pci@4/SUNW,ispw@4
        6 " scsi-initiator-id" integer-property
        device-end
        cd /pci@1f,4000/pci@2/SUNW,ispw@4
        6 " scsi-initiator-id" integer-property
        device-end
        install-console
        banner
{0} ok
```

10. Instruct the OpenBoot PROM Monitor to use the `nvrarc` script.

```
{0} ok setenv use-nvramrc? true
use-nvramrc? = true
{0} ok
```

11. Without allowing the node to boot, power on the second node. If necessary, abort the system to continue with OpenBoot PROM Monitor tasks.

12. Verify that the `scsi-initiator-id` for the host adapter on the second node is set to 7.

Use the `show-disks` command to find the paths to the host adapters connected to these enclosures. Select each host adapter's device tree node, and display the node's properties to confirm that the `scsi-initiator-id` for each host adapter is set to 7.

```
{0} ok cd /pci@1f,4000/pci@4/SUNW,isptwo@4
{0} ok .properties
...
scsi-initiator-id      00000007
...
{0} ok cd /pci@1f,4000/pci@2/SUNW,isptwo@4
{0} ok .properties
...
scsi-initiator-id      00000007
```

13. Continue with the Solaris operating environment, Sun Cluster software, and volume management software installation tasks.

For software installation procedures, see *Sun Cluster 3.0 Installation Guide*.

Maintaining a StorEdge MultiPack

This section provides the procedures for maintaining a StorEdge MultiPack enclosure. The following table lists these procedures.

TABLE 4-2 Task Map:Maintaining a StorEdge MultiPack Enclosure

Task	For Instructions, Go To...
Add a disk drive Use the <code>cfgadm(1M)</code> , <code>devfsadm(1M)</code> , <code>scgdevs(1M)</code> , and <code>scdidadm(1M)</code> commands to add a disk drive to an enclosure.	“How to Add a StorEdge MultiPack Disk Drive” on page 55
Replace a disk drive	“How to Replace a StorEdge MultiPack Disk Drive” on page 57
Remove a disk drive	“How to Remove a StorEdge MultiPack Disk Drive” on page 61
Add an enclosure to an existing cluster Use the <code>nvedit(1M)</code> , <code>shutdown(1M)</code> , <code>scswitch(1M)</code> , and <code>scdidadm</code> commands to add an enclosure to an existing cluster.	“How to Add a StorEdge MultiPack Enclosure” on page 62
Replace an enclosure in an existing cluster Use this procedure to replace an enclosure with a new enclosure. This procedure assumes that you want to retain the disk drives in the enclosure you want to replace.	“How to Replace a StorEdge MultiPack Enclosure” on page 68
Remove an enclosure from an existing cluster Use the <code>cfgadm</code> , <code>devfsadm</code> , and <code>scdidadm</code> commands to remove an enclosure without replacing it with another enclosure.	“How to Remove a StorEdge MultiPack Enclosure” on page 69

▼ How to Add a StorEdge MultiPack Disk Drive

Use this procedure to add a disk drive to an existing cluster. Perform the steps in this procedure in conjunction with the procedures in *Sun Cluster 3.0 System Administration Guide* and your server hardware manual.

For conceptual information on quorum, quorum devices, global devices, and device IDs, see *Sun Cluster 3.0 Concepts*.

1. Locate an empty disk slot in the enclosure for the disk drive you want to add.

Identify the empty slots either by observing the disk drive LEDs on the front of the enclosure, or by removing the left side cover of the unit. The target address IDs corresponding to the slots appear on the middle partition of the drive bay.

2. Install the disk drive.

For more information, see the documentation that shipped with your enclosure.

3. On all nodes attached to the enclosure, configure the disk drive.

```
# cfmadm -c configure cN
# devfsadm
```

4. On all nodes, ensure that entries for the disk drive have been added to the /dev/rdisk directory.

```
# ls -l /dev/rdisk
```

5. If needed, use the format(1M) command or the fmthard(1M) command to partition the disk drive.

6. From any node, update the global device namespace.

If a volume management daemon such as vold is running on your node, and you have a CD-ROM drive connected to the node, a device busy error might be returned even if no disk is in the drive. This error is an expected behavior.

```
# scgdevs
```

7. On all nodes, verify that a device ID (DID) has been assigned to the disk drive.

```
# scdidadm -l
```

As shown in “Example—Adding a StorEdge MultiPack Disk Drive” on page 57, the DID 35 assigned to the new disk drive might not be in sequential order in the disk array.

8. Perform volume management administration to add the new disk drive to the configuration.

For more information, see your Solstice DiskSuite™ or VERITAS Volume Manager documentation.

9. If you want this new disk drive to be a quorum device, add the quorum device.

Refer to *Sun Cluster 3.0 System Administration Guide* for the procedure on adding a quorum device.

Example—Adding a StorEdge MultiPack Disk Drive

The following example shows how to apply the procedure for adding a StorEdge MultiPack enclosure disk drive.

```
# scdidadm -l
16 phys-circinus-3:/dev/rdisk/c2t0d0 /dev/did/rdsk/d16
17 phys-circinus-3:/dev/rdisk/c2t1d0 /dev/did/rdsk/d17
18 phys-circinus-3:/dev/rdisk/c2t2d0 /dev/did/rdsk/d18
19 phys-circinus-3:/dev/rdisk/c2t3d0 /dev/did/rdsk/d19
...
26 phys-circinus-3:/dev/rdisk/c2t12d0 /dev/did/rdsk/d26
30 phys-circinus-3:/dev/rdisk/c1t2d0 /dev/did/rdsk/d30
31 phys-circinus-3:/dev/rdisk/c1t3d0 /dev/did/rdsk/d31
32 phys-circinus-3:/dev/rdisk/c1t10d0 /dev/did/rdsk/d32
33 phys-circinus-3:/dev/rdisk/c0t0d0 /dev/did/rdsk/d33
34 phys-circinus-3:/dev/rdisk/c0t6d0 /dev/did/rdsk/d34
8190 phys-circinus-3:/dev/rmt/0 /dev/did/rmt/2
# cfgadm -c configure c1
# devfsadm
# scgdevs
Configuring DID devices
Could not open /dev/rdisk/c0t6d0s2 to verify device id.
Device busy
Configuring the /dev/global directory (global devices)
obtaining access to all attached disks
reservation program successfully exiting
# scdidadm -l
16 phys-circinus-3:/dev/rdisk/c2t0d0 /dev/did/rdsk/d16
17 phys-circinus-3:/dev/rdisk/c2t1d0 /dev/did/rdsk/d17
18 phys-circinus-3:/dev/rdisk/c2t2d0 /dev/did/rdsk/d18
19 phys-circinus-3:/dev/rdisk/c2t3d0 /dev/did/rdsk/d19
...
26 phys-circinus-3:/dev/rdisk/c2t12d0 /dev/did/rdsk/d26
30 phys-circinus-3:/dev/rdisk/c1t2d0 /dev/did/rdsk/d30
31 phys-circinus-3:/dev/rdisk/c1t3d0 /dev/did/rdsk/d31
32 phys-circinus-3:/dev/rdisk/c1t10d0 /dev/did/rdsk/d32
33 phys-circinus-3:/dev/rdisk/c0t0d0 /dev/did/rdsk/d33
34 phys-circinus-3:/dev/rdisk/c0t6d0 /dev/did/rdsk/d34
35 phys-circinus-3:/dev/rdisk/c2t13d0 /dev/did/rdsk/d35
8190 phys-circinus-3:/dev/rmt/0 /dev/did/rmt/2
#
```

▼ How to Replace a StorEdge MultiPack Disk Drive

Use this procedure to replace a StorEdge MultiPack enclosure disk drive.

“Example—Replacing a StorEdge MultiPack Disk Drive” on page 60 shows you how to apply this procedure. Perform the steps in this procedure in conjunction with the

procedures in *Sun Cluster 3.0 System Administration Guide* and your server hardware manual. Use the procedures in your server hardware manual to identify a failed disk drive.

For conceptual information on quorum, quorum devices, global devices, and device IDs, see *Sun Cluster 3.0 Concepts*.

1. Identify the disk drive that needs replacement.

If the disk error message reports the drive problem by device ID (DID), use the `scdidadm -l` command to determine the Solaris logical device name. If the disk error message reports the drive problem by the Solaris physical device name, use your Solaris documentation to map the Solaris physical device name to the Solaris logical device name. Use this Solaris logical device name and DID throughout this procedure.

```
# scdidadm -l deviceID
```

2. Determine if the disk drive you want to replace is a quorum device.

```
# scstat -q
```

If the disk drive you want to replace is a quorum device, remove the quorum device before you proceed. Otherwise, proceed to Step 3 on page 58.

Refer to *Sun Cluster 3.0 System Administration Guide* for procedures on replacing a quorum device and putting a quorum device into maintenance state.

3. If possible, back up the metadvice or volume.

For more information, see your Solstice DiskSuite or VERITAS Volume Manager documentation.

4. Perform volume management administration to remove the disk drive from the configuration.

For more information, see your Solstice DiskSuite or VERITAS Volume Manager documentation.

5. Identify the failed disk drive's physical DID.

Use this physical DID in Step 12 on page 60 to verify that the failed disk drive has been replaced with a new disk drive.

```
# scdidadm -o diskid -l cNtXdY
```

6. If you are using Solstice DiskSuite as your volume manager, save the disk partitioning for use when partitioning the new disk drive.

If you are using VERITAS Volume Manager, proceed to Step 7 on page 59.

```
# prtvtoc /dev/rdisk/cNtXdYsZ > filename
```



Caution - Do not save this file under /tmp because you will lose this file when you reboot. Instead, save this file under /usr/tmp.

7. Replace the failed disk drive.

For more information, see the documentation that shipped with your enclosure.

8. On one node attached to the enclosure, run the devfsadm(1M) command to probe all devices and to write the new disk drive to the /dev/rdisk directory.

Depending on the number of devices connected to the node, the devfsadm command can take at least five minutes to complete.

```
# devfsadm
```

9. If you are using Solstice DiskSuite as your volume manager, from any node connected to the enclosure, partition the new disk drive, using the partitioning you saved in Step 6 on page 58.

If you are using VERITAS Volume Manager, proceed to Step 10 on page 59.

```
# fmthard -s filename /dev/rdisk/cNtXdYsZ
```

10. One at a time, shut down and reboot the nodes connected to the enclosure.

```
# scswitch -S -h nodename  
# shutdown -y -g 0 -i 6
```

For more information, see *Sun Cluster 3.0 System Administration Guide*.

11. From any node connected to the disk drive, update the DID database.

```
# scdidadm -R deviceID
```

12. From any node, confirm that the failed disk drive has been replaced by comparing the new physical DID to the physical DID in Step 5.

If the new physical DID is different from the physical DID in Step 5 on page 58, you successfully replaced the failed disk drive with a new disk drive.

```
# scdidadm -o diskid -l cNtXdY
```

13. On all connected nodes, upload the new information to the DID driver.

If a volume management daemon such as vold is running on your node, and you have a CD-ROM drive connected to the node, a device busy error might be returned even if no disk is in the drive. This error is an expected behavior.

```
# scdidadm -ui
```

14. Perform volume management administration to add the disk drive back to its diskset or disk group.

For more information, see your Solstice DiskSuite or VERITAS Volume Manager documentation.

15. If you want this new disk drive to be a quorum device, add the quorum device.

Refer to *Sun Cluster 3.0 System Administration Guide* for the procedure on adding a quorum device.

Example—Replacing a StorEdge MultiPack Disk Drive

The following example shows how to apply the procedure for replacing a StorEdge MultiPack enclosure disk drive.

```
# scdidadm -l d20
20      phys-schost-2:/dev/rdisk/c3t2d0 /dev/did/rdsk/d20
# scdidadm -o diskid -l c3t2d0
5345414741544520393735314336343734310000
# prtvtoc /dev/rdsk/c3t2d0s2 > /usr/tmp/c3t2d0.vtoc
...
# devfsadm
# fmthard -s /usr/tmp/c3t2d0.vtoc /dev/rdsk/c3t2d0s2
# scswitch -S -h node1
# shutdown -y -g 0 -i 6
...
# scdidadm -R d20
# scdidadm -o diskid -l c3t2d0
```

(continued)

```
5345414741544520393735314336363037370000
# scdidadm -ui
```

▼ How to Remove a StorEdge MultiPack Disk Drive

Use this procedure to remove a disk drive from a StorEdge MultiPack enclosure. Perform the steps in this procedure in conjunction with the procedures in *Sun Cluster 3.0 System Administration Guide* and your server hardware manual.

For conceptual information on quorum, quorum devices, global devices, and device IDs, see *Sun Cluster 3.0 Concepts*.

1. Determine if the disk drive you want to remove is a quorum device.

```
# scstat -q
```

If the disk drive you want to remove is a quorum device, remove the quorum device before you proceed. Otherwise, proceed to Step 2 on page 61.

Refer to *Sun Cluster 3.0 System Administration Guide* for procedures on removing a quorum device and putting a quorum device into maintenance state.

2. Perform volume management administration to remove the disk drive from the configuration.

For more information, see your Solstice DiskSuite or VERITAS Volume Manager documentation.

3. Identify the disk drive that needs to be removed and the slot that the disk drive needs to be removed from.

If the disk error message reports the drive problem by DID, use the `scdidadm -l` command to determine the Solaris device name.

```
# scdidadm -l deviceID
# cfgadm -al
```

4. Remove the disk drive.

For more information, see the documentation that shipped with your enclosure.

5. On all nodes, remove references to the disk drive.

