



Sun Cluster 3.0 U1 Hardware Guide

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

Part No. 806-7070-10
July 2001, Revision A

Copyright 2001 Sun Microsystems, Inc., 901 San Antonio Road, Palo Alto, CA 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, docs.sun.com, Solaris, Solstice DiskSuite, Sun StorEdge, SunSolve, SunSwift, 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 2001 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, docs.sun.com, Solaris, Solstice DiskSuite, Sun StorEdge, SunSolve, SunSwift, 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.



Adobe PostScript

Contents

Preface ix

- 1. Introduction to Sun Cluster Hardware** 1
 - Overview of Sun Cluster Hardware 2
 - Installing Sun Cluster Hardware 3
 - Maintaining Sun Cluster Hardware 5
 - Powering On and Off Sun Cluster Hardware 6
 - Local and Multihost Disks in a Sun Cluster 6
 - Removable Media in a Sun Cluster 7

- 2. Installing and Configuring the Terminal Concentrator** 9
 - Installing the Terminal Concentrator 10
 - ▼ How to Install the Terminal Concentrator in a Cabinet 10
 - ▼ How to Cable the Terminal Concentrator 15
 - Configuring the Terminal Concentrator 16
 - ▼ How to Configure the Terminal Concentrator 16
 - ▼ How to Set Terminal Concentrator Port Parameters 19
 - ▼ How to Correct a Port Configuration Access Error 21
 - ▼ How to Establish a Default Route for the Terminal Concentrator 23
 - Using the Terminal Concentrator 26

- ▼ How to Connect to a Node's Console Through the Terminal Concentrator 26
- ▼ How to Reset a Terminal Concentrator Port 28
- 3. Installing and Maintaining Cluster Interconnect and Public Network Hardware 31**
 - Installing Cluster Interconnect and Public Network Hardware 32
 - ▼ How to Install Cluster Transport Cables 32
 - Maintaining Cluster Interconnect and Public Network Hardware 35
 - Adding Cluster Transport Adapters 36
 - Replacing Cluster Transport Adapters 37
 - Removing Cluster Transport Adapters 39
 - Adding Public Network Adapters 41
 - Replacing Public Network Adapters 41
 - Removing Public Network Adapters 42
 - Sun Gigabit Ethernet Adapter Considerations 43
- 4. Installing and Maintaining a Sun StorEdge MultiPack Enclosure 45**
 - Installing a StorEdge MultiPack Enclosure 46
 - ▼ How to Install a StorEdge MultiPack Enclosure 46
 - Maintaining a StorEdge MultiPack 51
 - ▼ How to Add Disk Drive to StorEdge Multipack Enclosure in a Running Cluster 52
 - ▼ How to Replace a Disk Drive in StorEdge MultiPack Enclosure in a Running Cluster 55
 - ▼ How to Remove a Disk Drive From a StorEdge MultiPack Enclosure in Running Cluster 59
 - ▼ How to Add a StorEdge MultiPack Enclosure to a Running Cluster 60
 - ▼ How to Replace a StorEdge MultiPack Enclosure in a Running Cluster 67
 - ▼ How to Remove a StorEdge MultiPack Enclosure From a Running Cluster 69

5. Installing and Maintaining a Sun StorEdge D1000 Disk Array 71

Installing a StorEdge D1000 Disk Array 72

- ▼ How to Install a StorEdge D1000 Disk Array 72

Maintaining a StorEdge D1000 Disk Array 77

- ▼ How to Add a Disk Drive in a StorEdge D1000 Disk Array in a Running Cluster 78
- ▼ How to Replace a Disk Drive in a StorEdge D1000 Disk Array in a Running Cluster 81
- ▼ How to Remove a Disk Drive From a StorEdge D1000 Disk Array in a Running Cluster 85
- ▼ How to Add a StorEdge D1000 Disk Array to a Running Cluster 87
- ▼ How to Replace a StorEdge D1000 Disk Array in a Running Cluster 94
- ▼ How to Remove a StorEdge D1000 Disk Array From a Running Cluster 96

6. Installing and Maintaining a Sun StorEdge A5x00 Disk Array 99

Installing a StorEdge A5x00 Disk Array 100

- ▼ How to Install a StorEdge A5x00 Disk Array 100

Maintaining a StorEdge A5x00 Disk Array 102

- ▼ How to Add a Disk Drive to a StorEdge A5x00 Disk Array in a Running Cluster 103
- ▼ How to Replace a Disk Drive in a StorEdge A5x00 Disk Array in a Running Cluster 105
- ▼ How to Remove a Disk Drive From a StorEdge A5x00 Disk Array in a Running Cluster 110
- ▼ How to Add the First StorEdge A5x00 Disk Array to a Running Cluster 112
- ▼ How to Add a StorEdge A5x00 Disk Array to a Running Cluster That Has Existing StorEdge A5x00 Disk Arrays 115
- ▼ How to Replace a StorEdge A5x00 Disk Array in a Running Cluster 117
- ▼ How to Remove a StorEdge A5x00 Disk Array From a Running Cluster 120

7. Installing and Maintaining a Sun StorEdge A3500/A3500FC System	123
Installing a Sun StorEdge A3500/A3500FC System	124
▼ How to Install a StorEdge A3500/A3500FC System	124
Configuring a Sun StorEdge A3500/A3500FC System	132
▼ How to Create a LUN	133
▼ How to Delete a LUN	136
▼ How to Reset StorEdge A3500/A3500FC LUN Configuration	139
▼ How to Correct Mismatched DID Numbers	142
Maintaining a StorEdge A3500/A3500FC System	144
▼ How to Add a StorEdge A3500/A3500FC System to a Running Cluster	148
▼ How to Remove a StorEdge A3500/A3500FC System From a Running Cluster	158
▼ How to Replace a Failed Controller or Restore an Offline Controller	161
How to Upgrade Controller Module Firmware in a Running Cluster	163
▼ How to Add a Disk Drive in a Running Cluster	165
▼ How to Replace a Failed Disk Drive in a Running Cluster	166
▼ How to Remove a Disk Drive From a Running Cluster	167
How to Upgrade Disk Drive Firmware in a Running Cluster	167
▼ How to Replace a Host Adapter in a Node (Connected to a StorEdge A3500 System)	168
▼ How to Replace a Host Adapter in a Node (Connected to a StorEdge A3500FC System)	170
8. Installing and Maintaining a Sun StorEdge T3 Disk Tray Single-Controller Configuration	173
Installing a StorEdge T3 Disk Tray	174
▼ How to Install a StorEdge T3 Disk Tray	174
Configuring a StorEdge T3 Disk Tray	177
▼ How to Create a Sun StorEdge T3 Disk Tray Logical Volume	177

- ▼ How to Remove a Sun StorEdge T3 Disk Tray Logical Volume 179
- Maintaining a StorEdge T3 Disk Tray 182
 - ▼ How to Upgrade StorEdge T3 Disk Tray Firmware 184
 - ▼ How to Replace a Disk Drive 185
 - ▼ How to Add a StorEdge T3 Disk Tray 186
 - ▼ How to Remove a StorEdge T3 Disk Tray 195
 - ▼ How to Replace a Host-to-Hub Component 198
 - ▼ How to Replace a Sun StorEdge FC-100 Hub and Hub-to-Disk Tray Component 199
 - ▼ How to Replace a StorEdge T3 Disk Tray Controller 200
 - ▼ How to Replace a StorEdge T3 Disk Tray Chassis 201
 - ▼ How to Replace a Host Adapter 202
- A. Verifying Sun Cluster Hardware Redundancy 205**
 - Testing Node Redundancy 206
 - ▼ How to Test Nodes Using a Power-off Method 206
 - Testing Cluster Interconnect and Network Adapter Failover Group Redundancy 207
 - ▼ How to Test Cluster Interconnects 207
 - ▼ How to Test Network Adapter Failover Groups 209
- B. NVRAMRC Editor and NVEDIT Keystroke Commands 211**
- C. Recabling Disk Devices 213**
 - Moving a Disk Cable 214
 - ▼ How to Move a Disk Cable to a New Host Adapter 214
 - ▼ How to Move a Disk Cable From One Node to Another 216
 - ▼ How to Update Sun Cluster Software to Reflect Proper Device Configuration 218

Preface

Sun Cluster 3.0 U1 Hardware Guide contains the procedures for installing and maintaining 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. Some procedures in this document are for trained service providers only, as noted.

Using UNIX Commands

This document might 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:

- 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 U1 Concepts</i>	806-1424
Software installation	<i>Sun Cluster 3.0 U1 Installation Guide</i>	806-1419
Data services	<i>Sun Cluster 3.0 U1 Data Services Installation and Configuration Guide</i>	805-1421
API development	<i>Sun Cluster 3.0 U1 Data Services Developer's Guide</i>	805-1422
System administration	<i>Sun Cluster 3.0 U1 System Administration Guide</i>	806-1423
Sun Cluster release notes	<i>Sun Cluster 3.0 U1 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
Sun StorEdge A5x00 product notes	<i>Sun StorEdge A5000 Product Notes</i>	805-1018
Sun StorEdge A5x00 installation	<i>Sun StorEdge A5000 Installation and Documentation Guide</i>	805-1903
Sun StorEdge A5x00 installation and service	<i>Sun StorEdge A5000 Installation and Service Manual</i>	802-7573
Sun StorEdge A5x00 hardware configuration	<i>Sun StorEdge A5000 Configuration Guide</i>	805-0264
Sun StorEdge RAID Manager installation	<i>Sun StorEdge RAID Manager Installation and Support Guide</i>	805-7756

Application	Title	Part Number
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 A3500/A3500FC hardware configuration	<i>Sun StorEdge A3500/A3500FC Hardware Configuration Guide</i>	805-4981
Sun StorEdge A3500 controller module configuration	<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
FC Hub installation and service	<i>FC-100 Hub Installation and Service Manual</i>	806-0315
Sun StorEdge T3 hardware installation and setup	<i>Sun StorEdge T3 Installation, Operation, and Service Manual</i>	806-1062
Sun StorEdge T3 hardware configuration.	<i>Sun StorEdge T3 Configuration Guide</i>	806-4210
Sun StorEdge T3 hardware administration	<i>Sun StorEdge T3 Disk Tray Administrator's Guide</i>	806-1063
Sun Gigabit Ethernet adapter installation and usage	<i>Sun Gigabit Ethernet/P 2.0 Adapter Installation and User's Guide</i>	805-2785

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.com`SM 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. The chapter also provides overviews of the tasks that are involved in installing and maintaining this hardware *specifically in a Sun Cluster environment*.

This chapter contains the following information:

- [“Overview of Sun Cluster Hardware” on page 2](#)
- [“Installing Sun Cluster Hardware” on page 3](#)
- [“Maintaining Sun Cluster Hardware” on page 5](#)
- [“Powering On and Off Sun Cluster Hardware” on page 6](#)
- [“Local and Multihost Disks in a Sun Cluster” on page 6](#)
- [“Removable Media in a Sun Cluster” on page 7](#)

Overview of Sun Cluster Hardware

The procedures in this document discuss the installation, configuration, and maintenance of cluster hardware. [FIGURE 1-1](#) shows an overview of cluster hardware components.

For conceptual information on these hardware components, see the *Sun Cluster 3.0 U1 Concepts* document.

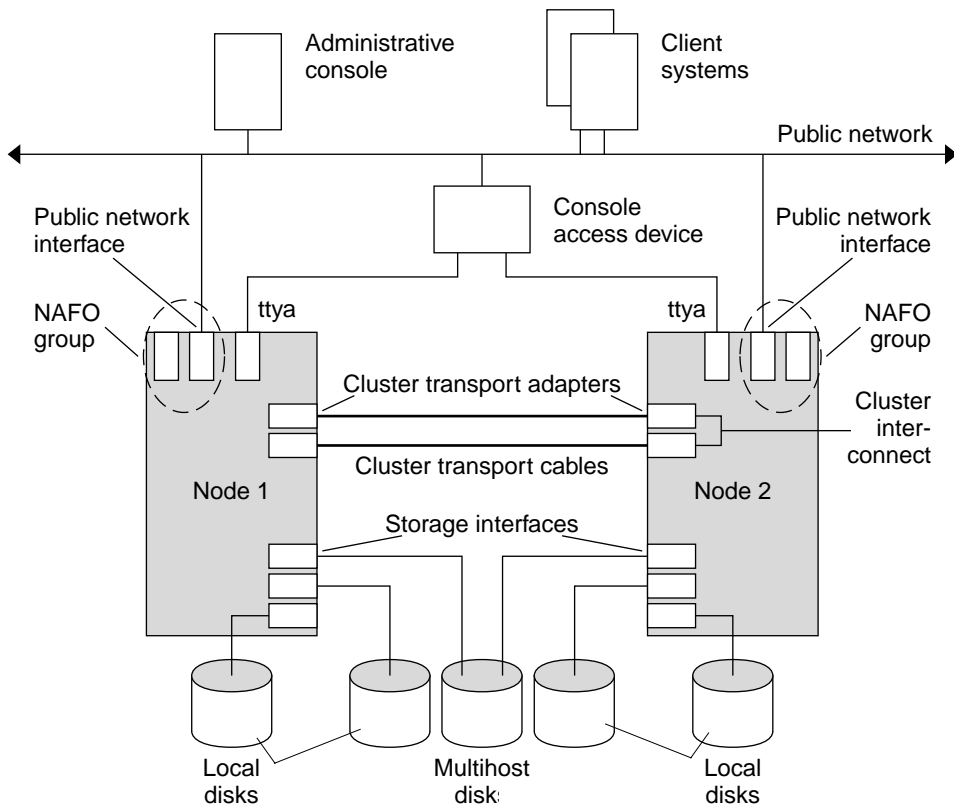


FIGURE 1-1 Cluster Hardware in a Sample Two-Node Cluster Configuration

Installing Sun Cluster Hardware

[TABLE 1-1](#) lists the tasks for installing a cluster and the sources for instructions. Perform these tasks in the order they are listed.