```
# cfigadm -c unconfigure cN::dsk/cNtXdY
# devfsadm -C
# scdidadm -C
```

▼ How to Add a StorEdge MultiPack Enclosure

Use this procedure to install a StorEdge MultiPack enclosure in an existing cluster. Perform the steps in this procedure in conjunction with the procedures in *Sun Cluster 3.0 Installation Guide* and your server hardware manual.

For conceptual information on multi-initiator SCSI and device IDs, see *Sun Cluster 3.0 Concepts*.

1. Ensure that each device in the SCSI chain has a unique SCSI address.

The default SCSI address for host adapters is 7. Reserve SCSI address 7 for one host adapter in the SCSI chain. This procedure refers to the host adapter you choose for SCSI address 7 as the host adapter on the **second** node. To avoid conflicts, in Step 7 on page 52 you will change the `scsi-initiator-id` of the remaining host adapter in the SCSI chain to an available SCSI address. This procedure refers to the host adapter with an available SCSI address as the host adapter on the **first** node. Depending on the device and configuration settings of the device, either SCSI address 6 or 8 is usually available.



Caution - Even though a slot in the enclosure might not be in use, you should avoid setting the `scsi-initiator-id` for the first node to the SCSI address for that disk slot. This precaution minimizes future complications if you install additional disk drives.

For more information, see the *OpenBoot 3.x Command Reference Manual* and the labels inside the storage device.

2. Shut down and power off the first node.

```
# scswitch -S -h nodename  
# shutdown -y -g 0
```

For more information, see *Sun Cluster 3.0 System Administration Guide*.

3. Install the host adapters in the node.

For the procedure on installing host adapters, see the documentation that shipped with your host adapters and nodes.

4. Connect the single-ended SCSI cable between the node and the enclosures, as shown in Figure 4-2.

Make sure that the *entire* SCSI bus length to each enclosure is less than 6 m. This measurement includes the cables to both nodes, as well as the bus length internal to each enclosure, node, and host adapter. Refer to the documentation that shipped with the enclosure for other restrictions regarding SCSI operation.




Figure 4-2 Example of a StorEdge MultiPack enclosure mirrored pair

5. Temporarily install a single-ended terminator on the SCSI IN port of the second enclosure, as shown in Figure 4-2.

6. Connect the enclosures to different power sources.

7. Power on the node and the enclosures.

8. Find the paths to the host adapters.

```
{0} ok show-disks
a) /pci@1f,4000/pci@4/SUNW,isptwo@4/sd
b) /pci@1f,4000/pci@2/SUNW,isptwo@4/sd
```

Identify and record the two controllers that will be connected to the storage devices, and record these paths. Use this information to change the SCSI addresses of these controllers in the `nvrarc` script. Do not include the `/sd` directories in the device paths.

9. Edit the `nvrarc` script to set the `scsi-initiator-id` for the host adapters on the first node.

For a list of `nvrarc` editor and `nvedit` keystroke commands, see Appendix B. The following example sets the `scsi-initiator-id` to 6. The OpenBoot PROM Monitor prints the line numbers (0:, 1:, and so on).



Caution - Insert exactly one space after the first double quote and before `scsi-initiator-id`.

```
{0} ok nvedit
0: probe-all
1: cd /pci@1f,4000/pci@4/SUNW,isptwo@4
2: 6 " scsi-initiator-id" integer-property
3: device-end
4: cd /pci@1f,4000/pci@2/SUNW,isptwo@4
5: 6 " scsi-initiator-id" integer-property
6: device-end
7: install-console
8: banner <Control C>
{0} ok
```

10. Store the changes.

The changes you make through the `nvedit` command are done on a temporary copy of the `nvrarc` script. You can continue to edit this copy without risk. After you complete your edits, save the changes. If you are not sure about the changes, discard them.

- To store the changes, type:

```
{0} ok nvstore  
{0} ok
```

- To discard the changes, type:

```
{0} ok nvquit  
{0} ok
```

11. Verify the contents of the nvramrc script you created in Step 9 on page 64.

If the contents of the nvramrc script are incorrect, use the nvedit command to make corrections.

```
{0} ok printenv nvramrc  
nvramrc =  
    probe-all  
    cd /pci@1f,4000/pci@4/SUNW,isp2wo@4  
    6 "scsi-initiator-id" integer-property  
    device-end  
    cd /pci@1f,4000/pci@2/SUNW,isp2wo@4  
    6 "scsi-initiator-id" integer-property  
    device-end  
    install-console  
    banner  
  
{0} ok
```

12. Instruct the OpenBoot PROM Monitor to use the nvramrc script.

```
{0} ok setenv use-nvramrc? true  
use-nvramrc? = true  
{0} ok
```

13. Boot the first node and wait for it to join the cluster.

```
{0} ok boot -r
```

For more information, see *Sun Cluster 3.0 System Administration Guide*.

14. On all nodes, verify that the DIDs have been assigned to the disk drives in the enclosure.

```
# scdidadm -l
```

15. Shut down the second node.

```
# scswitch -S -h nodename  
# shutdown -y -g 0
```

16. Power off the second node.

17. Install the host adapters in the second node.

For the procedure on installing host adapters, see the documentation that shipped with your nodes.

18. Connect the enclosures to the host adapters by using single-ended SCSI cables.




Figure 4-3 Example of a StorEdge MultiPack enclosure mirrored pair

19. Without allowing the node to boot, power on the second node. If necessary, abort the system to continue with OpenBoot PROM Monitor tasks.

20. Verify that the second node sees the new host adapters and disk drives.

```
{0} ok show-disks
```

21. Verify that the scsi-initiator-id for the host adapter on the second node is set to 7.

Use the show-disks command to find the paths to the host adapters connected to these enclosures. Select each host adapter's device tree node, and display the node's properties to confirm that the scsi-initiator-id for each host adapter is set to 7.

```
{0} ok cd /pci@1f,4000/pci@4/SUNW,ispw@4
{0} ok .properties
...
scsi-initiator-id      00000007
...
{0} ok cd /pci@1f,4000/pci@2/SUNW,ispw@4
{0} ok .properties
...
scsi-initiator-id      00000007
```

22. Boot the second node and wait for it to join the cluster.

```
{0} ok boot -r
```

23. On all nodes, verify that the DIDs have been assigned to the disk drives in the enclosure.

```
# scdidadm -l
```

24. Perform volume management administration to add the disk drives in the array to the volume management configuration.

For more information, see your Solstice DiskSuite or VERITAS Volume Manager documentation.

▼ How to Replace a StorEdge MultiPack Enclosure

Use this procedure to replace a StorEdge MultiPack enclosure. This procedure assumes that you want to retain the disk drives in the enclosure you want to replace and retain the references to these same disk drives.

If you want to replace your disk drives, see “How to Replace a StorEdge MultiPack Disk Drive” on page 57.

1. If possible, back up the metadevices or volumes that reside in the disk array.

For more information, see your Solstice DiskSuite or VERITAS Volume Manager documentation.

2. Perform volume management administration to remove the disk array from the configuration.

For more information, see your Solstice DiskSuite or VERITAS Volume Manager documentation.

3. Disconnect the SCSI cables from the enclosure, removing them in the order shown in Figure 4-4.




Figure 4-4 Disconnecting the SCSI cables

4. Power off, and disconnect the enclosure from the AC power source.

For more information, see the documentation that shipped with your enclosure and the labels inside the lid of the enclosure.

5. Connect the new enclosure to an AC power source.

Refer to the documentation that shipped with the enclosure as well as the labels inside the lid of the enclosure.

6. Connect the SCSI cables to the new enclosure by reversing the steps shown in Figure 4-4.

7. Move the disk drives one at a time from the old enclosure to the same slot in the new enclosure.

8. Power on the enclosure.

9. On all nodes attached to the enclosure, run the `devfsadm(1M)` command.

```
# devfsadm
```

10. One at a time, shut down and reboot the nodes connected to the enclosure.

```
# scswitch -S -h nodename  
# shutdown -y -g 0 -i 6
```

For more information on `shutdown(1M)`, see *Sun Cluster 3.0 System Administration Guide*.

11. Perform volume management administration to add the disk array to the configuration.

For more information, see your Solstice DiskSuite or VERITAS Volume Manager documentation.

▼ How to Remove a StorEdge MultiPack Enclosure

Use this procedure to remove a StorEdge MultiPack enclosure from a cluster. This procedure assumes that you want to remove the references to the disk drives in the enclosure.

1. Perform volume management administration to remove the enclosure from the configuration.

For more information, see your Solstice DiskSuite or VERITAS Volume Manager documentation.

2. Disconnect the SCSI cables from the enclosure, removing them in the order shown in Figure 4-5.




Figure 4-5 Disconnecting the SCSI cables

3. Power off, and disconnect the enclosure from the AC power source.

For more information, see the documentation that shipped with the enclosure and the labels inside the lid of the enclosure.

4. Remove the enclosure.

5. Identify the disk drives you need to remove.

```
# cfgadm -al
```

6. On all nodes, remove references to the disk drives in the enclosure you removed in Step 4 on page 70.

```
# cfgadm -c unconfigure cN::dsk/cNtXdY
# devfsadm -C
# scdidadm -C
```

7. If needed, remove any host adapters from the nodes.

For the procedure on removing host adapters, see the documentation that shipped with your nodes.

Installing and Maintaining a Sun StorEdge D1000 Disk Array

This chapter provides the procedures for installing and maintaining a Sun StorEdge™ D1000 disk array.

This chapter contains the following procedures:

- “How to Install a StorEdge D1000” on page 74
- “How to Add a StorEdge D1000 Disk Drive” on page 80
- “How to Replace a StorEdge D1000 Disk Drive” on page 82
- “How to Remove a StorEdge D1000 Disk Drive” on page 86
- “How to Add a StorEdge D1000 Disk Array” on page 87
- “How to Replace a StorEdge D1000 Disk Array” on page 93
- “How to Remove a StorEdge D1000 Disk Array” on page 95

For conceptual information on multihost disks, see *Sun Cluster 3.0 Concepts*.

Installing a StorEdge D1000

This section provides the procedure for an initial installation of a StorEdge D1000 disk array. The following table lists the steps involved in an initial installation of a StorEdge D1000 disk array.

TABLE 5-1 Task Map: Installing a StorEdge D1000

Task	For Instructions, Go To...
Install the host adapters	The documentation that shipped with your nodes
Cable the disk array	"How to Install a StorEdge D1000" on page 74
Ensure that each device in the SCSI chain has a unique SCSI address	"How to Install a StorEdge D1000" on page 74
Instruct the OpenBoot PROM Monitor to use the new nvramrc script	"How to Install a StorEdge D1000" on page 74
Continue with the Solaris operating environment, Sun Cluster software, and volume management software installation tasks	<i>Sun Cluster 3.0 Installation Guide</i>

▼ How to Install a StorEdge D1000

Use this procedure for an initial installation of a StorEdge D1000 disk array, prior to installing the Solaris operating environment and Sun Cluster software. Perform this procedure in conjunction with the procedures in *Sun Cluster 3.0 Installation Guide* and your server hardware manual.

Multihost storage in clusters uses the multi-initiator capability of the SCSI specification. For conceptual information on multi-initiator capability, see *Sun Cluster 3.0 Concepts*.

1. Ensure that each device in the SCSI chain has a unique SCSI address.

The default SCSI address for host adapters is 7. Reserve SCSI address 7 for one host adapter in the SCSI chain. This procedure refers to the host adapter you choose for SCSI address 7 as the host adapter on the second node. To avoid conflicts, in Step 7 on page 77 you will change the `scsi-initiator-id` of the remaining host adapter in the SCSI chain to an available SCSI address. This procedure refers to the host adapter with an available SCSI address as the host adapter on the first node. SCSI address 6 is usually available.

For more information, see the *OpenBoot 3.x Command Reference Manual* and the labels inside the storage device.



Caution - Even though a slot in the disk array might not be in use, you should avoid setting the `scsi-initiator-id` for the first node to the SCSI address for that disk slot. This precaution minimizes future complications if you install additional disk drives.

2. Install the host adapters in the node that will be connected to the disk array.

For the procedure on installing host adapters, see the documentation that shipped with your host adapters and nodes.

3. Connect the cables to the disk arrays, as shown in Figure 5–1.

Make sure that the *entire* bus length connected to each disk array is less than 25 m. This measurement includes the cables to both nodes, as well as the bus length internal to each disk array, node, and the host adapter.




Figure 5-1 Example of a StorEdge D1000 disk array mirrored pair

4. **Connect the AC power cord for each disk array of the pair to a different power source.**
5. **Power on the first node and the disk arrays.**
6. **Find the paths to the host adapters.**

```
{0} ok show-disks
```

Identify and record the two controllers that will be connected to the storage devices and record these paths. You will use this information to change the SCSI addresses of these controllers in the `nvrampc` script. Do not include the `/sd` directories in the device paths.

7. Edit the nvramrc script to change the scsi-initiator-id for the host adapters on the first node.

For a list of nvramrc editor and nvedit keystroke commands, see Appendix B.

The following example sets the scsi-initiator-id to 6. The OpenBoot PROM Monitor prints the line numbers (0:, 1:, and so on).



Caution - Insert exactly one space after the first double quote and before scsi-initiator-id.

```
{0} ok nvedit
0: probe-all
1: cd /sbus@1f,0/QLGC,isp@3,10000
2: 6 encode-int " scsi-initiator-id" property
3: device-end
4: cd /sbus@1f,0/
5: 6 encode-int " scsi-initiator-id" property
6: device-end
7: install-console
8: banner [Control C]
{0} ok
```

8. Store the changes.

The changes you make through the nvedit command are done on a temporary copy of the nvramrc script. You can continue to edit this copy without risk. After you complete your edits, save the changes. If you are not sure about the changes, discard them.

- To store the changes, type:

```
{0} ok nvstore
{0} ok
```

- To discard the changes, type:

```
{0} ok nvquit
{0} ok
```

9. Verify the contents of the nvramrc script you created in Step 7 on page 77.

If the contents of the nvramrc script are incorrect, use the nvedit command to make corrections.