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
Install the administrative console.	The documentation that shipped with your administrative console
Install a console access device.	“Installing the Terminal Concentrator” on page 10
Use the procedure that is indicated for your type of console access device. For example, Sun Enterprise E10000 servers use a System Service Processor (SSP) as a console access device, rather than a terminal concentrator.	or The documentation that shipped with your Sun Enterprise E10000 hardware
Install the cluster interconnect and public network hardware.	“Installing and Maintaining Cluster Interconnect and Public Network Hardware” on page 31

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

Task	For Instructions, Go To
Install and configure the storage devices. Use the procedure that is indicated for your type of storage hardware.	<p data-bbox="822 239 1182 319">“Installing and Maintaining a Sun StorEdge MultiPack Enclosure” on page 45</p> <p data-bbox="822 354 1210 434">“Installing and Maintaining a Sun StorEdge D1000 Disk Array” on page 71</p> <p data-bbox="822 468 1210 548">“Installing and Maintaining a Sun StorEdge A5x00 Disk Array” on page 99</p> <p data-bbox="822 583 1222 663">“Installing and Maintaining a Sun StorEdge A3500/A3500FC System” on page 123</p> <p data-bbox="822 697 1222 770">“Installing and Maintaining a Sun StorEdge T3 Disk Tray Single-Controller Configuration” on page 173</p>
Install the Solaris operating environment and Sun Cluster software.	<p data-bbox="822 788 1188 814"><i>Sun Cluster 3.0 U1 Installation Guide</i></p>
Configure the cluster interconnects.	<p data-bbox="822 866 1085 918"><i>Sun Cluster 3.0 U1 System Administration Guide</i></p>

Maintaining Sun Cluster Hardware

This guide augments documentation that ships with your hardware components by providing information on maintaining the hardware *specifically in a Sun Cluster environment*. TABLE 1-2 describes some of the differences between maintaining cluster hardware and maintaining standalone hardware.

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

Task	Standalone Hardware	Cluster Hardware
Shutting down a node	Use the <code>shutdown(1M)</code> command.	To perform an orderly node shutdown, first use the <code>scswitch(1M)</code> command to switch device groups and resource groups to another node. Then shut down the node by running the <code>shutdown(1M)</code> command.
Adding a disk	Run <code>boot -r</code> or <code>devfsadm(1M)</code> 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 <code>devfsadm(1M)</code> , <code>scgdevs(1M)</code> , and <code>scdidadm(1M)</code> 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 adapter	Perform an orderly node shutdown, then install the public network adapter. After you install the network adapter, update the <code>/etc/hostname.adapter</code> and <code>/etc/inet/hosts</code> files.	Perform an orderly node shutdown, then install the public network adapter. After you install the public network adapter, update the <code>/etc/hostname.adapter</code> and <code>/etc/inet/hosts</code> files. Finally, add this public network adapter to a NAFO group.

Powering On and Off Sun Cluster Hardware

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

- Use power-on and power-off procedures in *Sun Cluster 3.0 U1 System Administration Guide* for nodes in a running cluster.
- Use the power-on and power-off procedures in the manuals that shipped with the hardware *only* for systems that are newly installed or are in the process of being installed.



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 in a Sun Cluster

Two sets of storage devices reside 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 and local disks, see the *Sun Cluster 3.0 U1 Concepts* document.

Removable Media in a Sun Cluster

Removable media include tape and CD-ROM drives, which are local devices. This guide does not contain procedures for adding, removing, or replacing removable media as highly available storage devices. Although tape and CD-ROM drives are global devices, these drives do not have more than one port and do not have multi-initiator firmware support that would enable these devices as 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, remove, replace, and use tape and CD-ROM drives as you would in a non-cluster environment. For procedures on installing, 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 as a console access device in a Sun Cluster environment. This chapter also includes information on how to use a terminal concentrator.

This chapter contains the following procedures:

- [“How to Install the Terminal Concentrator in a Cabinet” on page 10](#)
- [“How to Cable the Terminal Concentrator” on page 15](#)
- [“How to Configure the Terminal Concentrator” on page 16](#)
- [“How to Set Terminal Concentrator Port Parameters” on page 19](#)
- [“How to Correct a Port Configuration Access Error” on page 21](#)
- [“How to Establish a Default Route for the Terminal Concentrator” on page 23](#)
- [“How to Connect to a Node’s Console Through the Terminal Concentrator” on page 26](#)
- [“How to Reset a Terminal Concentrator Port” on page 28](#)

For conceptual information on terminal concentrators, see the *Sun Cluster 3.0 U1 Concepts* document.

Installing the Terminal Concentrator

This section describes the procedure for installing the terminal concentrator hardware and for connecting cables from the terminal concentrator to the administrative console and to 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 in a cabinet. For convenience, you can rack-mount the terminal concentrator even if your cluster does not contain rack-mounted nodes.

- To rack-mount your terminal concentrator, go to the first step of the following procedure.
- If you do not want to rack-mount your terminal concentrator, place the terminal concentrator in its standalone location, connect the unit power cord into a utility outlet, and go to [“How to Cable the Terminal Concentrator”](#) on page 15.

1. Install the terminal concentrator bracket hinge onto the primary cabinet:

a. **Locate the bracket hinge portion of the terminal concentrator bracket assembly** (see [FIGURE 2-1](#)).

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

Thread the screws into holes 8 and 29, as shown in [FIGURE 2-1](#). The locator screws 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 into holes 7 and 28.**

Tighten these screws, and the screws in holes 8 and 29, as shown in [FIGURE 2-1](#).




FIGURE 2-1 Installing the Terminal Concentrator Bracket Hinge to the Cabinet

2. Install the terminal concentrator into 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, aligning the holes in the side pieces with the threaded studs on the bottom plate.
- c. Install and tighten three nuts on the three threaded studs that penetrate through each side plate.




FIGURE 2-2 Installing the Terminal Concentrator Into the Bracket

3. **Install the terminal concentrator bracket onto the bracket hinge that is already installed on the cabinet:**
 - a. Turn the terminal concentrator bracket on its side so the hinge holes and cable connectors face toward the bracket hinge (see [FIGURE 2-3](#)).
 - b. Align the bracket holes with the boss pins on the bracket hinge and install the bracket onto the hinge.
 - c. 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 Bracket Installed on the Hinge

4. Connect one end of the power cord to the terminal concentrator, as shown in [FIGURE 2-4](#). Connect the other end of the power cord to the power distribution unit.




FIGURE 2-4 Terminal Concentrator Cable Connector Locations

Where to Go From Here

To cable the terminal concentrator, go to [“How to Cable the Terminal Concentrator”](#) on page 15.

▼ How to Cable the Terminal Concentrator

1. **Connect a DB-25 to RJ-45 serial cable (part number 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](#).**

This cable connection from the administrative console enables you to configure the terminal concentrator. You can remove this connection after you set up the terminal concentrator.




FIGURE 2-5 Connecting the Administrative Console

2. **Connect the cluster nodes to the terminal concentrator by using DB-25 to RJ-45 serial cables.**

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 windows from the Cluster Control Panel (CCP). For more information on using the CCP, see the *Sun Cluster 3.0 U1 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 rear rail of the cabinet (see [FIGURE 2-3](#)).**

Where to Go From Here

Go to [“Configuring the Terminal Concentrator”](#) on page 16.

Configuring the Terminal Concentrator

This section describes 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
```

4. Hold down the terminal concentrator Test button (FIGURE 2-6) until the power LED flashes (about three seconds), then release the Test button.
5. Hold down the terminal concentrator Test button again for one second, then release it.

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 and LEDs

6. Observe the terminal concentrator front-panel LEDs:

- If the front-panel LEDs light up as shown in [TABLE 2-1](#), and the administrative console displays a `monitor:: prompt`, go to [Step 7](#).
- If the front-panel LEDs do not light up as shown in [TABLE 2-1](#), or the administrative console does not display a `monitor:: prompt`, use [TABLE 2-2](#) and the documentation that shipped with your terminal concentrator to troubleshoot the problem.

TABLE 2-1 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 2-2 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					

7. 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
```

8. After you finish the `addr` session, power-cycle the terminal concentrator.

The Load and Active LEDs should briefly blink, then the Load LED should turn off.

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

10. Exit the `tip` utility by pressing 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 19.

▼ 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 this 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.

1. Locate the Sun serial number label on the top panel of the terminal concentrator (FIGURE 2-7).
2. Check if the serial number 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. Go to [Step 4](#).
 - If the numbers after the dash start with 9519 or lower, you must change the port type variable. Go to [Step 3](#).

Sun label:

- 9520 or higher, the variable is correct
- 9519 or lower, the variable must be reset



Sun serial number label




FIGURE 2-7 Determining the Version From the Serial Number Label

3. Use the administrative console to change the port type variable to `dial_in` 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 is 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
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 the boot process before you proceed.

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 26.

▼ How to Correct a Port Configuration Access Error

A misconfigured port that does not accept network connections might 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 Return again after you make the connection, then 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 28.
- If you see an `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](#).

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 4.

```
admin-ws# telnet tcl
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 4
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 4
```

▼ How to Establish a Default Route for the Terminal Concentrator

Note – This procedure is optional. By setting a default route, you prevent possible problems with routing table overflows (see the following paragraphs). Routing table overflow is not a problem for connections that are made from a host that resides on the same network as 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.

1. Connect to the terminal concentrator.

```
# telnet tc_name
```

tc_name Specifies the name of the terminal concentrator

2. Press Return again after you make the connection, then 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
```

Note – 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, you cannot access the node consoles.

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 tcl
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 describes the procedures for using the terminal concentrator in a cluster.

TABLE 2-3 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 26
Reset a terminal concentrator port	"How to Reset a Terminal Concentrator Port" on page 28

▼ 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 terminal concentrator port that the node is 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 that is 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 telnet connection, the system prompts you for the login name and password.

3. Set the terminal type, based on the type of window that was used in [Step 1](#).

```
# 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 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)
pys-palindrome-1# TERM=xterm; export TERM
(for csh)
pys-palindrome-1# set term=xterm
```

▼ How to Reset a Terminal Concentrator Port

When a port on the terminal concentrator is busy, you can reset the port to disconnect its 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 21](#) 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 Return again after you make 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 tcl
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 describes the procedures for installing and maintaining cluster interconnect and public network hardware. Cluster interconnect components include cluster transport adapters, cluster transport junctions, and cluster transport cables. Use [FIGURE 1-1](#) to identify cluster interconnect and public network hardware.

This chapter contains the following procedures and information for maintaining cluster interconnect and public network hardware:

- [“How to Install Cluster Transport Cables” on page 32](#)
- [“How to Add Cluster Transport Adapters” on page 36](#)
- [“How to Replace Cluster Transport Adapters” on page 37](#)
- [“How to Remove Cluster Transport Adapters” on page 39](#)
- [“How to Add Public Network Adapters” on page 41](#)
- [“How to Replace Public Network Adapters” on page 41](#)
- [“How to Remove Public Network Adapters” on page 42](#)
- [“Sun Gigabit Ethernet Adapter Considerations” on page 43](#)

For conceptual information on cluster interconnects and public network interfaces, see the *Sun Cluster 3.0 U1 Concepts* document.

Installing Cluster Interconnect and Public Network Hardware

TABLE 3-1 lists procedures for installing cluster interconnect and public network hardware. Perform the procedures in the order that they are listed.

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

Task	For Instructions, Go To
Install the cluster transport adapters.	The documentation that shipped with your node and transport adapters
Install the cluster transport cables and (optionally) transport junctions.	“How to Install Cluster Transport Cables” on page 32
Install the public network adapters.	The documentation that shipped with your node and public network adapters

▼ How to Install Cluster Transport Cables

1. Locate the cluster transport adapters.

When you cable the cluster interconnect, be aware that the location of the cluster transport adapter might differ by node platform. For example, some platforms have the cluster transport adapters mounted on the motherboards. In other platforms, the cluster transport adapters are mounted on system and I/O boards.

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

2. Install the transport cables and (optionally) transport junctions, depending on how many nodes are in your cluster:

- A cluster with only two nodes can use a point-to-point connection, requiring no cluster transport junctions. You need to use a point-to-point (crossover) Ethernet cable if you are connecting 100BaseT or TPE ports of 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. See [FIGURE 3-1](#).



FIGURE 3-1 Typical Two-Node Sun Cluster Interconnect

- A cluster with more than two nodes requires two cluster transport junctions. These transport junctions are Ethernet-based switches (customer-supplied). See [FIGURE 3-2](#).

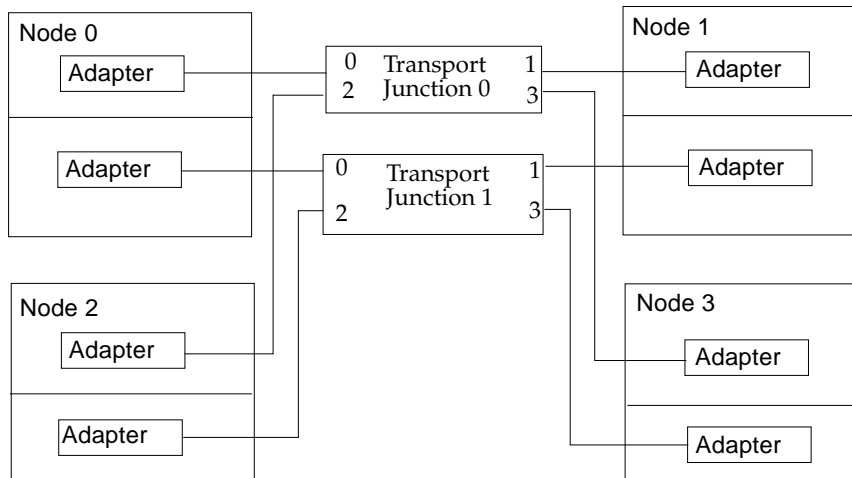


FIGURE 3-2 Typical Four-Node Cluster Interconnect

3. Install the public network adapters.

For the procedure on physically adding public network adapters, see the hardware documentation that shipped with your node and public network adapters.

Where to Go From Here

To install your storage, go to the section for your storage platform:

- [“Installing a StorEdge MultiPack Enclosure” on page 46](#)
- [“Installing a StorEdge D1000 Disk Array” on page 72](#)
- [“Installing a StorEdge A5x00 Disk Array” on page 100](#)
- [“Installing a Sun StorEdge A3500/A3500FC System” on page 124](#)
- [“Installing a StorEdge T3 Disk Tray” on page 174](#)

To review the task map for installing cluster hardware, see [“Installing Sun Cluster Hardware” on page 3](#).

Maintaining Cluster Interconnect and Public Network Hardware

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

TABLE 3-2 Task Map: Maintaining Cluster Interconnect and Public Network Hardware

Task	For Instructions, Go To
Add cluster transport adapters	"How to Add Cluster Transport Adapters" on page 36
Replace cluster transport adapters	"How to Replace Cluster Transport Adapters" on page 37
Remove cluster transport adapters	"How to Remove Cluster Transport Adapters" on page 39
Add public network adapters	"How to Add Public Network Adapters" on page 41
Remove public network adapters	"How to Remove Public Network Adapters" on page 42
Replace public network adapters	"How to Remove Public Network Adapters" on page 42

Adding Cluster Transport Adapters

This section provides the procedure for adding cluster transport adapters to a running cluster.

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

▼ How to Add Cluster Transport Adapters

1. **Shut down the node in which you are installing the cluster transport adapter.**

```
# scswitch -S -h nodename  
# shutdown -y -g0 -i0
```

For the procedure on shutting down a node, see the *Sun Cluster 3.0 U1 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 as the procedure in a non-cluster environment. For the procedure on installing network adapters, see the documentation that shipped with your network adapter and node hardware.

4. **Power on and boot the node.**

```
# boot -r
```

For the procedures on powering on and booting a node, see the *Sun Cluster 3.0 U1 System Administration Guide*.

Where to Go From Here

If the cluster transport adapter you add appears in the Sun Cluster software configuration, and you want to reconfigure Sun Cluster with the new cluster transport adapter, see the *Sun Cluster 3.0 U1 System Administration Guide*.

Replacing Cluster Transport Adapters

This section provides the procedure for replacing a failed cluster transport adapter in a running cluster.

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

▼ How to Replace Cluster Transport Adapters



Caution – You must maintain at least one cluster interconnect between the nodes of a cluster. The cluster does not function properly without a working cluster interconnect.

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

```
# scswitch -S -h nodename  
# shutdown -y -g0 -i0
```

For the procedure on shutting down a node, see the *Sun Cluster 3.0 U1 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 as the procedure 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 as the procedure in a non-cluster environment. For the procedure on replacing network 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 as the procedure 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 the *Sun Cluster 3.0 U1 System Administration Guide*.

Where to Go From Here

If 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 the *Sun Cluster 3.0 U1 System Administration Guide*.

Removing Cluster Transport Adapters

This section provides the procedure for removing an unused cluster transport adapter from a running cluster.

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

▼ How to Remove Cluster Transport Adapters



Caution – You must maintain at least one cluster interconnect between the nodes of a cluster. The cluster does not function properly without a working 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 remove appears in the Sun Cluster software configuration, remove the cluster transport adapter from the Sun Cluster configuration. To remove a cluster interconnect, follow the procedures in the *Sun Cluster 3.0 System Administration Guide* before going to [Step 2](#).
 - If the cluster transport adapter you want to remove does not appear in the Sun Cluster software configuration, go to [Step 2](#).
- 2. Shut down the node with the cluster transport adapter you want to remove.**

```
# scswitch -s -h nodename  
# shutdown -y -g0 -i0
```

For the procedure on shutting down a node, see the *Sun Cluster 3.0 U1 System Administration Guide*.

- 3. Power off the node.**

For the procedure on powering off a node, see the documentation that shipped with your node.
- 4. Disconnect the cluster transport cables from the cluster transport adapter you want to remove.**

This procedure is the same as the procedure 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.

5. Remove the cluster transport adapter.

This procedure is the same as the procedure in a non-cluster environment. For the procedure on removing network 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 the *Sun Cluster 3.0 U1 System Administration Guide*.

Adding Public Network Adapters

Physically adding public network adapters to a node in a cluster is no different from adding public network adapters in a non-cluster environment.

▼ How to Add Public Network Adapters

For the procedure on physically adding public network adapters, see the hardware documentation that shipped with your node and public network adapters.

Where to Go From Here

To add this public network adapter to a NAFO group, see the *Sun Cluster 3.0 U1 System Administration Guide*.

Replacing Public Network Adapters

Physically replacing public network adapters to a node in a cluster is no different from replacing public network adapters in a non-cluster environment.

▼ How to Replace Public Network Adapters

For the procedure on physically replacing public network adapters, see the hardware documentation that shipped with your node and public network adapters.

Where to Go From Here

To add this public network adapter to a NAFO group, see the *Sun Cluster 3.0 U1 System Administration Guide*.

Removing Public Network Adapters



Caution – You must maintain at least one public network connection for each node in the cluster. You cannot access the cluster without a public network connection.

Removing public network adapters from a node in a cluster is no different from removing public network adapters in a non-cluster environment. For procedures related to administering public network connections, see the *Sun Cluster 3.0 U1 System Administration Guide*.

▼ How to Remove Public Network Adapters

For the procedure on removing public network adapters, see the hardware documentation that shipped with your node and public network adapters.

Sun Gigabit Ethernet Adapter Considerations

Some Gigabit Ethernet switches require some device parameter values to be set differently than the defaults. Chapter 3 of the *Sun Gigabit Ethernet/P 2.0 Adapter Installation and User's Guide* describes the procedure for changing device parameters. The procedure used on nodes running Sun Cluster 3.0 software varies slightly from the procedure described in the guide. In particular, the difference is in how you derive parent names for use in the `ge.conf` file from the `/etc/path_to_inst` file.

Chapter 3 of the *Sun Gigabit Ethernet/P 2.0 Adapter Installation and User's Guide* describes the procedure for changing `ge` device parameter values through entries in the `/kernel/drv/ge.conf` file. The procedure to derive the parent name from the `/etc/path_to_inst` listing (to be used in `ge.conf` entries) appears in "Setting Driver Parameters Using a `ge.conf` File." For example, from the following `/etc/path_to_inst` line, you can derive the parent name for `ge2` to be `/pci@4,4000`.

```
"/pci@4,4000/network@4" 2 "ge"
```

On Sun Cluster 3.0 nodes, a `/node@nodeid` prefix appears in the `/etc/path_to_inst` line. Do *not* consider the `/node@nodeid` prefix when you derive the parent name. For example, on a cluster node, an equivalent `/etc/path_to_inst` entry would be the following:

```
"/node@l1/pci@4,4000/network@4" 2 "ge"
```

The parent name for `ge2`, to be used in the `ge.conf` file is still `/pci@4,4000` in this instance.

Installing and Maintaining a 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 Enclosure” on page 46](#)
- [“How to Add Disk Drive to StorEdge Multipack Enclosure in a Running Cluster” on page 52](#)
- [“How to Replace a Disk Drive in StorEdge MultiPack Enclosure in a Running Cluster” on page 55](#)
- [“How to Remove a Disk Drive From a StorEdge MultiPack Enclosure in Running Cluster” on page 59](#)
- [“How to Add a StorEdge MultiPack Enclosure to a Running Cluster” on page 60](#)
- [“How to Replace a StorEdge MultiPack Enclosure in a Running Cluster” on page 67](#)
- [“How to Remove a StorEdge MultiPack Enclosure From a Running Cluster” on page 69](#)

For conceptual information on multihost disks, see the *Sun Cluster 3.0 U1 Concepts* document.

Installing a StorEdge MultiPack Enclosure

This section describes the procedure for an initial installation of a StorEdge MultiPack enclosure.

▼ How to Install a StorEdge MultiPack Enclosure

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 the *Sun Cluster 3.0 U1 Installation Guide* and your server hardware manual.

Multihost storage in clusters uses the multi-initiator capability of the Small Computer System Interface (SCSI) specification. For conceptual information on multi-initiator capability, see the *Sun Cluster 3.0 U1 Concepts* document.



Caution – SCSI-reservations failures have been observed when clustering StorEdge MultiPack enclosures that contain a particular model of Quantum disk drive: SUN4.2G VK4550J. Avoid the use of this particular model of Quantum disk drive for clustering with StorEdge MultiPack enclosures. If you do use this model of disk drive, you must set the `scsi-initiator-id` of the “first node” to 6. If you are using a six-slot StorEdge MultiPack enclosure, you must also set the enclosure for the 9-through-14 SCSI target address range. For more information, see the *Sun StorEdge MultiPack Storage Guide*.

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 node that has SCSI address 7 as the “second node.”

To avoid conflicts, in [Step 7](#) you change the `scsi-initiator-id` of the remaining host adapter in the SCSI chain to an available SCSI address. This procedure refers to the node that has an available SCSI address as the “first node.”

For a partial list of `nvrामrc` editor and `nvedit` keystroke commands, see [Appendix B](#) of this guide. For a full list of commands, see the *OpenBoot 3.x Command Reference Manual*.

Note – Even though a slot in the StorEdge MultiPack enclosure might not be in use, do not set 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 nodes that will be connected to the StorEdge MultiPack enclosure.

For the procedure on installing a host adapter, see the documentation that shipped with your host adapter and node hardware.

3. Connect the cables to the StorEdge MultiPack enclosure, as shown in [FIGURE 4-1](#).

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

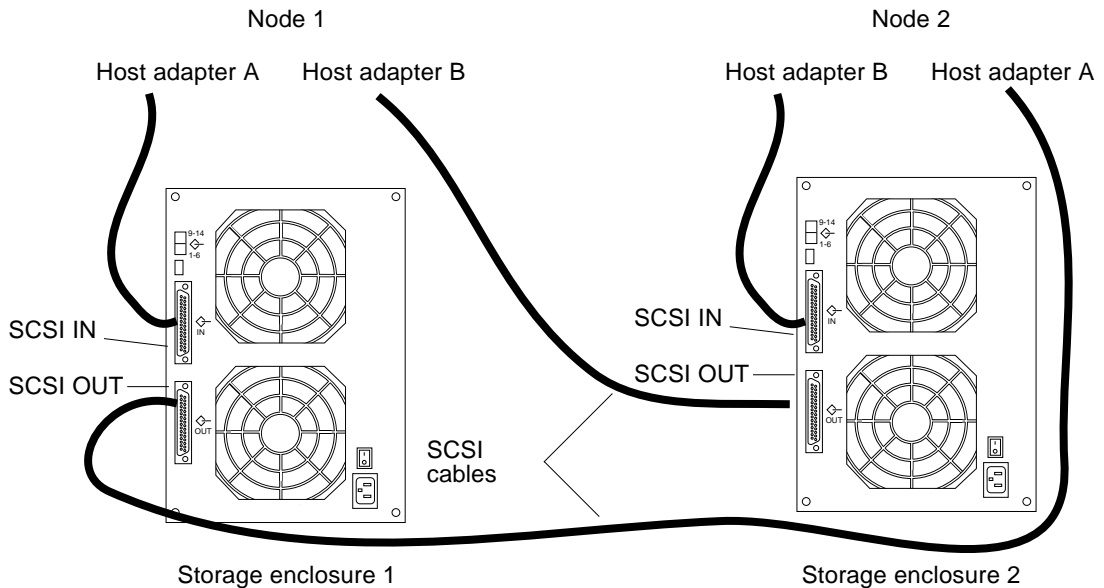


FIGURE 4-1 Example of a StorEdge MultiPack Enclosure Mirrored Pair

4. Connect the AC power cord for each StorEdge MultiPack enclosure of the mirrored pair to a different power source.

5. Power on the *first* node but do not allow it to boot. If necessary, halt the node to continue with OpenBoot™ PROM (OBP) Monitor tasks. The first node is the node with an available SCSI address.
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 are to 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 in [Step 7](#). 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 partial list of `nvrarc` editor and `nvedit` keystroke commands, see [Appendix B](#). For a full list of commands, see the *OpenBoot 3.x Command Reference Manual*.

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

Note – Insert exactly one space after the first quotation mark 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 recorded on a temporary copy of the `nvrामrc` 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 `nvrामrc` script you created in [Step 7](#), as shown in the following example.

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

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

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

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

11. Power on the *second* node but do not allow it to boot. If necessary, halt the node to continue with OpenBoot PROM Monitor tasks. The second node is the node that has SCSI address 7.

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 (as in [Step 6](#)). 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, as shown in the following example.