```
{0} ok printenv nvramrc
nvramrc =                probe-all
                           cd /sbus@1f,0/QLGC,isp@3,10000
                           6 encode-int " scsi-initiator-id" property
                           device-end
                           cd /sbus@1f,0/
                           6 encode-int " scsi-initiator-id" property
                           device-end
                           install-console
                           banner
{0} ok
```

10. Instruct the OpenBoot PROM Monitor to use the nvramrc script.

```
{0} ok setenv use-nvramrc? true
use-nvramrc? = true
{0} ok
```

11. Without allowing the node to boot, power on the second node. If necessary, abort the system to continue with OpenBoot PROM Monitor tasks.**12. Verify that the scsi-initiator-id for each host adapter on the second node is set to 7.**

Use the show-disks command to find the paths to the host adapters connected to these enclosures. Select each host adapter's device tree node, and display the node's properties to confirm that the scsi-initiator-id for each host adapter is set to 7.

```

{0} ok cd /sbus@1f,0/QLGC,isp@3,10000
{0} ok .properties
scsi-initiator-id      00000007
differential
isp-fcode              1.21 95/05/18
device_type           scsi
...

```

13. Continue with the Solaris operating environment, Sun Cluster software, and volume management software installation tasks.

For software installation procedures, see *Sun Cluster 3.0 Installation Guide*.

Maintaining a StorEdge D1000 Disk Array

This section provides the procedures for maintaining a StorEdge D1000 disk array. The following table list these procedures.

TABLE 5-2 Task Map: Maintaining a StorEdge D1000 Disk Array

Task	For Instructions, Go To...
Add a disk drive Use the <code>cfgadm(1M)</code> , <code>devfsadm(1M)</code> , <code>scgdevs(1M)</code> , and <code>scdidadm(1M)</code> commands to add a disk drive to a disk array.	“How to Add a StorEdge D1000 Disk Drive” on page 80
Replace a disk drive	“How to Replace a StorEdge D1000 Disk Drive” on page 82
Remove a disk drive	“How to Remove a StorEdge D1000 Disk Drive” on page 86

TABLE 5-2 Task Map: Maintaining a StorEdge D1000 Disk Array *(continued)*

Task	For Instructions, Go To...
Add a disk array to an existing cluster Use the <code>nvedit(1M)</code> , <code>shutdown(1M)</code> , <code>scswitch(1M)</code> , and <code>scdidadm</code> commands to add a disk array to an existing cluster.	“How to Add a StorEdge D1000 Disk Array” on page 87
Replace a disk array in an existing cluster Use this procedure to replace a disk array with a new disk array. This procedure assumes that you want to retain the disk drives in the disk array you want to replace.	“How to Replace a StorEdge D1000 Disk Array” on page 93
Remove a disk array from an existing cluster Use the <code>cfgadm</code> , <code>devfsadm</code> , and <code>scdidadm</code> commands to remove a disk array without replacing it with another disk array.	“How to Remove a StorEdge D1000 Disk Array” on page 95

▼ How to Add a StorEdge D1000 Disk Drive

Use this procedure to add a disk drive to an existing cluster. Perform the steps in this procedure in conjunction with the procedures in *Sun Cluster 3.0 System Administration Guide* and your server hardware manual.

For conceptual information on quorum, quorum devices, global devices, and device IDs, see *Sun Cluster 3.0 Concepts*.

1. Locate an empty disk slot in the disk array for the disk drive you want to add.

Identify the disk slot in the disk array for the disk drive being added, and note the target number. Refer to the documentation that shipped with your StorEdge D1000 disk array.

2. Install the disk drive.

For more information, see the documentation that shipped with the disk array.

3. On all nodes attached to the disk array, configure the disk drive.

```
# cfgadm -c configure cN
# devfsadm
```

4. On all nodes, ensure that entries for the disk drive have been added to the /dev/rdisk directory.

```
# ls -l /dev/rdisk
```

5. If needed, use the `format(1M)` command or the `fmthard(1M)` command to partition the disk drive.

6. From any node, update the global device namespace.

If a volume management daemon such as `vold` is running on your node, and you have a CD-ROM drive connected to the node, a device busy error might be returned even if no disk is in the drive. This error is an expected behavior.

```
# scgdevs
```

7. On all nodes, verify that a device ID (DID) has been assigned to the disk drive.

```
# sddidadm -l
```

Note - As shown in “Example—Adding a StorEdge D1000 Disk Drive” on page 81, the DID 35 assigned to the new disk drive might not be in sequential order in the disk array.

8. Perform volume management administration to add the new disk drive to the configuration.

For more information, see your Solstice DiskSuite or VERITAS Volume Manager documentation.

9. If you want this new disk drive to be a quorum device, add the quorum device.

Refer to *Sun Cluster 3.0 System Administration Guide* for the procedure on adding a quorum device.

Example—Adding a StorEdge D1000 Disk Drive

The following example shows how to apply the procedure for adding a StorEdge D1000 disk array disk drive.

```
# sddidadm -l
16      phys-circinus-3:/dev/rdisk/c2t0d0 /dev/did/rdisk/d16
```

(continued)

```

17      phys-circinus-3:/dev/rdisk/c2t1d0 /dev/did/rdsk/d17
18      phys-circinus-3:/dev/rdisk/c2t2d0 /dev/did/rdsk/d18
19      phys-circinus-3:/dev/rdisk/c2t3d0 /dev/did/rdsk/d19
...
26      phys-circinus-3:/dev/rdisk/c2t12d0 /dev/did/rdsk/d26
30      phys-circinus-3:/dev/rdisk/c1t2d0 /dev/did/rdsk/d30
31      phys-circinus-3:/dev/rdisk/c1t3d0 /dev/did/rdsk/d31
32      phys-circinus-3:/dev/rdisk/c1t10d0 /dev/did/rdsk/d32
33      phys-circinus-3:/dev/rdisk/c0t0d0 /dev/did/rdsk/d33
34      phys-circinus-3:/dev/rdisk/c0t6d0 /dev/did/rdsk/d34
8190    phys-circinus-3:/dev/rmt/0      /dev/did/rmt/2
# cftadm -c configure c1
# devfsadm
# scgdevs
Configuring DID devices
Could not open /dev/rdsk/c0t6d0s2 to verify device id.
Device busy
Configuring the /dev/global directory (global devices)
obtaining access to all attached disks
reservation program successfully exiting
# scdidadm -l
16      phys-circinus-3:/dev/rdisk/c2t0d0 /dev/did/rdsk/d16
17      phys-circinus-3:/dev/rdisk/c2t1d0 /dev/did/rdsk/d17
18      phys-circinus-3:/dev/rdisk/c2t2d0 /dev/did/rdsk/d18
19      phys-circinus-3:/dev/rdisk/c2t3d0 /dev/did/rdsk/d19
...
26      phys-circinus-3:/dev/rdisk/c2t12d0 /dev/did/rdsk/d26
30      phys-circinus-3:/dev/rdisk/c1t2d0 /dev/did/rdsk/d30
31      phys-circinus-3:/dev/rdisk/c1t3d0 /dev/did/rdsk/d31
32      phys-circinus-3:/dev/rdisk/c1t10d0 /dev/did/rdsk/d32
33      phys-circinus-3:/dev/rdisk/c0t0d0 /dev/did/rdsk/d33
34      phys-circinus-3:/dev/rdisk/c0t6d0 /dev/did/rdsk/d34
35      phys-circinus-3:/dev/rdsk/c2t13d0 /dev/did/rdsk/d35
8190    phys-circinus-3:/dev/rmt/0      /dev/did/rmt/2
#

```

▼ How to Replace a StorEdge D1000 Disk Drive

Use this procedure to replace a StorEdge D1000 disk array disk drive.

“Example—Replacing a StorEdge D1000 Disk Drive” on page 85 shows you how to apply this procedure. Perform the steps in this procedure in conjunction with the procedures in *Sun Cluster 3.0 System Administration Guide* and your server hardware manual. Use the procedures in your server hardware manual to identify a failed disk drive.

For conceptual information on quorum, quorum devices, global devices, and device IDs, see *Sun Cluster 3.0 Concepts*.

1. Identify the disk drive that needs replacement.

If the disk error message reports the drive problem by device ID (DID), use the `scdidadm -l` command to determine the Solaris logical device name. If the disk error message reports the drive problem by the Solaris physical device name, use your Solaris documentation to map the Solaris physical device name to the Solaris logical device name. Use this Solaris logical device name and DID throughout this procedure.

```
# sddidadm -l deviceID
```

2. Determine if the disk drive you want to replace is a quorum device.

```
# scstat -q
```

If the disk drive you want to replace is a quorum device, remove the quorum device before you proceed. Otherwise, proceed to Step 3 on page 83.

Refer to *Sun Cluster 3.0 System Administration Guide* for procedures on replacing a quorum device and putting a quorum device into maintenance state.

3. If possible, back up the metadvice or volume.

For more information, see your Solstice DiskSuite or VERITAS Volume Manager documentation.

4. Perform volume management administration to remove the disk drive from the configuration.

For more information, see your Solstice DiskSuite or VERITAS Volume Manager documentation.

5. Identify the failed disk drive's physical DID.

Use this physical DID in Step 12 on page 84 to verify that the failed disk drive has been replaced with a new disk drive.

```
# sddidadm -o diskid -l cNtXdY
```

6. If you are using Solstice DiskSuite as your volume manager, save the disk partitioning for use when partitioning the new disk drive.

If you are using VERITAS Volume Manager, proceed to Step 7 on page 84.

```
# prtvtoc /dev/rdisk/cNtXdYsZ > filename
```



Caution - Do not save this file under /tmp because you will lose this file when you reboot. Instead, save this file under /usr/tmp.

7. Replace the failed disk drive.

For more information, see the documentation that shipped with your disk array.

8. On one node attached to the disk array, run the devfsadm(1M) command to probe all devices and to write the new disk drive to the /dev/rdisk directory.

Depending on the number of devices connected to the node, the devfsadm command can take at least five minutes to complete.

```
# devfsadm
```

9. If you are using Solstice DiskSuite as your volume manager, from any node connected to the disk array, partition the new disk drive, using the partitioning you saved in Step 6 on page 83.

If you are using VERITAS Volume Manager, proceed to Step 10 on page 84.

```
# fmthard -s filename /dev/rdisk/cNtXdYsZ
```

10. One at a time, shut down and reboot the nodes connected to the disk array.

```
# scswitch -S -h nodename  
# shutdown -y -g 0 -i 6
```

For more information on shutdown, see *Sun Cluster 3.0 System Administration Guide*.

11. From any node connected to the disk drive, update the DID database.

```
# scdidadm -R deviceID
```

12. From any node, confirm that the failed disk drive has been replaced by comparing the new physical DID to the physical DID in Step 5 on page 83.

If the new physical DID is different from the physical DID in Step 5 on page 83, you successfully replaced the failed disk drive with a new disk drive.

```
# sddidadm -o diskid -l cNtXdY
```

13. On all nodes, upload the new information to the DID driver.

If a volume management daemon such as vold is running on your node, and you have a CD-ROM drive connected to the node, a device busy error might be returned even if no disk is in the drive. This error is an expected behavior.

```
# sddidadm -ui
```

14. Perform volume management administration to add the disk drive back to its diskset or disk group.

For more information, see your Solstice DiskSuite or VERITAS Volume Manager documentation.

15. If you want this new disk drive to be a quorum device, add the quorum device.

Refer to *Sun Cluster 3.0 System Administration Guide* for the procedure on adding a quorum device.

Example—Replacing a StorEdge D1000 Disk Drive

The following example shows how to apply the procedure for replacing a StorEdge D1000 disk array disk drive.

```
# sddidadm -l d20
20      phys-schost-2:/dev/rdisk/c3t2d0 /dev/did/rdisk/d20
# sddidadm -o diskid -l c3t2d0
5345414741544520393735314336343734310000
# prtvtoc /dev/rdisk/c3t2d0s2 > /usr/tmp/c3t2d0.vtoc
...
# devfsadm
# fmthard -s /usr/tmp/c3t2d0.vtoc /dev/rdisk/c3t2d0s2
# scswitch -S -h node1
# shutdown -y -g 0 -i 6
...
# sddidadm -R d20
# sddidadm -o diskid -l c3t2d0
5345414741544520393735314336363037370000
```

(continued)

```
# scdidadm -ui
```

▼ How to Remove a StorEdge D1000 Disk Drive

Use this procedure to remove a disk drive from a StorEdge D1000 disk array. Perform the steps in this procedure in conjunction with the procedures in *Sun Cluster 3.0 System Administration Guide* and your server hardware manual.

For conceptual information on quorum, quorum devices, global devices, and device IDs, see *Sun Cluster 3.0 Concepts*.

1. Determine if the disk drive you want to remove is a quorum device.

```
# scstat -q
```

If the disk drive you want to remove is a quorum device, remove the quorum device before you proceed. Otherwise, proceed to Step 2 on page 86.

Refer to *Sun Cluster 3.0 System Administration Guide* for procedures on removing a quorum device and putting a quorum device into maintenance state.

2. If possible, back up the metadvice or volume.

For more information, see your Solstice DiskSuite or VERITAS Volume Manager documentation.

3. Perform volume management administration to remove the disk drive from the configuration.

For more information, see your Solstice DiskSuite or VERITAS Volume Manager documentation.

4. Identify the disk drive that needs to be removed and the slot that the disk drive needs to be removed from.

If the disk error message reports the drive problem by DID, use the `scdidadm -l` command to determine the Solaris device name.