```
{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 the *Sun Cluster 3.0 U1 Installation Guide*.

Maintaining a StorEdge MultiPack

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

TABLE 4-1 Task Map: Maintaining a StorEdge MultiPack Enclosure

Task	For Instructions, Go To
Add a disk drive	“How to Add Disk Drive to StorEdge Multipack Enclosure in a Running Cluster” on page 52
Replace a disk drive	“How to Replace a Disk Drive in StorEdge MultiPack Enclosure in a Running Cluster” on page 55
Remove a disk drive	“How to Remove a Disk Drive From a StorEdge MultiPack Enclosure in a Running Cluster” on page 59
Add a StorEdge MultiPack enclosure	“How to Add a StorEdge MultiPack Enclosure to a Running Cluster” on page 60
Replace a StorEdge MultiPack enclosure	“How to Replace a StorEdge MultiPack Enclosure in a Running Cluster” on page 67
Remove a StorEdge MultiPack enclosure	“How to Remove a StorEdge MultiPack Enclosure From a Running Cluster” on page 69

▼ How to Add Disk Drive to StorEdge Multipack Enclosure in a Running Cluster

Use this procedure to add a disk drive to a running cluster. Perform the steps in this procedure in conjunction with the procedures in the *Sun Cluster 3.0 U1 System Administration Guide* and your server hardware manual. “[Example—Adding a StorEdge MultiPack Disk Drive](#)” on [page 54](#) shows how to apply this procedure.

For conceptual information on quorums, quorum devices, global devices, and device IDs, see the *Sun Cluster 3.0 U1 Concepts* document.



Caution – SCSI-reservations failures have been observed when clustering StorEdge MultiPack enclosures that contain a particular model of Quantum disk drive: SUN4.2G VK4550J. Avoid the use of this particular model of Quantum disk drive for clustering with StorEdge MultiPack enclosures. If you do use this model of disk drive, you must set the `scsi-initiator-id` of the “first node” to 6. If you are using a six-slot StorEdge MultiPack enclosure, you must also set the enclosure for the 9-through-14 SCSI target address range. For more information, see the *Sun StorEdge MultiPack Storage Guide*.

1. Locate an empty disk slot in the StorEdge MultiPack 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 StorEdge MultiPack enclosure, or by removing the side-cover of the unit. The target address IDs that correspond to the slots appear on the middle partition of the drive bay.

2. Install the disk drive.

For detailed instructions, see the documentation that shipped with your StorEdge MultiPack enclosure.

3. On all nodes that are attached to the StorEdge MultiPack enclosure, 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 necessary, 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 that is 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
```

Note – As shown in [“Example—Adding a StorEdge MultiPack Disk Drive”](#) on page 54, the DID 35 that is assigned to the new disk drive might not be in sequential order in the StorEdge MultiPack enclosure.

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.

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/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
```

Where to Go From Here

To configure a disk drive as a quorum device, see the *Sun Cluster 3.0 U1 System Administration Guide* for the procedure on adding a quorum device.

▼ How to Replace a Disk Drive in StorEdge MultiPack Enclosure in a Running Cluster

Use this procedure to replace a StorEdge MultiPack enclosure disk drive. “[Example—Replacing a StorEdge MultiPack Disk Drive](#)” on [page 58](#) shows how to apply this procedure. Perform the steps in this procedure in conjunction with the procedures in *Sun Cluster 3.0 U1 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 quorums, quorum devices, global devices, and device IDs, see the *Sun Cluster 3.0 U1 Concepts* document.



Caution – SCSI-reservations failures have been observed when clustering StorEdge MultiPack enclosures that contain a particular model of Quantum disk drive: SUN4.2G VK4550J. Avoid the use of this particular model of Quantum disk drive for clustering with StorEdge MultiPack enclosures. If you do use this model of disk drive, you must set the `scsi-initiator-id` of the “first node” to 6. If you are using a six-slot StorEdge MultiPack enclosure, you must also set the enclosure for the 9-through-14 SCSI target address range (for more information, see the *Sun StorEdge MultiPack Storage Guide*).

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, put the quorum device into maintenance state before you go to [Step 3](#). For the procedure on putting a quorum device into maintenance state, see the *Sun Cluster 3.0 U1 System Administration Guide*.
- If the disk is not a quorum device, go to [Step 3](#).

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](#) 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 you partition the new disk drive.

If you are using VERITAS Volume Manager, skip this step and go to [Step 7](#).

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

Note – 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 *Sun StorEdge MultiPack Storage Guide*.

8. On one node that is attached to the StorEdge MultiPack 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 require at least five minutes to complete.

```
# devfsadm
```

9. If you are using Solstice DiskSuite as your volume manager, from any node that is connected to the StorEdge MultiPack enclosure, partition the new disk drive by using the partitioning you saved in [Step 6](#).

If you are using VERITAS Volume Manager, skip this step and go to [Step 10](#).

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

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

```
# scswitch -S -h nodename  
# shutdown -y -g0 -i6
```

For more information, see the *Sun Cluster 3.0 U1 System Administration Guide*.

11. From any node that is 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 that was identified in [Step 5](#).

If the new physical DID is different from the physical DID in [Step 5](#), 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 that is 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.

For the procedure on adding a quorum device, see the *Sun Cluster 3.0 U1 System Administration Guide*.

Example—Replacing a StorEdge MultiPack Disk Drive

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

```
# sddidadm -l d20
20      phys-schost-2:/dev/rdisk/c3t2d0 /dev/did/rdsk/d20
# sddidadm -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 -g0 -i6
...
# sddidadm -R d20
# sddidadm -o diskid -l c3t2d0
5345414741544520393735314336363037370000
# sddidadm -ui
```


▼ How to Remove a Disk Drive From a StorEdge MultiPack Enclosure in Running Cluster

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 the *Sun Cluster 3.0 U1 System Administration Guide* and your server hardware manual.

For conceptual information on quorum, quorum devices, global devices, and device IDs, see the *Sun Cluster 3.0 U1 Concepts* document.

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

```
# scstat -q
```

- If the disk drive you want to replace is a quorum device, put the quorum device into maintenance state before you go to [Step 2](#). For the procedure on putting a quorum device into maintenance state, see the *Sun Cluster 3.0 U1 System Administration Guide*.
- If the disk is not a quorum device, go to [Step 2](#).

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 from which the disk drive 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  
# cfgadm -al
```

4. Remove the disk drive.

For more information on the procedure for removing a disk drive, see the *Sun StorEdge MultiPack Storage Guide*.

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

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

▼ How to Add a StorEdge MultiPack Enclosure to a Running Cluster

Use this procedure to install a StorEdge MultiPack enclosure in a running cluster. Perform the steps in this procedure in conjunction with the procedures in the *Sun Cluster 3.0 U1 Installation Guide* and your server hardware manual.

For conceptual information on multi-initiator SCSI and device IDs, see the *Sun Cluster 3.0 U1 Concepts* document.



Caution – Quorum failures have been observed when clustering StorEdge Multipack enclosures that contain a particular model of Quantum disk drive: SUN4.2G VK4550J. Avoid the use of this particular model of Quantum disk drive for clustering with StorEdge Multipack enclosures. If you do use this model of disk drive, you must set the `scsi-initiator-id` of the “first node” to 6. If you are using a six-slot StorEdge Multipack, you must also set the enclosure for the 9-through-14 SCSI target address range (for more information, see the *Sun StorEdge MultiPack Storage Guide*).

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 node with SCSI address 7 as the “second node.”

To avoid conflicts, in [Step 9](#) you change the `scsi-initiator-id` of the remaining host adapter in the SCSI chain to an available SCSI address. This procedure refers to the node with an available SCSI address as the “first node.”

For a partial list of `nvrampc` editor and `nvedit` keystroke commands, see [Appendix B](#) of this guide. For a full list of commands, see the *OpenBoot 3.x Command Reference Manual*.

Note – Even though a slot in the StorEdge MultiPack enclosure might not be in use, do not set 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 -g0 -i0
```

For more information, see the *Sun Cluster 3.0 U1 System Administration Guide*.

3. Install the host adapters in the first node.

For the procedure on installing a host adapter, see the documentation that shipped with your host adapter and node.

4. Connect the single-ended SCSI cable between the node and the StorEdge MultiPack enclosures, as shown in [FIGURE 4-2](#).

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

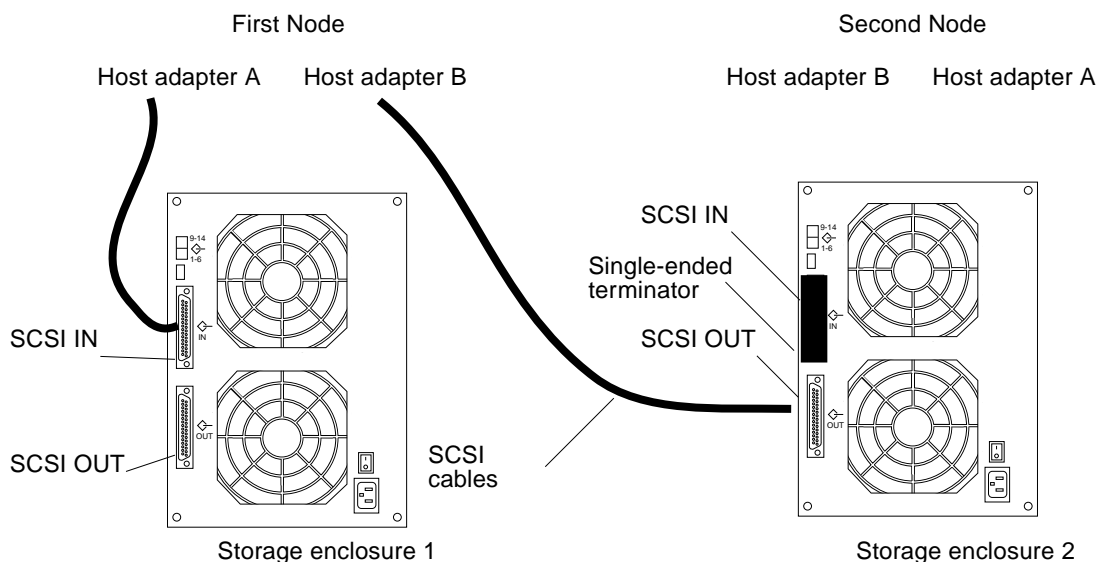


FIGURE 4-2 Example of a StorEdge MultiPack Enclosure Mirrored Pair

5. Temporarily install a single-ended terminator on the SCSI IN port of the *second* StorEdge MultiPack enclosure, as shown in [FIGURE 4-2](#).

6. Connect each StorEdge MultiPack enclosure of the mirrored pair to different power sources.

7. Power on the first node and the StorEdge MultiPack 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 are to 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 in [Step 9](#). 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 partial list of `nvrarc` editor and `nvedit` keystroke commands, see [Appendix B](#). For a full list of commands, see the *OpenBoot 3.x Command Reference Manual*.

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 quotation mark 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 recorded on a temporary copy of the `nvrामrc` 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 `nvrामrc` script you created in [Step 9](#), as shown in the following example.

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

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

12. Instruct the OpenBoot PROM Monitor to use the `nvrामrc` script, as shown in the following example.

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

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

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

For more information, see the *Sun Cluster 3.0 U1 System Administration Guide*.

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

```
# scdidadm -l
```

15. Shut down and power off the *second* node.

```
# scswitch -s -h nodename  
# shutdown -y -g0 -i0
```

16. Install the host adapters in the second node.

For the procedure on installing a host adapter, see the documentation that shipped with your host adapter and node.

17. Remove the SCSI terminator you installed in [Step 5](#).
18. Connect the StorEdge MultiPack enclosures to the host adapters by using single-ended SCSI cables.

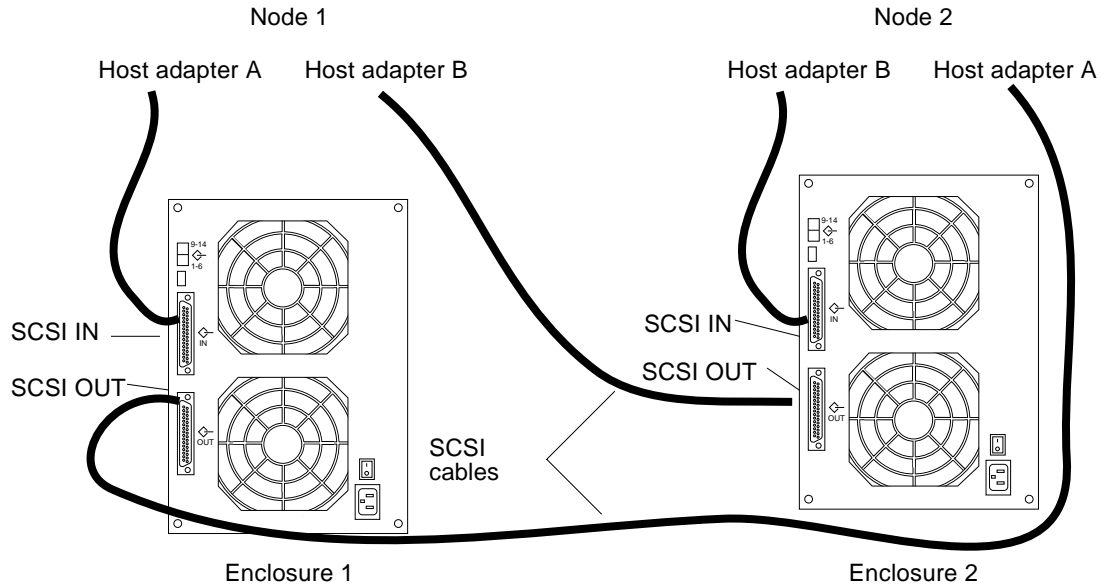


FIGURE 4-3 Example of a StorEdge MultiPack Enclosure Mirrored Pair

19. Power on the *second* node but do not allow it to boot. If necessary, halt the node to continue with OpenBoot PROM Monitor tasks.
20. Verify that the second node checks for 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 that are 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
```

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 StorEdge MultiPack enclosure.

```
# scdidadm -l
```

24. Perform volume management administration to add the disk drives in the StorEdge MultiPack enclosure to the volume management configuration.

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

▼ How to Replace a StorEdge MultiPack Enclosure in a Running Cluster

Use this procedure to replace a StorEdge MultiPack enclosure in a running cluster. This procedure assumes that you are retaining the disk drives in the StorEdge MultiPack enclosure that you are replacing and that you are retaining the references to these same disk drives.

If you want to replace your disk drives, see [“How to Replace a Disk Drive in StorEdge MultiPack Enclosure in a Running Cluster”](#) on page 55.

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 StorEdge MultiPack enclosure, disconnecting the cable on the SCSI OUT connector first, then the cable on the SCSI IN connector second (see [FIGURE 4-4](#)).




FIGURE 4-4 Disconnecting the SCSI Cables

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

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

5. **Connect the new StorEdge MultiPack enclosure to an AC power source.**

Refer to the documentation that shipped with the StorEdge MultiPack enclosure and the labels inside the lid of the StorEdge MultiPack enclosure.

6. **Connect the SCSI cables to the new StorEdge MultiPack enclosure, reversing the order in which you disconnected them (connect the SCSI IN connector first, then the SCSI OUT connector second). See [FIGURE 4-4](#).**

7. **Move the disk drives one at a time from the old StorEdge MultiPack enclosure to the same slots in the new StorEdge MultiPack enclosure.**

8. **Power on the StorEdge MultiPack enclosure.**

9. **On all nodes that are attached to the StorEdge MultiPack enclosure, run the `devfsadm(1M)` command.**

```
# devfsadm
```

10. **One at a time, shut down and reboot the nodes that are connected to the StorEdge MultiPack enclosure.**

```
# scswitch -S -h nodename  
# shutdown -y -g0 -i6
```

For more information on `shutdown(1M)`, see the *Sun Cluster 3.0 U1 System Administration Guide*.

11. **Perform volume management administration to add the StorEdge MultiPack enclosure to the configuration.**

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

▼ How to Remove a StorEdge MultiPack Enclosure From a Running Cluster

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 StorEdge MultiPack enclosure from the configuration.**

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

2. **Disconnect the SCSI cables from the StorEdge MultiPack enclosure, disconnecting them in the order that is shown in [FIGURE 4-5](#).**

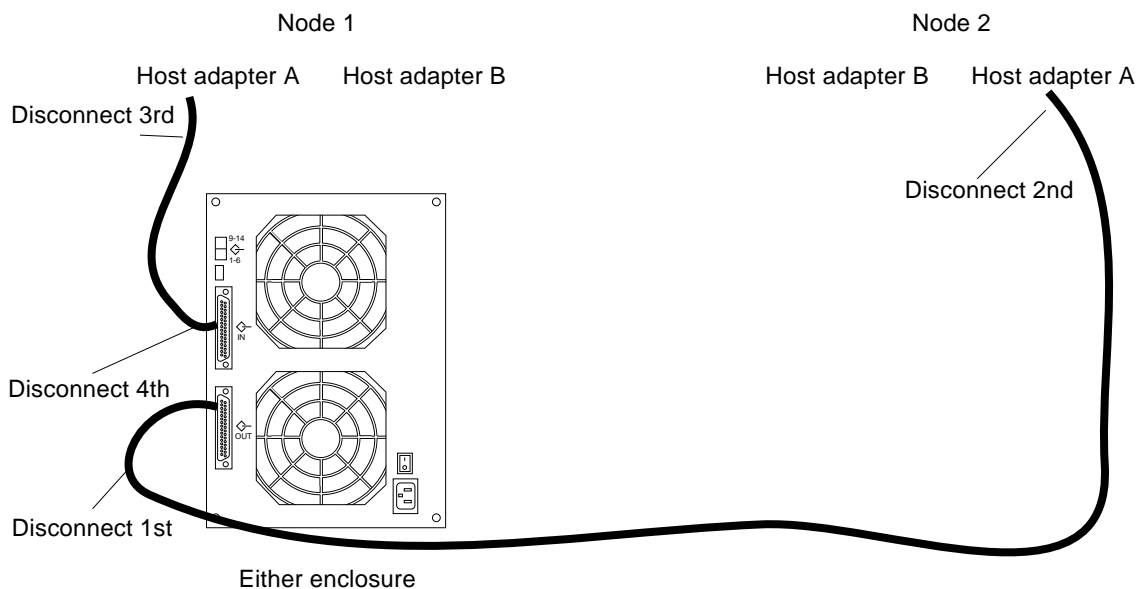


FIGURE 4-5 Disconnecting the SCSI Cables

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

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

4. Remove the StorEdge MultiPack enclosure.

For the procedure on removing an enclosure, see the *Sun StorEdge MultiPack Storage Guide*.

5. Identify the disk drives you need to remove from the cluster.

```
# cfigadm -al
```

6. On all nodes, remove references to the disk drives that were in the StorEdge MultiPack enclosure you removed in [Step 4](#).

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

7. If necessary, remove any unused host adapters from the nodes.

For the procedure on removing a host adapter, see the documentation that shipped with your host adapter and node.

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 Disk Array” on page 72](#)
- [“How to Add a Disk Drive in a StorEdge D1000 Disk Array in a Running Cluster” on page 78](#)
- [“How to Replace a Disk Drive in a StorEdge D1000 Disk Array in a Running Cluster” on page 81](#)
- [“How to Remove a Disk Drive From a StorEdge D1000 Disk Array in a Running Cluster” on page 85](#)
- [“How to Add a StorEdge D1000 Disk Array to a Running Cluster” on page 87](#)
- [“How to Replace a StorEdge D1000 Disk Array in a Running Cluster” on page 94](#)
- [“How to Remove a StorEdge D1000 Disk Array From a Running Cluster” on page 96](#)

For conceptual information on multihost disks, see the *Sun Cluster 3.0 U1 Concepts* document.

Installing a StorEdge D1000 Disk Array

This section provides the procedure for an initial installation of a StorEdge D1000 disk array.

▼ How to Install a StorEdge D1000 Disk Array

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 the *Sun Cluster 3.0 U1 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 the *Sun Cluster 3.0 U1 Concepts* document.

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](#) you 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.

Note – Even though a slot in the disk array might not be in use, do not set 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 node that you are connecting to the StorEdge D1000 disk array.

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

3. Connect the cables to the StorEdge D1000 disk arrays, as shown in [FIGURE 5-1](#).

Make sure that the *entire* bus length that is connected to each StorEdge D1000 disk array is less than 25 m. This measurement includes the cables to both nodes, as well as the bus length internal to each StorEdge D1000 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 StorEdge D1000 disk array of the mirrored pair to a different power source.

5. Power on the *first* node and the StorEdge D1000 disk arrays.

For the procedure on powering on a StorEdge D1000 disk array, see the *Sun StorEdge A1000 and D1000 Installation, Operations, and Service Manual*.

6. Find the paths to the host adapters.

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

Identify and record the two controllers that are to be connected to the storage devices and record these paths. Use this information to change the SCSI addresses of these controllers in the `nvrsrc` script in [Step 7](#). Do not include the `/sd` directories in the device paths.

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

For a list of `nvrsrc` 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 recorded on a temporary copy of the `nvrामrc` 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 `nvrामrc` script you created in [Step 7](#), as shown in the following example.

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

```
{0} ok printenv nvrामrc
nvrामrc =                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 (OBP) Monitor to use the `nvrामrc` script, as shown in the following example.

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

11. Power on the *second* node, but do not allow it to boot. If necessary, halt the system to continue with OBP 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 that are 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 the *Sun Cluster 3.0 U1 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 5-1 Task Map: Maintaining a StorEdge D1000 Disk Array

Task	For Instructions, Go To
Add a disk drive	“How to Add a Disk Drive in a StorEdge D1000 Disk Array in a Running Cluster” on page 78
Replace a disk drive	“How to Replace a Disk Drive in a StorEdge D1000 Disk Array in a Running Cluster” on page 81
Remove a disk drive	“How to Remove a Disk Drive From a StorEdge D1000 Disk Array in a Running Cluster” on page 85
Add a StorEdge D1000 disk array	“How to Add a StorEdge D1000 Disk Array to a Running Cluster” on page 87
Replace a StorEdge D1000 disk array	“How to Replace a StorEdge D1000 Disk Array in a Running Cluster” on page 94
Remove a StorEdge D1000 disk array	“How to Remove a StorEdge D1000 Disk Array From a Running Cluster” on page 96

▼ How to Add a Disk Drive in a StorEdge D1000 Disk Array in a Running Cluster

Use this procedure to add a disk drive to a running cluster. Perform the steps in this procedure in conjunction with the procedures in the *Sun Cluster 3.0 U1 System Administration Guide* and your server hardware manual. “[Example—Adding a StorEdge D1000 Disk Drive](#)” on page 80 shows how to apply this procedure.

For conceptual information on quorum, quorum devices, global devices, and device IDs, see the *Sun Cluster 3.0 U1 Concepts* document.

1. Locate an empty disk slot in the StorEdge D1000 disk array for the disk drive you are adding.

Identify the disk slot in the StorEdge D1000 disk array for the disk drive that you are adding and note the target number. Refer to the documentation that shipped with your StorEdge D1000 disk array.