```
# scdidadm -l deviceID
# cfgadm -al
```

5. Remove the disk drive.

For more information, see the documentation that shipped with the disk array.

6. On all nodes, remove references to the disk drive.

```
# cfgadm -c unconfigure cN::dsk/cNtXdY
# devfsadm -C
# scdidadm -C
```

▼ How to Add a StorEdge D1000 Disk Array

Use this procedure to install a disk array in an existing cluster. Perform the steps in this procedure in conjunction with the procedures in *Sun Cluster 3.0 System Administration Guide* and your server hardware manual.

For conceptual information on multi-initiator SCSI and device IDs, see *Sun Cluster 3.0 Concepts*.

1. Ensure that each device in the SCSI chain has a unique SCSI address.

The default SCSI address for host adapters is 7. Reserve SCSI address 7 for one host adapter in the SCSI chain. This procedure refers to the host adapter you choose for SCSI address 7 as the host adapter on the second node. To avoid conflicts, in Step 7 on page 77 you will change the `scsi-initiator-id` of the remaining host adapter in the SCSI chain to an available SCSI address. This procedure refers to the host adapter with an available SCSI address as the host adapter on the first node. SCSI address 6 is usually available.

For more information, see the *OpenBoot 3.x Command Reference Manual* and the labels inside the storage device.



Caution - Even though a slot in the disk array might not be in use, you should avoid setting the `scsi-initiator-id` for the first node to the SCSI address for that disk slot. This precaution minimizes future complications if you install additional disk drives.

2. Shut down and power off the first node.

```
# scswitch -S -h nodename  
# shutdown -y -g 0
```

For more information on shutdown, see *Sun Cluster 3.0 System Administration Guide*.

3. Install the host adapters in the node.

For the procedure on installing host adapters, see the documentation that shipped with your host adapters and nodes.

4. Connect the differential SCSI cable between the node and the disk array as shown in Figure 5-2.

Make sure that the *entire* SCSI bus length to each enclosure is less than 25 m. This measurement includes the cables to both nodes, as well as the bus length internal to each enclosure, node, and host adapter. Refer to the documentation that shipped with the enclosure for other restrictions regarding SCSI operation.




Figure 5-2 Example of a StorEdge D1000 disk array mirrored pair

5. Power on the first node and the disk arrays.

6. Find the paths to the SCSI host adapters.

```
{0} ok show-disks
```

Identify and record the two controllers that will be connected to the disk arrays, and record these paths. Use this information to change the SCSI addresses of these controllers in the `nvrarc` script. Do not include the `/sd` directories in the device paths.

7. Edit the `nvrarc` script to change the `scsi-initiator-id` for the host adapters of the first node.

For a list of Editor nvramrc editor and keystroke commands, see Appendix B.
The following example sets the `scsi-initiator-id` to 6. The OpenBoot PROM Monitor prints the line numbers (0:, 1:, and so on).



Caution - Insert exactly one space after the double quote and before `scsi-initiator-id`.

```
{0} ok nvedit
0: probe-all
1: cd /sbus@1f,0/QLGC,isp@3,10000
2: 6 encode-int " scsi-initiator-id" property
3: device-end
4: cd /sbus@1f,0/
5: 6 encode-int " scsi-initiator-id" property
6: device-end
7: install-console
8: banner [Control C]
{0} ok
```

8. Store the changes.

The changes you make through the `nvedit` command are done on a temporary copy of the `nvramrc` script. You can continue to edit this copy without risk. After you have completed your edits, save the changes. If you are not sure about the changes, discard them.

- To store the changes, type:

```
{0} ok nvstore
{0} ok
```

- To discard the changes, type:

```
{0} ok nvquit
{0} ok
```

9. Verify the contents of the `nvramrc` script you created in Step 7 on page 89.

If the contents of the `nvramrc` script are incorrect, use the `nvedit` command to make corrections.

```
{0} ok printenv nvramrc
nvramrc =                probe-all
                           cd /sbus@1f,0/QLGC,isp@3,10000
                           6 encode-int " scsi-initiator-id" property
                           device-end
                           cd /sbus@1f,0/
                           6 encode-int " scsi-initiator-id" property
                           device-end
                           install-console
                           banner
{0} ok
```

10. Instruct the OpenBoot PROM Monitor to use the nvramrc script:

```
{0} ok setenv use-nvramrc? true
use-nvramrc? = true
{0} ok
```

11. Boot the first node, and wait for it to join the cluster.

```
{0} ok boot -r
```

For more information, see *Sun Cluster 3.0 System Administration Guide*.

12. On all nodes, verify that the DIDs have been assigned to the disk drives in the disk array.

```
# scdidadm -l
```

13. Shut down the second node.

```
# scswitch -S -h nodename
# shutdown -y -g 0
```

14. Power off the second node.

15. Install the host adapters in the second node.

For the procedure on installing host adapters, see the documentation that shipped with your nodes.

16. Connect the disk array to the host adapters using differential SCSI cables.




Figure 5-3 Example of a StorEdge D1000 disk array mirrored pair

17. Without allowing the node to boot, power on the second node. If necessary, abort the system to continue with OpenBoot PROM Monitor tasks.

18. Verify that the second node sees the new host adapters and disk drives.

```
{0} ok show-disks
```

19. Verify that the scsi-initiator-id for the host adapters on the second node is set to 7.

Use the `show-disks` command to find the paths to the host adapters connected to these enclosures. Select each host adapter's device tree node, and display the node's properties to confirm that the `scsi-initiator-id` for each host adapter is set to 7.

```
{0} ok cd /sbus@1f,0/QLGC,isp@3,10000
{0} ok .properties
scsi-initiator-id      00000007
...
```

20. Boot the second node, and wait for it to join the cluster.

```
{0} ok boot -r
```

21. On all nodes, verify that the DIDs have been assigned to the disk drives in the disk array.

```
# sctidadm -l
```

22. Perform volume management administration to add the disk drives in the array to the volume management configuration.

For more information, see your Solstice DiskSuite or VERITAS Volume Manager documentation.

▼ How to Replace a StorEdge D1000 Disk Array

Use this procedure to replace a StorEdge D1000 disk array. This procedure assumes that you want to retain the disk drives in the disk array you want to replace and retain the references to these same disk drives.

If you want to replace your disk drives, see “How to Replace a StorEdge D1000 Disk Drive” on page 82.

1. If possible, back up the metadevices or volumes that reside in the disk array.

For more information, see your Solstice DiskSuite or VERITAS Volume Manager documentation.

2. Perform volume management administration to remove the disk array from the configuration.

For more information, see your Solstice DiskSuite or VERITAS Volume Manager documentation.

3. Disconnect the SCSI cables from the disk array, removing them in the order shown in Figure 5-4.




Figure 5-4 Disconnecting the SCSI cables

4. Power off, and disconnect the disk array from the AC power source.

For more information, see the documentation that shipped with your disk array.

5. Connect the new disk array to an AC power source.

For more information, see the documentation that shipped with your disk array.

6. Connect the SCSI cables to the new disk array by reversing the steps shown in Figure 5-4.

7. **Move the disk drives one at a time from the old disk array to the same slot in the new disk array.**
8. **Power on the disk array.**
9. **On all nodes attached to the disk array, run the `devfsadm(1M)` command.**

```
# devfsadm
```

10. **One at a time, shut down and reboot the nodes connected to the disk array.**

```
# scswitch -S -h nodename  
# shutdown -y -g 0 -i 6
```

For more information on shutdown, see *Sun Cluster 3.0 System Administration Guide*.

11. **Perform volume management administration to add the disk array to the configuration.**

For more information, see your Solstice DiskSuite or VERITAS Volume Manager documentation.

▼ How to Remove a StorEdge D1000 Disk Array

Use this procedure to remove a StorEdge D1000 disk array from a cluster. This procedure assumes that you want to remove the references to the disk drives in the enclosure.

1. **Perform volume management administration to remove the disk array from the configuration.**

For more information, see your Solstice DiskSuite or VERITAS Volume Manager documentation.

2. **Disconnect the SCSI cables from the disk array, removing them in the order shown in Figure 5–5.**




Figure 5-5 Disconnecting the SCSI cables

3. Power off, and disconnect the disk array from the AC power source.

For more information, see the documentation shipped with your disk array.

4. Remove the disk array.

5. Identify the disk drives you need to remove.

```
# cfgadm -al
```

6. On all nodes, remove references to the disk drives in the disk array you removed in Step 4 on page 96.

```
# cfgadm -c unconfigure cN::dsk/cNtXdY
# devfsadm -C
# scdidadm -C
```

7. If needed, remove any host adapters from the nodes.

For the procedure on removing host adapters, see the documentation that shipped with your nodes.

Installing and Maintaining a Sun StorEdge A5x00 Disk Array

This chapter provides the procedures for installing and maintaining a Sun StorEdge™ A5x00 disk array.

This chapter contains the following procedures:

- “How to Install a StorEdge A5x00” on page 100
- “How to Add a StorEdge A5x00 Disk Drive” on page 102
- “How to Replace a StorEdge A5x00 Disk Drive” on page 104
- “How to Remove a StorEdge A5x00 Disk Drive” on page 108
- “How to Add a StorEdge A5x00 Disk Array to an Existing Cluster” on page 109
- “How to Replace a StorEdge A5x00 Disk Array” on page 112
- “How to Remove a StorEdge A5x00 Disk Array” on page 114

For conceptual information on multihost disks, see *Sun Cluster 3.0 Concepts*.

Installing a StorEdge A5x00

This section provides the procedure for an initial installation of an StorEdge A5x00 disk array. The following table lists the steps involved in an initial installation of an StorEdge A5x00 disk array.

TABLE 6-1 Task Map: Installing a StorEdge A5x00

Task	For Instructions, Go To...
Install the host adapters	The documentation that shipped with your nodes
Cable and configure the disk array	<i>Sun StorEdge A5000 Installation and Service Manual</i>
Check the hardware firmware levels, and install any required firmware updates	<i>Sun Cluster 3.0 Release Notes</i>

▼ How to Install a StorEdge A5x00

Use this procedure to install an StorEdge A5x00 disk array. Perform the steps in this procedure in conjunction with the procedures in *Sun Cluster 3.0 Installation Guide* and your server hardware manual.

1. Install host adapters in the nodes that will be connected to the disk array.

For the procedure on installing host adapters, see the documentation that shipped with your network adapters and nodes.

Note - To ensure maximum redundancy, put each host adapter on a separate I/O board if possible.

2. Cable, power on, and configure the disk array.

Figure 6-1 shows a sample disk array configuration.

For more information on cabling and configuring disk arrays, see *Sun StorEdge A5000 Installation and Service Manual*.




Figure 6-1 Sample StorEdge A5x00 Disk Array Configuration

3. Check the hardware firmware levels, and install any required firmware updates.

For the location of patches and installation instructions, see *Sun Cluster 3.0 Release Notes*.

Where to Go From Here

To install software, follow the procedures in *Sun Cluster 3.0 Installation Guide*.

Maintaining a StorEdge A5x00 Disk Array

This section provides the procedures for maintaining an StorEdge A5x00 disk array. The following table lists these procedures.

TABLE 6-2 Task Map: Maintaining a StorEdge A5x00 Disk Array

Task	For Instructions, Go To...
Perform an initial installation Use the <code>scgdevs(1M)</code> command to update the global device namespace without a reconfiguration reboot.	“How to Install a StorEdge A5x00” on page 100
Add a disk drive Use the <code>scswitch(1M)</code> , <code>shutdown(1M)</code> , and <code>luxadm insert</code> commands to add a disk drive to a disk array.	“How to Add a StorEdge A5x00 Disk Drive” on page 102
Replace a disk drive	“How to Replace a StorEdge A5x00 Disk Drive” on page 104
Remove a disk drive	“How to Remove a StorEdge A5x00 Disk Drive” on page 108
Add a disk array to an existing cluster Use the <code>pkgadd(1M)</code> , <code>shutdown(1M)</code> , <code>scswitch(1M)</code> , and <code>luxadm insert</code> commands to add a disk array to an existing cluster.	“How to Add a StorEdge A5x00 Disk Array to an Existing Cluster” on page 109
Replace a disk array in an existing cluster Remove a disk array from the cluster configuration, and replace it with a new disk array.	“How to Replace a StorEdge A5x00 Disk Array” on page 112
Remove a disk array from an existing cluster Use <code>luxadm remove</code> and <code>devfsadm -C</code> commands to remove a disk array without replacing it with another disk array.	“How to Remove a StorEdge A5x00 Disk Array” on page 114

▼ How to Add a StorEdge A5x00 Disk Drive

Use this procedure to add a disk drive to an existing cluster. Perform the steps in this procedure in conjunction with the procedures in *Sun Cluster 3.0 System Administration Guide* and your server hardware manual.

For conceptual information on quorum, quorum devices, global devices, and device IDs, see *Sun Cluster 3.0 Concepts*.

- 1. On one node connected to the disk array, use the `luxadm(1M)` command to insert the new disk.**

Physically install the new disk drive, and press Return when prompted. Using the `luxadm insert` command, you can insert multiple disk drives at the same time.

```
# luxadm insert enclosure,slot
```

2. **On all other nodes attached to the disk array, run the `devfsadm(1M)` command to probe all devices and to write the new disk drive to the `/dev/rdisk` directory.**

Depending on the number of devices connected to the node, the `devfsadm` command can take at least five minutes to complete.

```
# devfsadm
```

3. **Ensure that entries for the disk drive have been added to the `/dev/rdisk` directory.**

```
# ls -l /dev/rdisk
```

4. **If needed, partition the disk drive.**

You can use either the `format(1M)` command or copy the partitioning from another disk drive in the disk array.

5. **From any node in the cluster, update the global device namespace.**

If a volume management daemon such as `vold` is running on your node, and you have a CD-ROM drive connected to the node, a device busy error might be returned even if no disk is in the drive. This error is an expected behavior.