2. Install the disk drive.

For the procedure on installing a disk drive, see the *Sun StorEdge D1000 Storage Guide*.

3. On all nodes attached to the StorEdge D1000 disk array, configure the disk drive.

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

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

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

5. If necessary, 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 that is 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.

```
# sctidadm -l
```

Note – As shown in [“Example—Adding a StorEdge D1000 Disk Drive”](#) on page 80, the DID 35 that is 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.

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.

```
# 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/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
```

Where to Go From Here

To configure a disk drive as a quorum device, see the *Sun Cluster 3.0 U1 System Administration Guide* for the procedure on adding a quorum device.

▼ How to Replace a Disk Drive in a StorEdge D1000 Disk Array in a Running Cluster

Use this procedure to replace a StorEdge D1000 disk array disk drive. Perform the steps in this procedure in conjunction with the procedures in the *Sun Cluster 3.0 U1 System Administration Guide* and your server hardware manual. Use the procedures in your server hardware manual to identify a failed disk drive.

“[Example—Replacing a StorEdge D1000 Disk Drive](#)” on page 84 shows how to apply this procedure.

For conceptual information on quorum, quorum devices, global devices, and device IDs, see the *Sun Cluster 3.0 U1 Concepts* document.

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 are replacing is a quorum device.

```
# scstat -q
```

- If the disk drive you are replacing is a quorum device, put the quorum device into maintenance state before you go to [Step 3](#). For the procedure on putting a quorum device into maintenance state, see the *Sun Cluster 3.0 U1 System Administration Guide*.
- If the disk is not a quorum device, go to [Step 3](#).

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](#) 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, go to [Step 7](#).

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

Note – 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 the procedure on replacing a disk drive, see the *Sun StorEdge D1000 Storage Guide*.

8. On one node attached to the StorEdge D1000 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 that are connected to the node, the `devfsadm` command can require at least five minutes to complete.

```
# devfsadm
```

9. If you are using Solstice DiskSuite as your volume manager, from any node that is connected to the StorEdge D1000 disk array, partition the new disk drive by using the partitioning you saved in [Step 6](#).

If you are using VERITAS Volume Manager, go to [Step 10](#).

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


10. One at a time, shut down and reboot the nodes that are connected to the StorEdge D1000 disk array.

```
# scswitch -S -h nodename
# shutdown -y -g0 -i6
```

For more information on shutdown procedures, see the *Sun Cluster 3.0 U1 System Administration Guide*.

11. From any node that is 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 that was identified in [Step 5](#).

If the new physical DID is different from the physical DID that was identified in [Step 5](#), you successfully replaced the failed disk drive with a new disk drive.

```
# scdidadm -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 that is 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.

For the procedure on adding a quorum device, see the *Sun Cluster 3.0 U1 System Administration Guide*.

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.

```
# 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 -g0 -i6
...
# scdidadm -R d20
# scdidadm -o diskid -l c3t2d0
5345414741544520393735314336363037370000
# scdidadm -ui
```

▼ How to Remove a Disk Drive From a StorEdge D1000 Disk Array in a Running Cluster

Use this procedure to remove a disk drive from a StorEdge D1000 disk array in a running cluster. Perform the steps in this procedure in conjunction with the procedures in the *Sun Cluster 3.0 U1 System Administration Guide* and your server hardware manual.

For conceptual information on quorum, quorum devices, global devices, and device IDs, see the *Sun Cluster 3.0 U1 Concepts* document.

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

```
# scstat -q
```

- If the disk drive you want to replace is a quorum device, put the quorum device into maintenance state before you go to [Step 2](#). For the procedure on putting a quorum device into maintenance state, see the *Sun Cluster 3.0 U1 System Administration Guide*.
- If the disk is not a quorum device, go to [Step 2](#).

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 the procedure on removing a disk drive, see the *Sun StorEdge D1000 Storage Guide*.

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

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

▼ How to Add a StorEdge D1000 Disk Array to a Running Cluster

Use this procedure to install a StorEdge D1000 disk array in a running cluster. Perform the steps in this procedure in conjunction with the procedures in the *Sun Cluster 3.0 U1 System Administration Guide* and your server hardware manual.

For conceptual information on multi-initiator SCSI and device IDs, see the *Sun Cluster 3.0 U1 Concepts* document.

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](#) you 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.

Note – Even though a slot in the disk array might not be in use, do not set 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 -g0 -i0
```

For more information on shutdown procedures, see the *Sun Cluster 3.0 U1 System Administration Guide*.

3. Install the host adapters in the first 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 StorEdge D1000 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 about SCSI operation.

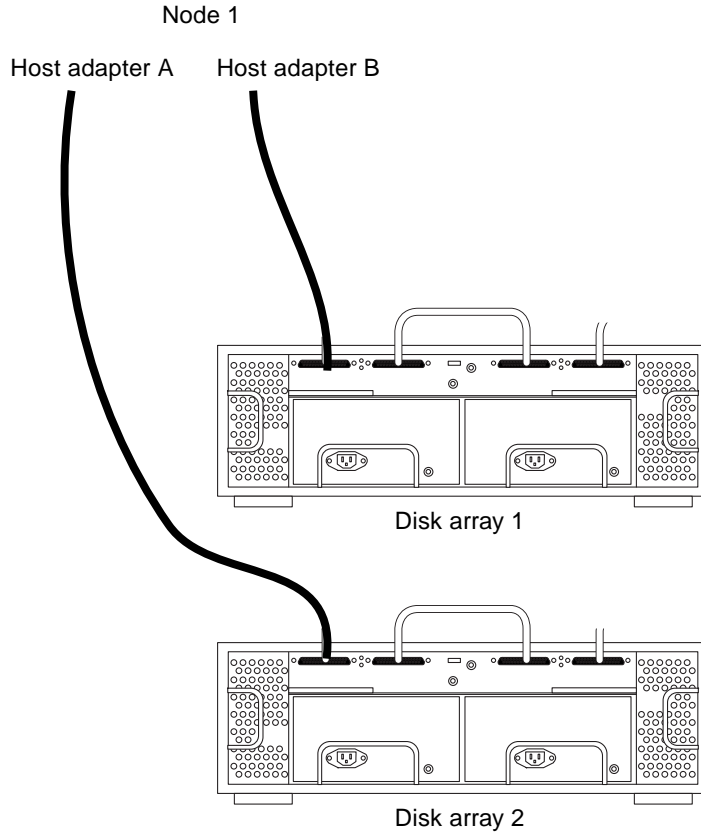


FIGURE 5-2 Example of a StorEdge D1000 Disk Array Mirrored Pair

5. Power on the first node and the StorEdge D1000 disk arrays.

For the procedure on powering on a StorEdge D1000 disk array, see the *Sun StorEdge A1000 and D1000 Installation, Operations, and Service Manual*.

6. Find the paths to the SCSI host adapters.

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

Identify and record the two controllers that are to be connected to the StorEdge D1000 disk arrays, and record these paths. Use this information to change the SCSI addresses of these controllers in the `nvrामrc` script in [Step 7](#). Do not include the `/sd` directories in the device paths.

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

For a list of Editor `nvrामrc` 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 recorded on a temporary copy of the `nvrामrc` 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 `nvrामrc` script you created in [Step 7](#), as shown in the following example.

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

```
{0} ok printenv nvrामrc
nvrामrc =
    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 `nvrामrc` script, as shown in the following example:

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


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

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

For more information, see the *Sun Cluster 3.0 U1 System Administration Guide*.

12. On all nodes, verify that the DIDs have been assigned to the disk drives in the StorEdge D1000 disk array.

```
# scdidadm -l
```

13. Shut down and power off the *second* node.

```
# scswitch -s -h nodename  
# shutdown -y -g0 -i0
```

For more information on shutdown procedures, see the *Sun Cluster 3.0 U1 System Administration Guide*.

14. Install the host adapters in the second node.

For the procedure on installing a host adapter, see the documentation that shipped with your host adapter and node.

15. Connect the StorEdge D1000 disk array to the host adapters by using differential SCSI cables (see [FIGURE 5-3](#)).




FIGURE 5-3 Example of a StorEdge D1000 Disk Array Mirrored Pair

16. Power on the second node, but do not allow it to boot. If necessary, halt the system to continue with OpenBoot PROM (OBP) Monitor tasks.
17. Verify that the second node checks for the new host adapters and disk drives.

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

18. 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 that are 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
...
```

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

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

20. On all nodes, verify that the DIDs have been assigned to the disk drives in the StorEdge D1000 disk array.

```
# scdidadm -l
```

21. 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 in a Running Cluster

Use this procedure to replace a StorEdge D1000 disk array in a running cluster. This procedure assumes that you are retaining the disk drives in the disk array you are replacing and you are retaining the references to these same disk drives.

If you are replacing your disk drives, see [“How to Replace a Disk Drive in a StorEdge D1000 Disk Array in a Running Cluster”](#) on page 81.

- 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 StorEdge D1000 disk array from the configuration.**
For more information, see your Solstice DiskSuite or VERITAS Volume Manager documentation.
- 3. Disconnect the SCSI cables from the StorEdge D1000 disk array, removing them in the order that is shown in [FIGURE 5-4](#).**




FIGURE 5-4 Disconnecting the SCSI Cables

4. **Power off and disconnect the StorEdge D1000 disk array from the AC power source.**

For the procedure on powering off a StorEdge D1000 disk array, see the *Sun StorEdge A1000 and D1000 Installation, Operations, and Service Manual*.

5. **Connect the new StorEdge D1000 disk array to an AC power source.**
6. **Connect the SCSI cables to the new StorEdge D1000 disk array, reversing the order they were disconnected, as shown in [FIGURE 5-4](#).**
7. **Move the disk drives one at a time from the old StorEdge D1000 disk array to the same slots in the new StorEdge D1000 disk array.**
8. **Power on the StorEdge D1000 disk array.**

For the procedure on powering on a StorEdge D1000 disk array, see the *Sun StorEdge A1000 and D1000 Installation, Operations, and Service Manual*.

9. **On all nodes that are attached to the StorEdge D1000 disk array, run the `devfsadm(1M)` command.**

```
# devfsadm
```

10. **One at a time, shut down and reboot the nodes that are connected to the StorEdge D1000 disk array.**

```
# scswitch -s -h nodename  
# shutdown -y -g0 -i6
```

For more information on shutdown procedures, see the *Sun Cluster 3.0 U1 System Administration Guide*.

11. **Perform volume management administration to add the StorEdge D1000 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 From a Running Cluster

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

1. **Perform volume management administration to remove the StorEdge D1000 disk array from the configuration.**

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

2. **Disconnect the SCSI cables from the StorEdge D1000 disk array, removing them in the order that is shown in [FIGURE 5-5](#).**




FIGURE 5-5 Disconnecting the SCSI cables

3. **Power off and disconnect the StorEdge D1000 disk array from the AC power source.**

For the procedure on powering off a StorEdge D1000 disk array, see the *Sun StorEdge A1000 and D1000 Installation, Operations, and Service Manual*.

4. **Remove the StorEdge D1000 disk array.**

For the procedure on removing a StorEdge D1000 disk array, see the *Sun StorEdge A1000 and D1000 Installation, Operations, and Service Manual*.

5. Identify the disk drives you need to remove.

```
# cfgadm -al
```

6. On all nodes, remove references to the disk drives in the StorEdge D1000 disk array you removed in [Step 4](#).

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

7. If necessary, remove any unused host adapters from the nodes.

For the procedure on removing a host adapter, see the documentation that shipped with your host adapter and node.

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 Disk Array” on page 100](#)
- [“How to Add a Disk Drive to a StorEdge A5x00 Disk Array in a Running Cluster” on page 103](#)
- [“How to Replace a Disk Drive in a StorEdge A5x00 Disk Array in a Running Cluster” on page 105](#)
- [“How to Remove a Disk Drive From a StorEdge A5x00 Disk Array in a Running Cluster” on page 110](#)
- [“How to Add the First StorEdge A5x00 Disk Array to a Running Cluster” on page 112](#)
- [“How to Add a StorEdge A5x00 Disk Array to a Running Cluster That Has Existing StorEdge A5x00 Disk Arrays” on page 115](#)
- [“How to Replace a StorEdge A5x00 Disk Array in a Running Cluster” on page 117](#)
- [“How to Remove a StorEdge A5x00 Disk Array From a Running Cluster” on page 120](#)

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

Installing a StorEdge A5x00 Disk Array

This section describes the procedure for an initial installation of a StorEdge A5x00 disk array.

▼ How to Install a StorEdge A5x00 Disk Array

Use this procedure to install a StorEdge A5x00 disk array. Perform the steps in this procedure in conjunction with the procedures in the *Sun Cluster 3.0 U1 Installation Guide* and your server hardware manual.

1. Install host adapters in the nodes that are to be connected to the StorEdge A5x00 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 StorEdge A5x00 disk array.

[FIGURE 6-1](#) shows a sample StorEdge A5x00 disk array configuration.

For more information on cabling and configuring StorEdge A5x00 disk arrays, see the *Sun StorEdge A5000 Installation and Service Manual*.

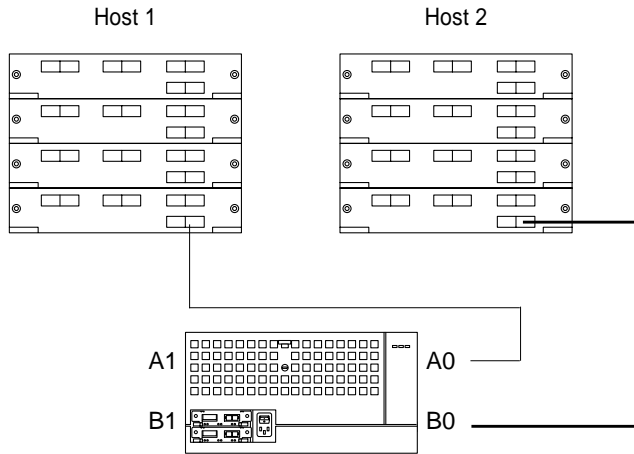


FIGURE 6-1 Sample StorEdge A5x00 Disk Array Configuration

- 3. Check the StorEdge A5x00 disk array controller firmware revision, and, if required, install the most recent firmware revision.**

For more information, see the *Sun StorEdge A5000 Product Notes*.

Where to Go From Here

To install software, follow the procedures in *Sun Cluster 3.0 U1 Installation Guide*.

Maintaining a StorEdge A5x00 Disk Array

This section describes the procedures for maintaining a StorEdge A5x00 disk array. [TABLE 6-1](#) lists these procedures.

TABLE 6-1 Task Map: Maintaining a Sun StorEdge A5x00 Disk Array

Task	For Instructions, Go To
Add a disk drive	“How to Add a Disk Drive to a StorEdge A5x00 Disk Array in a Running Cluster” on page 103
Replace a disk drive	“How to Replace a Disk Drive in a StorEdge A5x00 Disk Array in a Running Cluster” on page 105
Remove a disk drive	“How to Remove a Disk Drive From a StorEdge A5x00 Disk Array in a Running Cluster” on page 110
Add a StorEdge A5x00 disk array	“How to Add the First StorEdge A5x00 Disk Array to a Running Cluster” on page 112 or “How to Add a StorEdge A5x00 Disk Array to a Running Cluster That Has Existing StorEdge A5x00 Disk Arrays” on page 115
Replace a StorEdge A5x00 disk array	“How to Replace a StorEdge A5x00 Disk Array in a Running Cluster” on page 117
Remove a StorEdge A5x00 disk array	“How to Remove a StorEdge A5x00 Disk Array From a Running Cluster” on page 120

▼ How to Add a Disk Drive to a StorEdge A5x00 Disk Array in a Running Cluster

Use this procedure to add a disk drive to a running cluster. Perform the steps in this procedure in conjunction with the procedures in the *Sun Cluster 3.0 U1 System Administration Guide* and your server hardware manual.

For conceptual information on quorums, quorum devices, global devices, and device IDs, see the *Sun Cluster 3.0 U1 Concepts* document.

1. **On one node that is connected to the StorEdge A5x00 disk array, use the `luxadm(1M)` command to install 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 that are attached to the StorEdge A5x00 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 require 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 necessary, partition the disk drive.**

You can use either the `format(1M)` command or copy the partitioning from another disk drive in the StorEdge A5x00 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.

```
# sctidadm -l
```

Note – The DID that was assigned to the new disk drive might not be in sequential order in the StorEdge A5x00 disk array.

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

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

Where to Go From Here

To configure a disk drive as a quorum device, see the *Sun Cluster 3.0 U1 System Administration Guide* for the procedure on adding a quorum device.

▼ How to Replace a Disk Drive in a StorEdge A5x00 Disk Array in a Running Cluster

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

“[Example—Replacing 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 the *Sun Cluster 3.0 U1 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 quorums, quorum devices, global devices, and device IDs, see the *Sun Cluster 3.0 U1 Concepts* document.

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 are replacing is a quorum device.

```
# scstat -q
```

- If the disk drive you are replacing is a quorum device, put the quorum device into maintenance state before you go to [Step 3](#). For the procedure on putting a quorum device into maintenance state, see the *Sun Cluster 3.0 U1 System Administration Guide*.
- If the disk you are replacing is not a quorum device, go to [Step 3](#).

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](#) 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.

```
# 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, go to [Step 7](#).

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

7. On any node that is connected to the StorEdge A5x00 disk array, run the luxadm remove command.

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

8. Replace the failed disk drive.

For the procedure on replacing a disk drive, see the *Sun StorEdge A5000 Installation and Service Manual*.

9. On any node that is connected to the StorEdge A5x00 disk array, run the luxadm insert command.

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

If you are inserting a front disk drive, use the `fslotnumber` parameter. If you are inserting a rear disk drive, use the `rslotnumber` parameter.

10. On all other nodes that are attached to the StorEdge A5x00 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 that are connected to the node, the `devfsadm` command can require at least five minutes to complete.

```
# devfsadm
```


11. If you are using Solstice DiskSuite as your volume manager, on one node that is connected to the StorEdge A5x00 disk array, partition the new disk drive by using the partitioning you saved in [Step 6](#).

If you are using VERITAS Volume Manager, go to [Step 12](#).

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

12. One at a time, shut down and reboot the nodes that are connected to the StorEdge A5x00 disk array.

```
# scswitch -S -h nodename  
# shutdown -y -g0 -i6
```

For more information on shutdown procedures, see the *Sun Cluster 3.0 U1 System Administration Guide*.

13. On any of the nodes that are connected to the StorEdge A5x00 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](#).

If the following physical DID is different from the physical DID in [Step 5](#), you successfully replaced the failed disk drive with a new disk drive.

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

15. On all nodes that are connected to the StorEdge A5x00 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 that is 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.

For the procedure on adding a quorum device, see the *Sun Cluster 3.0 U1 System Administration Guide*.

Example—Replacing a StorEdge A5x00 Disk Drive

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

```
# scstat -q
# scdidadm -l d4
4      phys-schost-2:/dev/rdisk/clt32d0 /dev/did/rdisk/d4
# scdidadm -o diskid -l c1t32d0
2000002037000edf
# prtvtoc /dev/rdisk/clt32d0s2 > /usr/tmp/clt32d0.vtoc
# luxadm remove -F /dev/rdisk/clt32d0s2
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" 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/clt32d0.vtoc /dev/rdisk/clt32d0s2
# scswitch -S -h node1
# shutdown -y -g0 -i6
# scdidadm -R d4
# scdidadm -o diskid -l c1t32d0
200000203700bf955
# scdidadm -ui
```

▼ How to Remove a Disk Drive From a StorEdge A5x00 Disk Array in a Running Cluster

Use this procedure to remove a disk drive from a StorEdge A5x00 disk array. “[Example—Removing a StorEdge A5x00 Disk Drive](#)” on page 111 shows you how to apply this procedure. Perform the steps in this procedure in conjunction with the procedures in the *Sun Cluster 3.0 U1 System Administration Guide* and your server hardware manual.

For conceptual information on quorum, quorum devices, global devices, and device IDs, see the *Sun Cluster 3.0 U1 Concepts* document.

1. Determine if the disk drive you are removing is a quorum device.

```
# scstat -q
```

- If the disk drive you are replacing is a quorum device, put the quorum device into maintenance state before you go to [Step 2](#). For the procedure on putting a quorum device into maintenance state, see the *Sun Cluster 3.0 U1 System Administration Guide*.
- If the disk you are replacing is not a quorum device, go to [Step 2](#).

2. If possible, back up the metadevice 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 that is connected to the StorEdge A5x00 disk array, run the `luxadm remove` command.

Physically remove the disk drive, then 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 a StorEdge A5x00 disk array disk drive.

```
# scdidadm -l d4
4      phys-schost-2:/dev/rdisk/clt32d0 /dev/did/rdisk/d4
# luxadm remove -F /dev/rdisk/clt32d0s2

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" 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:
  clt32d0s0
  clt32d0s1
  clt32d0s2
  clt32d0s3
  clt32d0s4
  clt32d0s5
  clt32d0s6
  clt32d0s7
# devfsadm -C
# scdidadm -C
```

▼ How to Add the First StorEdge A5x00 Disk Array to a Running Cluster

Use this procedure to install a StorEdge A5x00 disk array in a running cluster that does not yet have an existing StorEdge A5x00 installed.

If you are installing a StorEdge A5x00 disk array in a running cluster that already has StorEdge A5x00 disk arrays installed and configured with hubs, use the procedure in [“How to Add a StorEdge A5x00 Disk Array to a Running Cluster That Has Existing StorEdge A5x00 Disk Arrays”](#) on page 115.

Perform the steps in this procedure in conjunction with the procedures in the *Sun Cluster 3.0 U1 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 StorEdge A5x00 disk array. This product requires the following packages.**

```
# 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 social Device Driver
system      SUNWluxlx  Sun Enterprise Network Array social Device Driver (64-bit)
system      SUNWluxop  Sun Enterprise Network Array firmware and utilities
```

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

The StorEdge A5x00 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 that is connected to the StorEdge A5x00 disk array.

```
# scswitch -S -h nodename
# shutdown -y -g0 -i0
```

For more information on shutdown procedures, see the *Sun Cluster 3.0 U1 System Administration Guide*.

4. Install host adapters in the node that is to be connected to the StorEdge A5x00 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 StorEdge A5x00 disk array.

For more information, see the *Sun StorEdge A5000 Installation and Service Manual* and the *Sun StorEdge A5000 Configuration Guide*.

FIGURE 6-2 shows a sample StorEdge A5x00 disk array configuration.

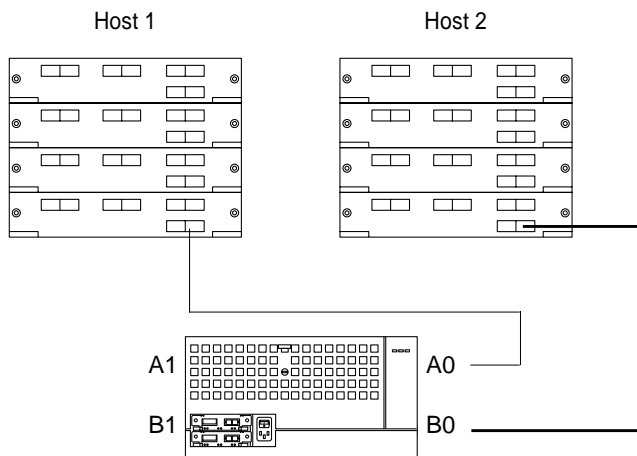


FIGURE 6-2 Sample StorEdge A5x00 Disk Array Configuration

6. Power on and boot the node.

```
# boot -r
```

For the procedures on powering on and booting a node, see the *Sun Cluster 3.0 U1 System Administration Guide*.

7. **Determine if any patches need to be installed on the node(s) that are to be connected to the StorEdge A5x00 disk array.**

For a list of patches specific to Sun Cluster, see the *Sun Cluster 3.0 U1 Release Notes*.

8. **Obtain and install any necessary patches on the nodes that are to be connected to the StorEdge A5x00 disk array.**

For procedures on applying patches, see the *Sun Cluster 3.0 U1 System Administration Guide*.

Note – 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 -g0 -i6
```

For more information on shutdown procedures, see the *Sun Cluster 3.0 U1 System Administration Guide*.

10. **Perform [Step 3](#) through [Step 9](#) for each node that is attached to the StorEdge A5x00 disk array.**

11. **Perform volume management administration to add the disk drives in the StorEdge A5x00 disk array to the volume management configuration.**

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

▼ How to Add a StorEdge A5x00 Disk Array to a Running Cluster That Has Existing StorEdge A5x00 Disk Arrays

Use this procedure to install a StorEdge A5x00 disk array in a running cluster that already has StorEdge A5x00 disk arrays installed and configured with hubs.

If you are installing the first StorEdge A5x00 disk array to a running cluster that does not yet have a StorEdge A5x00 disk array installed, use the procedure in [“How to Add the First StorEdge A5x00 Disk Array to a Running Cluster”](#) on page 112.

Perform the steps in this procedure in conjunction with the procedures in the *Sun Cluster 3.0 U1 System Administration Guide* and your server hardware manual.

1. Configure the new StorEdge A5x00 disk array.

Note – Each disk array in the loop must have a unique box ID. If necessary, use the front-panel module (FPM) to change the box ID for the new StorEdge A5x00 disk array you are adding. For more information about StorEdge A5x00 loops and general configuration, see the *Sun StorEdge A5000 Configuration Guide* and the *Sun StorEdge A5000 Installation and Service Manual*.