```
# scgdevs
```

6. **Verify that a device ID (DID) has been assigned to the disk drive.**

```
# sddidadm -l
```

Note - The DID assigned to the new disk drive might not be in sequential order in the disk array.

7. **Perform the usual volume management administration actions on the new disk drive.**

For more information, see your Solstice DiskSuite or VERITAS Volume Manager documentation.

8. If you want this new disk drive to be a quorum device, add the quorum device.

Refer to *Sun Cluster 3.0 System Administration Guide* for the procedure on adding a quorum device.

▼ How to Replace a StorEdge A5x00 Disk Drive

Use this procedure to replace an StorEdge A5x00 disk array disk drive.

“Example—Replacing a StorEdge A5x00 Disk Drive” on page 107 shows you how to apply this procedure. Perform the steps in this procedure in conjunction with the procedures in *Sun Cluster 3.0 System Administration Guide* and your server hardware manual. Use the procedures in your server hardware manual to identify a failed disk drive.

For conceptual information on quorum, quorum devices, global devices, and device IDs, see *Sun Cluster 3.0 Concepts*.

1. Identify the disk drive that needs replacement.

If the disk error message reports the drive problem by device ID (DID), use the `scdidadm -l` command to determine the Solaris logical device name. If the disk error message reports the drive problem by the Solaris physical device name, use your Solaris documentation to map the Solaris physical device name to the Solaris logical device name. Use this Solaris logical device name and DID throughout this procedure.

```
# scdidadm -l deviceID
```

2. Determine if the disk drive you want to replace is a quorum device.

```
# scstat -q
```

If the disk drive you want to replace is a quorum device, remove the quorum device before you proceed. Otherwise, proceed to Step 3 on page 104.

Refer to *Sun Cluster 3.0 System Administration Guide* for procedures on replacing a quorum device and putting a quorum device into maintenance state.

3. If possible, back up the metadvice or volume.

For more information, see your Solstice DiskSuite or VERITAS Volume Manager documentation.

4. Perform volume management administration to remove the disk drive from the configuration.

For more information, see your Solstice DiskSuite or VERITAS Volume Manager documentation.

5. Identify the failed disk drive's physical DID.

Use this physical DID in Step 14 on page 106 to verify that the failed disk drive has been replaced with a new disk drive. The DID and the World Wide Name (WWN) for the disk drive should be the same.

```
# sddidadm -o diskid -l cNtXdY
```

6. If you are using Solstice DiskSuite as your volume manager, save the disk partitioning for use when partitioning the new disk drive.

If you are using VERITAS Volume Manager, proceed to Step 7 on page 105.

```
# prtvtoc /dev/rdisk/cNtXdYsZ > filename
```

7. On any node connected to the disk array, run the luxadm remove command.

```
# luxadm remove -F /dev/rdisk/cNtXdYsZ
```

8. Replace the failed disk drive.

For more information, see the documentation that shipped with your disk array.

9. On any node connected to the disk array, run the luxadm insert command.

```
# luxadm insert boxname,rslotnumber  
# luxadm insert boxname,fslotnumber
```

If you want to insert a front disk drive, use the fslotnumber parameter. If you want to insert a rear disk drive, use the rslotnumber parameter.

10. On all other nodes attached to the disk array, run the devfsadm(1M) command to probe all devices and to write the new disk drive to the /dev/rdisk directory.

Depending on the number of devices connected to the node, the devfsadm command can take at least five minutes to complete.

```
# devfsadm
```

- 11. If you are using Solstice DiskSuite as your volume manager, on one node connected to the disk array, partition the new disk drive, using the partitioning you saved in Step 6 on page 105.**

If you are using VERITAS Volume Manager, proceed to Step 12 on page 106.

```
# fmthard -s filename /dev/rdisk/cNtXdYsZ
```

- 12. One at a time, shut down and reboot the nodes connected to the disk array.**

```
# scswitch -S -h nodename  
# shutdown -y -g 0 -i 6
```

For more information, see *Sun Cluster 3.0 System Administration Guide*.

- 13. On any of the nodes connected to the disk array, update the DID database.**

```
# scdidadm -R deviceID
```

- 14. On any node, confirm that the failed disk drive has been replaced by comparing the following physical DID to the physical DID in Step 5 on page 105.**

If the following physical DID is different from the physical DID in Step 5 on page 105, you successfully replaced the failed disk drive with a new disk drive.

```
# scdidadm -o diskid -l cNtXdY
```

- 15. On all nodes connected to the disk array, upload the new information to the DID driver.**

If a volume management daemon such as vold is running on your node, and you have a CD-ROM drive connected to the node, a device busy error might be returned even if no disk is in the drive. This error is an expected behavior.

```
# scdidadm -ui
```

- 16. Perform volume management administration to add the disk drive back to its diskset or disk group.**

For more information, see your Solstice DiskSuite or VERITAS Volume Manager documentation.

17. If you want this new disk drive to be a quorum device, add the quorum device.

Refer to *Sun Cluster 3.0 System Administration Guide* for the procedure on adding a quorum device.

Example—Replacing a StorEdge A5x00 Disk Drive

The following example shows how to apply the procedure for replacing an StorEdge A5x00 disk array disk drive.

```
# scstat -q
# sddidadm -l d4
4      phys-schost-2:/dev/rdisk/c1t32d0 /dev/did/rdisk/d4
# sddidadm -o diskid -l c1t32d0
20000002037000edf
# prtvtoc /dev/rdisk/c1t32d0s2 > /usr/tmp/c1t32d0.vtoc
# luxadm remove -F /dev/rdisk/c1t32d0s2
WARNING!!! Please ensure that no filesystems are mounted on these device(s). All data on these devices should have been

The list of devices that will be removed is:  1: Box Name "venus1" front slot 0

Please enter 'q' to Quit or <Return> to Continue: <Return>

stopping: Drive in "venus1" front slot 0....Done
offlining: Drive in "venus1" front slot 0....Done

Hit <Return> after removing the device(s). <Return>

Drive in Box Name "venus1" front slot 0
Logical Nodes being removed under /dev/dsk/ and /dev/rdisk:
    c1t32d0s0
    c1t32d0s1
    c1t32d0s2
    c1t32d0s3
    c1t32d0s4
    c1t32d0s5
    c1t32d0s6
    c1t32d0s7

# devfsadm
# fmthard -s /usr/tmp/c1t32d0.vtoc /dev/rdisk/c1t32d0s2
# scswitch -S -h node1
# shutdown -y -g 0 -i 6
# sddidadm -R d4
# sddidadm -o diskid -l c1t32d0
200000020370bf955
# sddidadm -ui
```

▼ How to Remove a StorEdge A5x00 Disk Drive

Use this procedure to remove a disk drive from a disk array. “Example—Removing a StorEdge A5x00 Disk Drive” on page 109 shows you how to apply this procedure. Perform the steps in this procedure in conjunction with the procedures in *Sun Cluster 3.0 System Administration Guide* and your server hardware manual.

For conceptual information on quorum, quorum devices, global devices, and device IDs, see *Sun Cluster 3.0 Concepts*.

1. Determine if the disk drive you want to remove is a quorum device.

```
# scstat -q
```

If the disk drive you want to remove is a quorum device, remove the quorum device before you proceed. Otherwise, proceed to Step 2 on page 108.

Refer to *Sun Cluster 3.0 System Administration Guide* for procedures on removing a quorum device and putting a quorum device into maintenance state.

2. If possible, back up the metadvice or volume.

For more information, see your Solstice DiskSuite or VERITAS Volume Manager documentation.

3. Perform volume management administration to remove the disk drive from the configuration.

For more information, see your Solstice DiskSuite or VERITAS Volume Manager documentation.

4. Identify the disk drive that needs to be removed.

If the disk error message reports the drive problem by DID, use the `scdidadm -l` command to determine the Solaris device name.

```
# scdidadm -l deviceID
```

5. On any node connected to the disk array, run the `luxadm remove` command.

Physically remove the disk drive, and press Return when prompted.

```
# luxadm remove -F /dev/rdisk/cNtXdYsZ
```

6. On all connected nodes, remove references to the disk drive.

```
# devfsadm -C
# scdidadm -C
```

Example—Removing a StorEdge A5x00 Disk Drive

The following example shows how to apply the procedure for removing an StorEdge A5x00 disk array disk drive.

```
# scdidadm -l d4
4      phys-schost-2:/dev/rdisk/c1t32d0 /dev/did/rdisk/d4
# luxadm remove -F /dev/rdisk/c1t32d0s2

WARNING!!! Please ensure that no filesystems are mounted on these device(s). All data on these devices should have been

The list of devices that will be removed is:  1: Box Name "venus1" front slot 0

Please enter 'q' to Quit or <Return> to Continue: <Return>

stopping: Drive in "venus1" front slot 0....Done
offlining: Drive in "venus1" front slot 0....Done

Hit <Return> after removing the device(s). <Return>

Drive in Box Name "venus1" front slot 0
Logical Nodes being removed under /dev/dsk/ and /dev/rdisk:
    c1t32d0s0
    c1t32d0s1
    c1t32d0s2
    c1t32d0s3
    c1t32d0s4
    c1t32d0s5
    c1t32d0s6
    c1t32d0s7
# devfsadm -C
# scdidadm -C
```

▼ How to Add a StorEdge A5x00 Disk Array to an Existing Cluster

Use this procedure to install a StorEdge A5x00 disk array in an existing cluster. Perform the steps in this procedure in conjunction with the procedures in *Sun Cluster 3.0 System Administration Guide* and your server hardware manual.

1. **Determine if the StorEdge A5x00 disk array packages need to be installed on the nodes to which you are connecting the disk array. The following packages are required.**

```
# pkginfo | egrep Wlux
system    SUNWluxd    Sun Enterprise Network Array sf Device Driver
system    SUNWluxdx    Sun Enterprise Network Array sf Device Driver (64-bit)
system    SUNWluxl    Sun Enterprise Network Array socal Device Driver
system    SUNWluxlx    Sun Enterprise Network Array socal Device Driver (64-bit)
system    SUNWluxop    Sun Enterprise Network Array firmware and utilities
```

2. **On each node, install any needed packages for the Solaris operating environment.**

The disk array packages are located in the Product directory of the CD-ROM. Use the pkgadd command to add any necessary packages.

```
# pkgadd -d path_to_Solaris/Product Pkg1 Pkg2 Pkg3 ... PkgN
```

path_to_Solaris Path to the Solaris operating environment

Pkg1 Pkg2 The packages to be added

3. **Shut down and power off any node connected to the disk array.**

```
# scswitch -S -h nodename
# shutdown -y -g 0
```

Refer to *Sun Cluster 3.0 System Administration Guide* for more information.

4. **Install host adapters in the node that will be connected to the disk array.**

For the procedure on installing host adapters, see the documentation that shipped with your network adapters and nodes.

5. **Cable, configure, and power on the disk array.**

For more information, see the documentation that shipped with your disk array. Figure 6–2 shows a sample disk array configuration.




Figure 6-2 Sample StorEdge A5x00 Disk Array Configuration

6. Power on and boot the node.

```
# boot -r
```

For the procedures on powering on and booting a node, see *Sun Cluster 3.0 System Administration Guide*.

7. Determine if any patches need to be installed on the node(s) that will be connected to the disk array.

For a list of Sun Cluster-specific patches, see *Sun Cluster 3.0 Release Notes*.

8. Obtain and install any needed patches on the nodes that will be connected to the disk array.

For procedures on applying patches, see *Sun Cluster 3.0 System Administration Guide*.



Caution - Read any README files that accompany the patches before you begin this installation. Some patches must be installed in a specific order.

9. If required by the patch README instructions, shut down and reboot the node.

```
# scswitch -S -h nodename  
# shutdown -y -g 0 -i 6
```

10. Perform Step 3 on page 110 through Step 9 on page 111 for each node attached to the disk array.

11. Perform volume management administration to add the disk drives in the array to the volume management configuration.

For more information, see your Solstice DiskSuite or VERITAS Volume Manager documentation.

▼ How to Replace a StorEdge A5x00 Disk Array

Use this procedure to replace a failed StorEdge A5x00 disk array.

“Example—Replacing a StorEdge A5x00 Disk Array” on page 113 shows you how to apply this procedure. This procedure assumes that you want to retain the disk drives.

If you want to replace your disk drives, see “How to Replace a StorEdge A5x00 Disk Drive” on page 104.

1. If possible, back up the metadevices or volumes that reside in the disk array.

For more information, see your Solstice DiskSuite or VERITAS Volume Manager documentation.

2. Perform volume management administration to remove the disk array from the configuration.

For more information, see your Solstice DiskSuite or VERITAS Volume Manager documentation.

3. On all nodes connected to the disk array, run the `luxadm remove` command.

```
# luxadm remove -F boxname
```

4. Disconnect the fiber optic cables from the disk array.

5. Power off, and disconnect the disk array from the AC power source.

For more information, see the documentation that shipped with your disk array.

6. Connect the fiber optic cables to the new disk array.

7. Connect the new disk array to an AC power source.

For more information, see the documentation that shipped with your disk array.

8. One at a time, move the disk drives from the old disk array to the same slot in the new disk array.

9. Power on the disk array.

10. Use the `luxadm insert` command to find the new disk array.

Repeat this step for each node connected to the disk array.

```
# luxadm insert
```

11. On all nodes connected to the new disk array, upload the new information to the DID driver.

If a volume management daemon such as `vold` is running on your node, and you have a CD-ROM drive connected to the node, a device busy error might be returned even if no disk is in the drive. This error is an expected behavior.

```
# scgdevs
```

12. Perform volume management administration to add the new disk array to the configuration.

For more information, see your Solstice DiskSuite or VERITAS Volume Manager documentation.

Example—Replacing a StorEdge A5x00 Disk Array

The following example shows how to apply the procedure for replacing an StorEdge A5x00 disk array.

```
# luxadm remove -F venus1
```

```
WARNING!!! Please ensure that no filesystems are mounted on these device(s).  
All data on these devices should have been backed up.
```

```
The list of devices that will be removed is:
```

```
1: Box name:      venus1  
   Node WWN:     123456789abcdeff  
   Device Type:  SENA (SES device)  
   SES Paths:  
     /devices/nodes@1/sbus@1f,0/SUNW,socal@1,0/sf@0,0/ses@w123456789abcdf03,0:0  
     /devices/nodes@1/sbus@1f,0/SUNW,socal@1,0/sf@1,0/ses@w123456789abcdf00,0:0
```

(continued)

```
Please verify the above list of devices and
then enter 'c' or <CR> to Continue or 'q' to Quit. [Default: c]: <Return>
```

```
Hit <Return> after removing the device(s). <Return>
# luxadm insert
Please hit <RETURN> when you have finished adding Fibre Channel
Enclosure(s)/Device(s): <Return>
# scgdevs
```

▼ How to Remove a StorEdge A5x00 Disk Array

Use this procedure to remove an StorEdge A5x00 disk array from a cluster. “Example—Removing a StorEdge A5x00 Disk Array” on page 115 shows you how to apply this procedure. Use the procedures in your server hardware manual to identify the disk array.

1. Perform volume management administration to remove the disk array from the configuration.

For more information, see your Solstice DiskSuite or VERITAS Volume Manager documentation.

2. On all nodes connected to the disk array, run the `luxadm remove` command.

```
# luxadm remove -F boxname
```

3. Remove the disk array and the fiber optic cables connected to the disk array.

Refer to *Sun StorEdge A5000 Installation and Service Manual*.

4. On all nodes, remove references to the disk array.

```
# devfsadm -C
# scdidadm -C
```

5. If needed, remove any host adapters from the nodes.

For the procedure on removing host adapters, see the documentation that shipped with your nodes.

Example—Removing a StorEdge A5x00 Disk Array

The following example shows how to apply the procedure for removing an StorEdge A5x00 disk array.

```
# luxadm remove -F venus1
WARNING!!! Please ensure that no filesystems are mounted on these device(s).
All data on these devices should have been backed up.

The list of devices that will be removed is:
1: Box name:      venus1
   Node WWN:      123456789abcdeff
   Device Type:   SENA (SES device)
   SES Paths:
     /devices/nodes@1/sbus@1f,0/SUNW,socal@1,0/sf@0,0/ses@w123456789abcdf03,0:0
     /devices/nodes@1/sbus@1f,0/SUNW,socal@1,0/sf@1,0/ses@w123456789abcdf00,0:0

Please verify the above list of devices and
then enter 'c' or <CR> to Continue or 'q' to Quit. [Default: c]: <Return>

Hit <Return> after removing the device(s). <Return>

# devfsadm -C
# scdidadm -C
```


Installing, Configuring, and Maintaining a Sun StorEdge A3500 Disk Array

This chapter provides the procedures for installing, configuring, and maintaining a Sun StorEdge™ A3500 disk array.

This chapter contains the following procedures:

- “How to Install a StorEdge A3500” on page 118
- “How to Create a LUN” on page 125
- “How to Remove a LUN” on page 126
- “How to Reset StorEdge A3500 Configuration” on page 126
- “How to Add a StorEdge A3500” on page 130
- “How to Remove a StorEdge A3500” on page 136
- “How to Replace a Failed StorEdge A3500 Controller or Restore an Offline StorEdge A3500 Controller” on page 137
- “How to Replace a Host Adapter” on page 139

Installing a StorEdge A3500

This section provides the procedure for an initial installation of a StorEdge A3500 disk array. The following table lists the steps involved in an initial installation of an StorEdge A3500 disk array. Perform these procedures in the order that they are listed.

TABLE 7-1 Task Map: Installing a StorEdge A3500

Task	For Instructions, Go To...
Install the host adapters	The documentation that shipped with your host adapters and nodes
Ensure that each device in the SCSI chain has a unique SCSI address	“How to Install a StorEdge A3500” on page 118
Cable, power on, and configure the disk array	<i>Sun StorEdge A3500/A3500FC Hardware Configuration Guide</i> <i>Sun StorEdge A3500/A3500FC Controller Module Guide</i>
Install the Solaris operating environment	<i>Sun Cluster 3.0 Installation Guide</i>
Apply the required Solaris patches	<i>Sun Cluster 3.0 Release Notes</i>
Install the RAID Manager	<i>Sun StorEdge RAID Manager Installation and Support Guide</i>
Install StorEdge A3500 patch(es)	<i>Sun Cluster 3.0 Release Notes</i>
Upgrade the StorEdge A3500 controller firmware	<i>Sun StorEdge RAID Manager User's Guide</i>
Set up the StorEdge A3500 with the desired LUNs and configuration	<i>Sun StorEdge RAID Manager User's Guide</i>
Continue with Sun Cluster software and data services installation tasks	<i>Sun Cluster 3.0 Installation Guide</i> <i>Sun Cluster 3.0 Data Services Installation and Configuration Guide</i>

▼ How to Install a StorEdge A3500

Use this procedure for an initial installation and configuration of an StorEdge A3500 disk array, prior to installing the Solaris operating environment and Sun Cluster software. Perform the steps in this procedure in conjunction with the procedures in *Sun Cluster 3.0 Installation Guide* and your server hardware manual.

1. Ensure that each device in the SCSI chain has a unique SCSI address.

The default SCSI address for host adapters is 7. Reserve SCSI address 7 for one host adapter in the SCSI chain. This procedure refers to the host adapter you choose for SCSI address 7 as the host adapter on the second node. To avoid conflicts, in Step 5 on page 120 you will change the `scsi-initiator-id` of the remaining host adapter in the SCSI chain to an available SCSI address. This procedure refers to the host adapter with an available SCSI address as the host adapter on the first node. Depending on the device and configuration settings of the device, either SCSI address 6 or 8 is usually available.



Caution - Even though a slot in the enclosure might not be in use, you should avoid setting the `scsi-initiator-id` for the first node to the SCSI address for that disk slot. This precaution minimizes future complications if you install additional disk drives.

For more information, see the *OpenBoot 3.x Command Reference Manual* and the labels inside the storage device.

2. Install the host adapters in the nodes that will be connected to the disk array.

For the procedure on installing host adapters, see the documentation that shipped with your host adapters and nodes.

3. Cable, power on, and configure the disk array as shown in Figure 7-1.

For the procedure on installing the SCSI cables, see *Sun StorEdge A3500/A3500FC Hardware Configuration Guide*. For the procedure on powering on the disk array, see *Sun StorEdge A3500/A3500FC Controller Module Guide*.




Figure 7-1 Example of a StorEdge A3500 disk array

4. Find the paths to the host adapters.

```
{0} ok show-disks
```

Identify and record the two controllers that will be connected to the storage devices and record these paths. You will use this information to change the SCSI addresses of these controllers in the nvramrc script. Do not include the /sd directories in the device paths.

5. Edit the nvramrc script to change the scsi-initiator-id for the host adapters on the first node.

For a list of nvramrc editor and nvedit keystroke commands, see Appendix B.

The following example sets the scsi-initiator-id to 6. The OpenBoot PROM Monitor prints the line numbers (0:, 1:, and so on).



Caution - Insert exactly one space after the first double quote and before scsi-initiator-id.

```
{0} ok nvedit
0: probe-all
1: cd /sbus@1f,0/QLGC,isp@3,10000
2: 6 encode-int " scsi-initiator-id" property
3: device-end
4: cd /sbus@1f,0/
5: 6 encode-int " scsi-initiator-id" property
6: device-end
7: install-console
8: banner [Control C]
{0} ok
```

6. Store the changes.

The changes you make through the `nvedit` command are done on a temporary copy of the `nvrarc` script. You can continue to edit this copy without risk. After you complete your edits, save the changes. If you are not sure about the changes, discard them.

- To store the changes, type:

```
{0} ok nvstore
{0} ok
```

- To discard the changes, type:

```
{0} ok nvquit
{0} ok
```

7. Verify the contents of the `nvrarc` script you created in Step 5 on page 120.

If the contents of the `nvrarc` script are incorrect, use the `nvedit` command to make corrections.

```
{0} ok printenv nvrarc
nvrarc =
        probe-all
        cd /sbus@1f,0/QLGC,isp@3,10000
        6 encode-int " scsi-initiator-id" property
        device-end
        cd /sbus@1f,0/
        6 encode-int " scsi-initiator-id" property
        device-end
        install-console
        banner
{0} ok
```

8. Instruct the OpenBoot PROM Monitor to use the nvramrc script.

```
{0} ok setenv use-nvramrc? true
use-nvramrc? = true
{0} ok
```

9. Without allowing the node to boot, power on the second node. If necessary, abort the system to continue with OpenBoot PROM Monitor tasks.**10. Verify that the scsi-initiator-id for each host adapter on the second node is set to 7.**

Use the `show-disks` command to find the paths to the host adapters connected to these enclosures. Select each host adapter's device tree node, and display the node's properties to confirm that the `scsi-initiator-id` for each host adapter is set to 7.

```
{0} ok cd /sbus@1f,0/QLGC,isp@3,10000
{0} ok .properties
scsi-initiator-id      00000007
```

11. Install the Solaris™ operating environment, and apply the required Solaris patches.

For the procedure on installing the Solaris operating environment, see *Sun Cluster 3.0 Installation Guide*. For the location of patches and installation instructions, see *Sun Cluster 3.0 Release Notes*.

12. Install the RAID Manager.

For the procedure on installing the RAID Manager, see *Sun StorEdge RAID Manager Installation and Support Guide*.

13. Install StorEdge A3500 disk array patches.

For the location of patches and installation instructions, see *Sun Cluster 3.0 Release Notes*.

14. Upgrade the StorEdge A3500 disk array controller firmware.

For the StorEdge A3500 disk array controller firmware version number and boot level, see *Sun Cluster 3.0 Release Notes*. For the procedure on upgrading the StorEdge A3500 disk array controller firmware, see *Sun StorEdge RAID Manager User's Guide*.

15. Set up the A3500 disk array with the desired LUNs and hot spares.

For the procedure on setting up the StorEdge A3500 with LUNs and hot spares, see *Sun StorEdge RAID Manager User's Guide*.

Note - The RAID Manager 6.x graphical user interface does not consistently display Solaris logical device names. Use the `format` command to verify Solaris logical device names.

Where to Go From Here

To continue with Sun Cluster software and data services installation tasks, see *Sun Cluster 3.0 Installation Guide* and *Sun Cluster 3.0 Data Services Installation and Configuration Guide*.

Configuring StorEdge A3500 Disk Drives After Sun Cluster Install

This section provides the procedures for configuring an StorEdge A3500 disk array. The following table lists these procedures. Perform these procedures after installing Sun Cluster software.

Configuring StorEdge A3500 disk array disk drives before installing Sun Cluster software is no different than doing so in a non-cluster environment. For procedures on configuring StorEdge A3500 disk drives before installing Sun Cluster, see *Sun StorEdge RAID Manager User's Guide*.

TABLE 7-2 Task Map: Configuring StorEdge A3500 Disk Drives

Task	For Instructions, Go To...
Create a logical unit number (LUN) On one node, follow the same procedure used in a non-cluster environment, then use the <code>scgdevs</code> command to update the global device namespace.	<i>Sun StorEdge RAID Manager User's Guide</i> "How to Create a LUN" on page 125
Remove a LUN On one node, follow the same procedure used in a non-cluster environment, then use the <code>scdidadm -C</code> command to remove obsolete DIDs.	<i>Sun StorEdge RAID Manager User's Guide</i> "How to Remove a LUN" on page 126
Rebalance existing LUNs Follow the same procedure used in a non-cluster environment.	<i>Sun StorEdge RAID Manager User's Guide</i>
Reset StorEdge A3500 configuration On one node, follow the same procedure used in a non-cluster environment, then use the <code>scgdevs</code> command to update the global device namespace, and the <code>scdidadm -C</code> command to remove obsolete DIDs.	<i>Sun StorEdge RAID Manager User's Guide</i> "How to Reset StorEdge A3500 Configuration" on page 126
Create a hot spare Follow the same procedure used in a non-cluster environment.	<i>Sun StorEdge RAID Manager User's Guide</i>
Delete a hot spare Follow the same procedure used in a non-cluster environment.	<i>Sun StorEdge RAID Manager User's Guide</i>
Increase the drive capacity of a LUN If a volume manager manages the LUN, run the appropriate Solstice DiskSuite or VERITAS Volume Manager commands to remove the LUN from any diskset or disk group, then follow the same procedure used in a non-cluster environment. Afterward, you can return the LUN to a diskset or disk group.	<i>Sun StorEdge RAID Manager User's Guide</i>

▼ How to Create a LUN

Use this procedure to create a logical unit number (LUN) from unassigned disk drives or remaining capacity.

This product supports the use of hardware RAID and host-based software RAID. For host-based software RAID, this product only supports RAID levels 0+1 and 1+0. You must use hardware RAID for Oracle Parallel Server (OPS) data stored on the StorEdge A3500. Also, do not place OPS data under volume management control. You must place all other data stored on the StorEdge A3500 under volume management control. Use either hardware RAID, host-based software RAID, or both to manage your non-OPS data.

Hardware RAID uses the StorEdge A3500's hardware redundancy to ensure that independent hardware failures do not impact data availability. Host-based software RAID ensures that independent hardware failures do not impact data availability when an entire StorEdge A3500 is offline by mirroring across separate StorEdge A3500 disk arrays. Although you can use hardware RAID and host-based software RAID concurrently, you only need one RAID solution to maintain a high degree data availability.

If you use host-based software RAID, it is important to identify the impact of StorEdge A3500 LUN configuration and hardware maintenance procedures on the availability of LUNs. If you use StorEdge A3500's hardware RAID, you can perform most maintenance procedures in "Maintaining a StorEdge A3500 Disk Array" on page 127 without volume management disruptions. However, if you use host-based software RAID without hardware RAID, some maintenance procedures in "Maintaining a StorEdge A3500 Disk Array" on page 127 will require additional volume management administration because the availability of the LUNs being used by the volume manager may be impacted.

1. On one node, with all cluster nodes booted and attached to the StorEdge A3500, create a LUN.

Shortly after the LUN formatting completes, a logical name for the new LUN appears in `/dev/rdisk` on all cluster nodes attached to the StorEdge A3500 array.

For the procedure on creating a LUN, see *Sun StorEdge RAID Manager User's Guide*.

Note - The RAID Manager 6.x graphical user interface does not consistently display Solaris logical device names. Use the `format(1M)` command to verify Solaris logical device names.

2. After the new logical name you created in Step 1 on page 125 appears in `/dev/rdisk` on all cluster nodes, on one node, update the global device namespace.

```
# scgdevs
```

3. If you want a volume manager to manage the new LUN you created in Step 1 on page 125, run the appropriate Solstice DiskSuite and VERITAS Volume Manager commands to incorporate the new LUN into a diskset or disk group. For more information, see your Solstice DiskSuite or VERITAS Volume Manager documentation.

▼ How to Remove a LUN



Caution - This procedure removes all data on the LUN you remove.

1. Migrate all Oracle Parallel Server (OPS) tables, data services, and volumes off of the logical unit number (LUN) you want to remove.
2. If a volume manager does not manage the LUN you want to remove, proceed to Step 3 on page 126. Otherwise, run the appropriate Solstice DiskSuite or VERITAS Volume Manager commands to remove the LUN from any diskset or disk group.

For more information, see your Solstice DiskSuite or VERITAS Volume Manager documentation.

3. From one node, remove a LUN

For the procedure on removing a LUN, see *Sun StorEdge RAID Manager User's Guide*.

4. On all nodes, run the `scdidadm` command to remove all obsolete DIDs.

```
# scdidadm -C
```

▼ How to Reset StorEdge A3500 Configuration



Caution - This procedure creates a default logical unit number (LUN) that removes all LUN configuration. All LUN configuration and associated data will be lost on the LUNs you remove.

1. **Migrate all Oracle Parallel Server (OPS) tables, data services, and volumes off of the LUNs you want to remove.**
2. **If a volume manager does not manage any LUNs on the StorEdge A3500 you want to reset, proceed to Step 3 on page 127. Otherwise, run the appropriate Solstice DiskSuite or VERITAS Volume Manager commands to remove the LUN(s) from any diskset or disk group.**

For more information, see your Solstice DiskSuite or VERITAS Volume Manager documentation.

3. **On one node, reset the StorEdge A3500 configuration.**

For the procedure for resetting StorEdge A3500 configuration, see *Sun StorEdge RAID Manager User's Guide*.

Note - The RAID Manager 6.x graphical user interface does not consistently display Solaris logical device names. Use the `format` command to verify Solaris logical device names.

4. **Using the `format` command, specify the Solaris logical device name, then label the new LUN 0 disk.**
5. **On all nodes, run the `scdidadm` command to remove all obsolete DIDs.**

```
# scdidadm -C
```

6. **One at a time, boot each node.**

```
# boot -r
```

LUN 0 receives a new DID.

Maintaining a StorEdge A3500 Disk Array

This section provides the procedures for maintaining an StorEdge A3500 disk array. The following table lists these procedures.

TABLE 7-3 Task Map: Maintaining a StorEdge A3500 Disk Array

Task	For Instructions, Go To...
Add a disk drive Follow the same procedure used in a non-cluster environment.	<i>Sun StorEdge D1000 Storage Guide</i> for the hardware procedure <i>Sun StorEdge RAID Manager User's Guide</i> for the software procedure
Remove a disk drive Follow the same procedure used in a non-cluster environment. Before you physically remove a disk drive from an Redundant Disk Array Controller (rdac) disk group or disk group, you need to remove the logical unit number (LUN).	<i>Sun StorEdge D1000 Storage Guide</i> for the hardware procedure "How to Remove a LUN" on page 126
Replace a disk drive Follow the same procedure used in a non-cluster environment if you have only one failed disk drive. If you have more than one failed disk drive, you might need to back up and restore your data. If disk drive replacement affects any LUN's availability, remove the LUN(s) from volume management control using Solstice DiskSuite or VERITAS Volume Manager. After you replace the disk drive and you can access all LUNs, you can return the LUN(s) to volume management control.	<i>Sun StorEdge D1000 Storage Guide</i> for the hardware procedure. <i>Sun StorEdge RAID Manager User's Guide</i> for the software procedure.
Add an StorEdge A3500 to an existing cluster Install host adapters and SCSI cables, then power on the disk array. Use the hardware install procedure to perform an initial installation and configuration of an StorEdge A3500 disk array or to install an StorEdge A3500 disk array to an existing cluster.	"How to Add a StorEdge A3500" on page 130
Remove an StorEdge A3500 from an existing cluster Use the software install to install an StorEdge A3500 disk array to an existing cluster.	"How to Remove a StorEdge A3500" on page 136
Replace an StorEdge A3500 controller and restore an offline StorEdge A3500 controller This procedure requires that you set the <code>System_LunReDistribution</code> parameter in the <code>/etc/raid/rmparams</code> file to <code>false</code>. You must also verify that the new controller has the correct SCSI reservation state before assigning LUNs.	"How to Replace a Failed StorEdge A3500 Controller or Restore an Offline StorEdge A3500 Controller" on page 137

TABLE 7-3 Task Map: Maintaining a StorEdge A3500 Disk Array *(continued)*

Task	For Instructions, Go To...
Replace a host adapter This procedure requires that you halt the node that contains the failed SCSI host adapter and prepare the SCSI connections for continued operation.	“How to Replace a Host Adapter” on page 139
Replace a failed SCSI cable from the controller to the drive tray Follow the same procedure used in a non-cluster environment.	<i>Sun StorEdge A3500/A3500FC Controller Module Guide</i> <i>Sun StorEdge RAID Manager User's Guide</i>
Replace a StorEdge A3500-to-host SCSI cable Follow the same procedure used in a non-cluster environment.	<i>Sun StorEdge A3500/A3500FC Controller Module Guide</i> <i>Sun StorEdge RAID Manager User's Guide</i>
Replace a controller fan canister Follow the same procedure used in a non-cluster environment.	<i>Sun StorEdge A3500/A3500FC Controller Module Guide</i>
Replace the power supply fan canister Follow the same procedure used in a non-cluster environment.	<i>Sun StorEdge A3500/A3500FC Controller Module Guide</i>
Replace DC power or battery harness Shut down the cluster, then follow the non-cluster procedure for replacing the power supply housing.	<i>Sun Cluster 3.0 System Administration Guide</i> for procedures on shutting down a cluster <i>Sun StorEdge A3500/A3500FC Controller Module Guide</i> for replacement procedures
Replace the battery unit Shut down the cluster, then follow the non-cluster procedure for replacing the power supply housing.	<i>Sun Cluster 3.0 System Administration Guide</i> for procedures on shutting down a cluster <i>Sun StorEdge A3500/A3500FC Controller Module Guide</i> for replacement procedures

TABLE 7-3 Task Map: Maintaining a StorEdge A3500 Disk Array *(continued)*

Task	For Instructions, Go To...
Replace the controller card cage Shut down the cluster, then follow the non-cluster procedure for replacing the power supply housing.	<i>Sun Cluster 3.0 System Administration Guide</i> for procedures on shutting down a cluster <i>Sun StorEdge A3500/A3500FC Controller Module Guide</i> for replacement procedures
Replace the entire controller assembly Shut down the cluster, then follow the non-cluster procedure for replacing the power supply housing.	<i>Sun Cluster 3.0 System Administration Guide</i> for procedures on shutting down a cluster <i>Sun StorEdge A3500/A3500FC Controller Module Guide</i> for replacement procedures
Replace the power supply housing Shut down the cluster, then follow the non-cluster procedure for replacing the power supply housing.	<i>Sun Cluster 3.0 System Administration Guide</i> for procedures on shutting down a cluster <i>Sun StorEdge A3500/A3500FC Controller Module Guide</i>

▼ How to Add a StorEdge A3500

Use this procedure to add an StorEdge A3500 disk array to an existing cluster.

1. Ensure that each device in the SCSI chain has a unique SCSI address.

The default SCSI address for host adapters is 7. Reserve SCSI address 7 for one host adapter in the SCSI chain. This procedure refers to the host adapter you choose for SCSI address 7 as the host adapter on the second node. To avoid conflicts, in Step 7 on page 132 you will change the `scsi-initiator-id` of the remaining host adapter in the SCSI chain to an available SCSI address. This procedure refers to the host adapter with an available SCSI address as the host adapter on the first node. SCSI address 6 is usually available.

For more information, see the *OpenBoot 3.x Command Reference Manual* and the labels inside the storage device.

2. Shut down and power off the first node.

```
# scswitch -S -h nodename
# shutdown -y -g 0
```

For more information on shutdown, see *Sun Cluster 3.0 System Administration Guide*.

3. Install the host adapters in the node.

For the procedure on installing host adapters, see the documentation that shipped with your host adapters and nodes.

4. Connect the differential SCSI cable between the node and the disk array as shown in Figure 7-2.

Make sure that the *entire* SCSI bus length to each enclosure is less than 25 m. This measurement includes the cables to both nodes, as well as the bus length internal to each enclosure, node, and host adapter. Refer to the documentation that shipped with the disk array for other restrictions regarding SCSI operation.




Figure 7-2 Example of a StorEdge A3500 disk array

5. Power on the first node and the disk arrays.

6. Find the paths to the SCSI host adapters.

```
{0} ok show-disks
```

Identify and record the two controllers that will be connected to the disk arrays, and record these paths. Use this information to change the SCSI addresses of these controllers in the `nvrarc` script. Do not include the `/sd` directories in the device paths.

7. Edit the `nvrarc` script to change the `scsi-initiator-id` for the host adapters of the first node.

For a list of Editor `nvrarc` editor and keystroke commands, see Appendix B.

The following example sets the `scsi-initiator-id` to 6. The OpenBoot PROM Monitor prints the line numbers (0:, 1:, and so on).



Caution - Insert exactly one space after the double quote and before `scsi-initiator-id`.

```
{0} ok nvedit
0: probe-all
1: cd /sbus@1f,0/QLGC,isp@3,10000
2: 6 encode-int " scsi-initiator-id" property
3: device-end
4: cd /sbus@1f,0/
5: 6 encode-int " scsi-initiator-id" property
6: device-end
7: install-console
8: banner [Control C]
{0} ok
```

8. Store the changes.

The changes you make through the `nvedit` command are done on a temporary copy of the `nvrarc` script. You can continue to edit this copy without risk. After you have completed your edits, save the changes. If you are not sure about the changes, discard them.

- To store the changes, type:

```
{0} ok nvstore
{0} ok
```

- To discard the changes, type:

```
{0} ok nvquit  
{0} ok
```

9. Verify the contents of the nvramrc script you created in Step 7 on page 132.

If the contents of the nvramrc script are incorrect, use the nvedit command to make corrections.

```
{0} ok printenv nvramrc  
nvramrc =  
    probe-all  
    cd /sbus@1f,0/QLGC,isp@3,10000  
    6 encode-int " scsi-initiator-id" property  
    device-end  
    cd /sbus@1f,0/  
    6 encode-int " scsi-initiator-id" property  
    device-end  
    install-console  
    banner  
  
{0} ok
```

10. Instruct the OpenBoot PROM Monitor to use the nvramrc script:

```
{0} ok setenv use-nvramrc? true  
use-nvramrc? = true  
{0} ok
```

11. Boot the first node, and wait for it to join the cluster.

```
{0} ok boot -r
```

For more information, see *Sun Cluster 3.0 System Administration Guide*.

12. On all nodes, verify that the DIDs have been assigned to the StorEdge A3500 LUNs.

```
# scdidadm -l
```

13. Shut down the second node.

```
# scswitch -S -h nodename  
# shutdown -y -g 0
```

14. Power off the second node.

15. Install the host adapters in the second node.

For the procedure on installing host adapters, see the documentation that shipped with your nodes.

16. Connect the disk array to the host adapters using differential SCSI cables as shown in Figure 7-3.




Figure 7-3 Example of a StorEdge A3500 disk array

17. Without allowing the node to boot, power on the second node. If necessary, abort the system to continue with OpenBoot PROM Monitor tasks.

18. Verify that the second node sees the new host adapters and disk drives.

```
{0} ok show-disks
```

19. Verify that the scsi-initiator-id for the host adapters on the second node is set to 7.

Use the show-disks command to find the paths to the host adapters connected to these enclosures. Select each host adapter's device tree node, and display the node's properties to confirm that the scsi-initiator-id for each host adapter is set to 7.

```
{0} ok cd /sbus@1f,0/QLGC,isp@3,10000
{0} ok .properties
scsi-initiator-id      00000007
...
```

20. Boot the second node, and wait for it to join the cluster.

```
{0} ok boot -r
```

21. On all nodes, verify that the DIDs have been assigned to the StorEdge A3500 LUNs.

```
# scdidadm -l
```

22. Install the RAID Manager.

For the procedure on installing RAID Manager, see *Sun StorEdge RAID Manager Installation and Support Guide*.

23. Install StorEdge A3500 disk array patches.

For the location of patches and installation instructions, see *Sun Cluster 3.0 Release Notes*.

24. One at a time, reboot each node into cluster mode.

```
# boot
```

25. Upgrade the StorEdge A3500 disk array controller firmware.

For the StorEdge A3500 disk array controller firmware version number and boot level, see *Sun Cluster 3.0 Release Notes*. For the procedure on upgrading the StorEdge A3500 controller firmware, see *Sun StorEdge RAID Manager User's Guide*.

Where to Go From Here

To create a LUN from unassigned disk drives, see “How to Create a LUN” on page 125.

▼ How to Remove a StorEdge A3500

Use this procedure to remove a StorEdge A3500 from an existing cluster.



Caution - This procedure removes all data on the disk array you remove.

1. **Migrate all Oracle Parallel Server (OPS) tables, data services, and volumes off of the StorEdge A3500 disk array volumes.**
2. **Stop all I/O activity to the StorEdge A3500.**
3. **If a volume manager does not manage any of the logical unit numbers (LUN) on the StorEdge A3500, proceed to Step 4 on page 136. Otherwise, run the appropriate Solstice DiskSuite or VERITAS Volume Manager commands to remove the LUN(s) from any diskset or disk group.**

For more information, see your Solstice DiskSuite or VERITAS Volume Manager documentation.

4. **If you are removing the last StorEdge A3500 in your cluster, remove StorEdge A3500 packages.**

For the procedure on removing StorEdge A3500 packages, see the documentation that shipped with your disk array.

5. **One at a time, reboot each node into cluster mode.**

`# boot`

6. **Disconnect the SCSI cables from the disk array.**
7. **On all cluster nodes, remove references to the disk array**

```
# devfsadm -C
# scdidadm -C
```

8. If needed, remove any host adapters from the nodes.

For the procedure on removing host adapters, see the documentation that shipped with your nodes.

▼ How to Replace a Failed StorEdge A3500 Controller or Restore an Offline StorEdge A3500 Controller

Use this procedure to remove a failed StorEdge A3500 disk array controller or restore an offline StorEdge A3500 disk array controller.

For conceptual information on SCSI reservations and failure fencing, see *Sun Cluster 3.0 Concepts*.

- 1. On all nodes, to prevent LUNs from being automatically assigned to the new controller, set the `System_LunReDistribution` parameter in the `/etc/raid/rmparams` file to `false`.**



Caution - You must set the `System_LunReDistribution` parameter in the `/etc/raid/rmparams` file to `false` so that no LUNs are assigned to the controller being brought online. After you verify in Step 4 on page 138 that the controller has the correct SCSI reservation state, you can balance LUNs between both controllers.

For the procedure on modifying the `rmparams` file, see *Sun StorEdge RAID Manager Installation and Support Guide*.

- 2. If you have a failed controller, replace the failed controller with a new controller, but do not bring the controller online. If you do not have a failed controller, proceed to Step 3 on page 137.**

For the procedure on replacing StorEdge A3500 controllers, see *Sun StorEdge RAID Manager Installation and Support Guide*.

- 3. On one node, use the RAID Manager 6.x graphical user interface's Recovery application to bring the controller online.**



Caution - You must use the RAID Manager 6.x graphical user interface's Recovery application to bring the controller online. Do not use the Redundant Disk Array Controller Utility (rdacutil) because it ignores the value of the System_LunReDistribution parameter in the /etc/raid/rmparams file.

For information on the Recovery application, see *Sun StorEdge RAID Manager User's Guide*. For the procedure on replacing StorEdge A3500 controllers, see *Sun StorEdge RAID Manager User's Guide*. If you have problems bringing the controller online, see *Sun StorEdge RAID Manager Installation and Support Guide*.

4. On one node connected to the disk array, verify that the controller has the correct SCSI reservation state.

Run the `sddidadm` repair procedure (-R) on LUN 0 of the controller you want to restore.

```
# sddidadm -R /dev/dsk/cNtXdY
```

5. Set the new controller to active/active mode and assign LUNs to the new controller.

For more information on controller modes, see *Sun StorEdge RAID Manager Installation and Support Guide* and *Sun StorEdge RAID Manager User's Guide*.

6. Reset the System_LunReDistribution parameter in the /etc/raid/rmparams file to true.

For the procedure on changing the `rmparams` file, see *Sun StorEdge RAID Manager Installation and Support Guide*.

Replacing a Host Adapter

This section provides the procedure for replacing a failed host adapter in an existing cluster. This procedure requires halting the node with the host adapter that needs replacement and preparing the SCSI connections for continued operation. This procedure requires one or two differential SCSI terminators.

▼ How to Replace a Host Adapter

In this procedure, node 1's host adapter on SCSI bus A needs replacement. node 2 remains in service.

Several steps in this procedure require that you stop I/O activity. To stop I/O activity, take the controller offline using the RAID Manager 6.x graphical user interface's manual recovery procedure in *Sun StorEdge RAID Manager User's Guide*.

1. Without powering off the node, shut down node 1.

For the procedure on shutting down a node, see *Sun Cluster 3.0 System Administration Guide*.

2. From node 2, stop I/O activity to SCSI bus A.

3. From the controller end of the SCSI cable, disconnect the SCSI bus A cable that connects the StorEdge A3500 to node 1, then replace this cable with a differential SCSI terminator.

4. Restart I/O activity on SCSI bus A.

5. Determine if you need to stop I/O activity to SCSI bus B.

If servicing of the failed host adapter will affect SCSI bus B, from node 1, stop I/O activity to the StorEdge A3500 controller on SCSI bus B. Otherwise, skip to Step 8 on page 139.

6. From the controller end of the SCSI cable, disconnect the SCSI bus B cable that connects the StorEdge A3500 to node 1, then replace this cable with a differential SCSI terminator.

7. Restart I/O activity on SCSI bus B.

8. Power off node 1.

9. Replace node 1's host adapter.

10. Without booting the node, power on node 1.

11. From node 2, stop I/O activity to the StorEdge A3500 controller on SCSI bus A.

12. Remove the differential SCSI terminator from SCSI bus A, and reinstall the SCSI cable to connect the StorEdge A3500 controller to node 1.

13. Restart I/O activity on SCSI bus A.

14. Stop I/O activity on SCSI bus B.

15. Determine if you need to remove the differential SCSI terminator from SCSI bus B, and reinstall the SCSI cable to connect node 1 and StorEdge A3500 controller.

If, in Step 6 on page 139, you disconnected the SCSI bus B cable that connects the StorEdge A3500 to node 1 and replaced the cable with a differential SCSI terminator, replace the differential SCSI terminator with the SCSI cable. Otherwise, skip to Step 18 on page 140.

16. Restart I/O activity on SCSI bus B.

17. Rebalance all logical unit numbers (LUN).

18. Boot node 1 into cluster mode.

```
# boot
```

Verifying Sun Cluster Hardware Redundancy

This appendix provides the tests for verifying and validating the high availability (HA) of your Sun Cluster configuration. The tests in this appendix assume that you installed Sun Cluster hardware, the Solaris operating environment, and Sun Cluster software. All nodes should be booted as cluster members.

This appendix contains the following procedures:

- “How to Test Nodes Using a Power-off Method” on page 142
- “How to Test Cluster Interconnects” on page 143
- “How to Test Network Adapter Failover Groups” on page 144

If your cluster passes these tests, your hardware has adequate redundancy: Nodes, cluster transport cables, and Network Adapter Failover (NAFO) groups are not single points of failure.

To perform the tests in “How to Test Nodes Using a Power-off Method” on page 142 and “How to Test Cluster Interconnects” on page 143, you must first identify the device groups that each node masters. Perform these tests on all cluster pairs that share a disk device group. Each pair will have a primary and a secondary for a particular device group. Use the `scstat(1M)` command to determine the initial primary and secondary.

For conceptual information on primary, secondary, failover, device groups, or cluster hardware, see *Sun Cluster 3.0 Concepts*.

Testing Node Redundancy

This section provides the procedure for testing node redundancy and high availability of device groups. Perform the following procedure to confirm that the secondary takes over the device group mastered by the primary when the primary fails.

▼ How to Test Nodes Using a Power-off Method

1. Power off the primary node.

Cluster interconnect error messages appear on the console of the existing nodes.

2. On another node, run the `scstat` command to verify that the secondary took ownership of the device group mastered by the primary.

Look for the output that shows the device group ownership.

```
# scstat
```

3. Power on the initial primary and boot into cluster mode.

Wait for the system to boot. The system automatically starts the membership monitor software. The node then rejoins the configuration.

If you have the device group failback option enabled, skip Step 4 on page 142 because the system boot process moves ownership of the device group back to the initial primary. Otherwise, proceed to Step 4 on page 142 to move ownership of the device group back to the initial primary. Use the `SCCONF -p` command to determine if your device group has the device group failback option enabled.

4. If you do not have the device group failback option enabled, from the initial primary, run the `scswitch(1M)` command to move ownership of the device group back to the initial primary.

```
# scswitch -S -h nodename
```

5. Verify that the initial primary has ownership of the device group.

Look for the output that shows the device group ownership.

```
# scstat
```

Testing Cluster Interconnect and Network Adapter Failover Group Redundancy

This section provides the procedure for testing cluster interconnect and Network Adapter Failover (NAFO) group redundancy.

▼ How to Test Cluster Interconnects

1. **Disconnect one of the cluster transport cables from a node that masters a device group.**

Messages appear on the consoles of each node, and error messages appear in the `/var/adm/messages` file. If you run the `scstat(1M)` command, the Sun Cluster software assigns the cluster transport path you disconnected a faulted status. This fault does not result in a failover.

2. **Disconnect the remaining cluster transport cable from the primary node you identified in Step 1 on page 143.**

Messages appear on the consoles of each node, and error messages appear in the `/var/adm/messages` file. If you run the `scstat` command, the Sun Cluster software assigns the cluster transport path you disconnected a faulted status. This action causes the primary node to go down, resulting in a partitioned cluster. For conceptual information on failure fencing or *split brain*, see *Sun Cluster 3.0 Concepts*.

3. **On another node, run the `scstat` command to verify that the secondary node took ownership of the device group mastered by the primary.**

```
# scstat
```

4. **Reconnect all cluster transport cables.**

5. **Boot the initial primary, which you identified in Step 1 on page 143, into cluster mode.**

```
{0} ok boot
```

If you have the device group failback option enabled, skip Step 7 on page 144 because the system boot process moves ownership of the device group back to the initial primary. Otherwise, proceed to Step 7 on page 144 to move ownership

of the device group back to the initial primary. Use the `scconf -p` command to determine if your device group has the device group failback option enabled.

6. **Verify that the Sun Cluster software assigned each cluster transport path you reconnected in Step 4 on page 143 a path online status.**

```
# scstat
```

7. **If you do not have the device group failback option enabled, move ownership of the device group back to the initial primary.**

```
# scswitch -S -h nodename
```

▼ How to Test Network Adapter Failover Groups

Perform this procedure on each node.

1. **Identify the current active network adapter.**

```
# pnmstat -l
```

2. **Disconnect one public network cable from the current active network adapter.**

Error messages appear in the node's console. This action causes a NAFO failover to a backup network adapter.

3. **From the master console, verify that the Sun Cluster software failed over to the backup NAFO adapter.**

A NAFO failover occurred if the backup NAFO adapter displays an active status.

```
# pnmstat -l
```

4. **Reconnect the public network cable, and wait for the initial network adapter to come online.**
5. **Switch over all IP addresses hosted by the active network adapter to the initial network adapter, and make the initial network adapter the active network adapter.**

```
# pnmset switch adapter
```


NVRAMRC Editor and NVEDIT Keystroke Commands

This section provides useful `nvrsrc` editor and `nvedit` keystroke commands.

An `nvrsrc` script contains a series of OBP commands that are executed during the boot sequence. The procedures in this guide assume that this script is empty. If your `nvrsrc` script contains data, add the entries to the end of the script. To edit an `nvrsrc` script or merge new lines in an `nvrsrc` script, use `nvrsrc` editor and `nvedit` keystroke commands.

Table B-1 and Table B-2 list useful `nvrsrc` editor and `nvedit` keystroke commands. For an entire list of `nvrsrc` editor and `nvedit` keystroke commands, see *OpenBoot 3.x Command Reference Manual*.

TABLE B-1 NVRAMRC Editor Commands

Command	Description
<code>nvedit</code>	Enter the <code>nvrsrc</code> editor. If the data remains in the temporary buffer from a previous <code>nvedit</code> session, resume editing previous contents. Otherwise, read the contents of <code>nvrsrc</code> into the temporary buffer and begin editing it. This command works on a buffer, and you can save the contents of this buffer by using the <code>nvstore</code> command.
<code>nvstore</code>	Copy the contents of the temporary buffer to <code>nvrsrc</code> and discard the contents of the temporary buffer.
<code>nvquit</code>	Discard the contents of the temporary buffer, without writing it to <code>nvrsrc</code> . Prompt for confirmation.

TABLE B-2 NVEDIT Keystroke Commands

Keystroke	Description
^A	Move to the beginning of the line.
^B	Move backward one character.
^C	Exit the script editor.
^F	Move forward one character.
^K	Delete until end of line.
^L	List all lines.
^N	Move to the next line of the <code>nvrnmrc</code> editing buffer.
^O	Insert a new line at the cursor position and stay on the current line.
^P	Move to the previous line of the <code>nvrnmrc</code> editing buffer.
^R	Replace the current line.
Delete	Delete previous character.
Return	Insert a new line at the cursor position and advance to the next line.