2. On both nodes, use the `luxadm insert` command to insert the new disk array to the cluster and to add paths to its disk drives.

```
# luxadm insert
Please hit <RETURN> when you have finished adding Fibre Channel
Enclosure(s)/Device(s):
```

Note – Do not press Return until after you have completed [Step 3](#).

3. Cable the new StorEdge A5x00 disk array to a spare port in the existing hub or host adapter in your cluster.

For cabling instructions and diagrams, see the *Sun StorEdge A5000 Configuration Guide*.

4. After you have finished cabling the new disk array, press Return to complete the `luxadm insert` operation (sample output shown below).

```
Waiting for Loop Initialization to complete...
New Logical Nodes under /dev/dsk and /dev/rdisk :
c4t98d0s0
c4t98d0s1
c4t98d0s2
c4t98d0s3
c4t98d0s4
c4t98d0s5
c4t98d0s6
...
New Logical Nodes under /dev/es:
ses12
ses13
```

5. On both nodes, use the `luxadm probe` command to verify that the new StorEdge A5x00 disk array is recognized by both cluster nodes.

```
# luxadm probe
```

6. On one node, use the `scgdevs` command to update the DID database.

```
# scgdevs
```

▼ How to Replace a StorEdge A5x00 Disk Array in a Running Cluster

Use this procedure to replace a failed StorEdge A5x00 disk array in a running cluster. “[Example—Replacing a StorEdge A5x00 Disk Array](#)” on page 119 shows you how to apply this procedure. This procedure assumes that you are retaining the disk drives.

If you are replacing your disk drives, see “[How to Replace a Disk Drive in a StorEdge A5x00 Disk Array in a Running Cluster](#)” on page 105.

1. **If possible, back up the metadevices or volumes that reside in the StorEdge A5x00 disk array.**

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

2. **Perform volume management administration to remove the StorEdge A5x00 disk array from the configuration.**

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

3. **On all nodes that are connected to the StorEdge A5x00 disk array, run the `luxadm remove` command.**

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

4. **Disconnect the fiber optic cables from the StorEdge A5x00 disk array.**

5. **Power off and disconnect the StorEdge A5x00 disk array from the AC power source.**

For more information, see the *Sun StorEdge A5000 Installation and Service Manual* and the *Sun StorEdge A5000 Configuration Guide*.

6. **Connect the fiber optic cables to the new StorEdge A5x00 disk array.**

7. **Connect the new StorEdge A5x00 disk array to an AC power source.**

8. **One at a time, move the disk drives from the old StorEdge A5x00 disk array to the same slots in the new StorEdge A5x00 disk array.**

9. **Power on the StorEdge A5x00 disk array.**

10. Use the `luxadm insert` command to find the new StorEdge A5x00 disk array.

Repeat this step for each node that is connected to the StorEdge A5x00 disk array.

```
# luxadm insert
```

11. On all nodes that are connected to the new StorEdge A5x00 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 that is 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 StorEdge A5x00 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 a 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@w123456789abcdef03,0:0

/devices/nodes@1/sbus@1f,0/SUNW,socal@1,0/sf@1,0/ses@w123456789abcdef00,0:0

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>
# 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 From a Running Cluster

Use this procedure to remove a StorEdge A5x00 disk array from a cluster.

[“Example—Removing a StorEdge A5x00 Disk Array” on page 121](#) shows you how to apply this procedure. Use the procedures in your server hardware manual to identify the StorEdge A5x00 disk array.

1. **Perform volume management administration to remove the StorEdge A5x00 disk array from the configuration.**

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

2. **On all nodes that are connected to the StorEdge A5x00 disk array, run the `luxadm remove` command.**

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

3. **Remove the StorEdge A5x00 disk array and the fiber optic cables that are connected to the StorEdge A5x00 disk array.**

For more information, see the *Sun StorEdge A5000 Installation and Service Manual*.

4. **On all nodes, remove references to the StorEdge A5x00 disk array.**

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

5. **If necessary, remove any unused 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 a 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@w123456789abcdef03,0:0

/devices/nodes@1/sbus@1f,0/SUNW,socal@1,0/sf@1,0/ses@w123456789abcdef00,0:0

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
# sctidadm -C
```


Installing and Maintaining a Sun StorEdge A3500 / A3500FC System

This chapter describes the procedures for installing, configuring, and maintaining both Sun StorEdge™ A3500 (SCSI-based) and Sun StorEdge A3500FC (Fibre Channel-based) systems in a Sun Cluster environment.

This chapter contains the following procedures:

- [“How to Install a StorEdge A3500/A3500FC System” on page 124](#)
- [“How to Create a LUN” on page 133](#)
- [“How to Delete a LUN” on page 136](#)
- [“How to Reset StorEdge A3500/A3500FC LUN Configuration” on page 139](#)
- [“How to Correct Mismatched DID Numbers” on page 142](#)
- [“How to Add a StorEdge A3500/A3500FC System to a Running Cluster” on page 148](#)
- [“How to Remove a StorEdge A3500/A3500FC System From a Running Cluster” on page 158](#)
- [“How to Replace a Failed Controller or Restore an Offline Controller” on page 161](#)
- [“How to Upgrade Controller Module Firmware in a Running Cluster” on page 163](#)
- [“How to Add a Disk Drive in a Running Cluster” on page 165](#)
- [“How to Replace a Failed Disk Drive in a Running Cluster” on page 166](#)
- [“How to Remove a Disk Drive From a Running Cluster” on page 167](#)
- [“How to Upgrade Disk Drive Firmware in a Running Cluster” on page 167](#)
- [“How to Replace a Host Adapter in a Node \(Connected to a StorEdge A3500 System\)” on page 168](#)
- [“How to Replace a Host Adapter in a Node \(Connected to a StorEdge A3500FC System\)” on page 170](#)

Installing a Sun StorEdge A3500 / A3500FC System

This section describes the procedure for an *initial installation* of a StorEdge A3500/A3500FC system.

▼ How to Install a StorEdge A3500 / A3500FC System

Use this procedure for an initial installation and configuration, *before* installing the Solaris operating environment and Sun Cluster software.

1. Install the host adapters in the nodes that are to be connected to the StorEdge A3500/A3500FC system.

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

2. Cable the StorEdge A3500/A3500FC system:

- See [FIGURE 7-1](#) for a sample StorEdge A3500 system cabling.
- See [FIGURE 7-2](#) for a sample StorEdge A3500FC system cabling.

For more sample configurations, see the *Sun StorEdge A3500/A3500FC Hardware Configuration Guide*.

For the procedure on installing the cables, see the *Sun StorEdge A3500/A3500FC Controller Module Guide*.

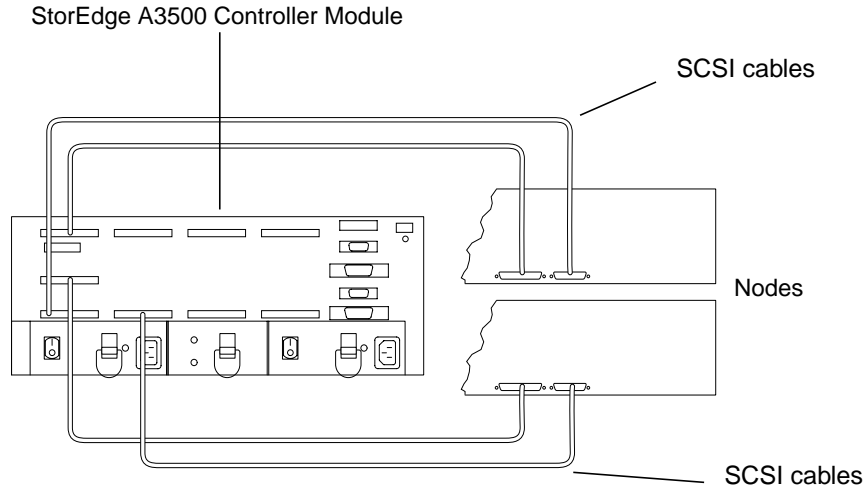


FIGURE 7-1 Sample StorEdge A3500 System Cabling

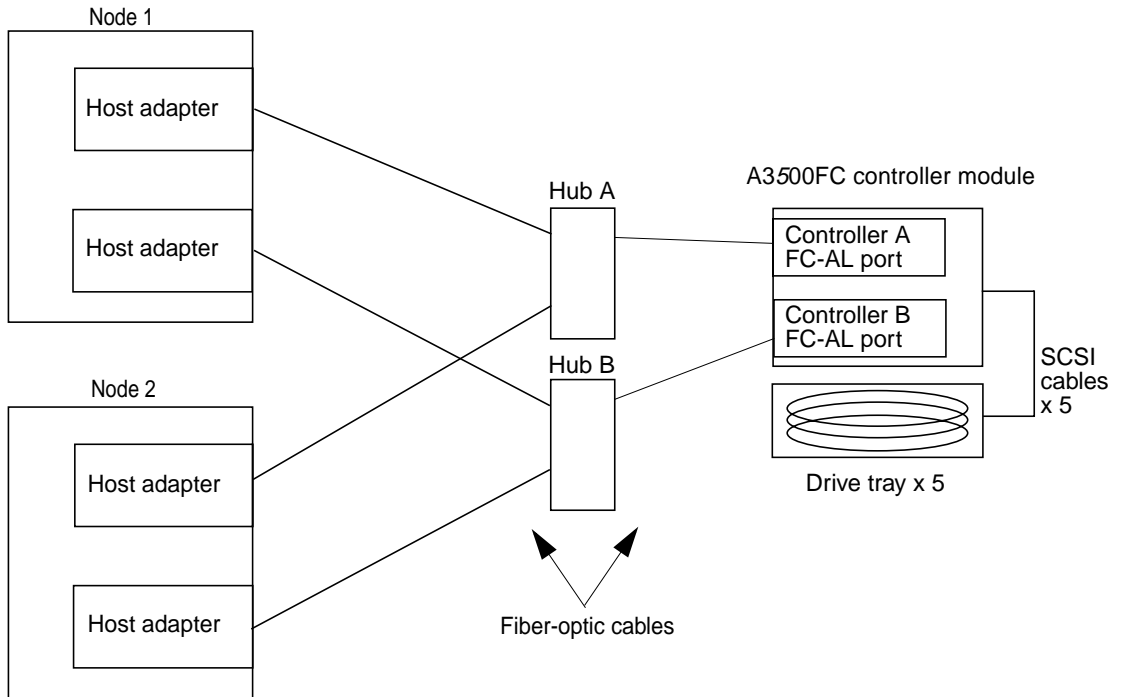


FIGURE 7-2 Sample StorEdge A3500FC System Cabling

3. Depending on which type of controller module you are installing:

- If you are installing a StorEdge A3500 controller module, go to [Step 4](#).
- If you are installing a StorEdge A3500FC controller module, set the loop ID of the controller module by installing jumpers to the appropriate pins on the rear of the controller module.

For diagrams and information about setting FC-AL ID settings, see the *Sun StorEdge A3500/A3500FC Controller Module Guide*.

4. Power on the StorEdge A3500/A3500FC system and cluster nodes.

Note – For StorEdge A3500 controller modules only: When you power on the nodes, do not allow them to boot. If necessary, halt the nodes so that you can perform OpenBoot™ PROM (OBP) Monitor tasks at the ok prompt.

For the procedure on powering on the StorEdge A3500/A3500FC system, see the *Sun StorEdge A3500/A3500FC Controller Module Guide*.

5. Depending on which type of controller module you are installing:

- For a StorEdge A3500FC controller module, go to [Step 13](#).
- For a StorEdge A3500 controller module, find the paths to the host adapters in the *first* node:

```
{0} ok show-disks
...
b) /sbus@6,0/QLGC,isp@2,10000/sd
...
d) /sbus@2,0/QLGC,isp@2,10000/sd
...
```

Note – Use this information to change the SCSI addresses of the host adapters in the `nvrarc` script in [Step 6](#), but do not include the `sd` directories in the device paths.

6. Edit the `nvrामrc` script to change the `scsi-initiator-id` for the host adapters on the *first* node.

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 node that has a host adapter with SCSI address 7 as the “second node.”

To avoid conflicts, you must change the `scsi-initiator-id` of the remaining host adapter in the SCSI chain to an available SCSI address. This procedure refers to the node that has a host adapter with an available SCSI address as the “first node.”

For a partial list of `nvrामrc` editor and `nvedit` keystroke commands, see [Appendix B](#) of this guide. For a full list of commands, see the *OpenBoot 3.x Command Reference Manual*.

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

Note – Insert exactly one space after the first quotation mark and before `scsi-initiator-id`.

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

7. Store the changes.

The changes you make through the `nvedit` command are recorded on a temporary copy of the `nvrामrc` 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
```

- To discard the changes, type:

```
{0} ok nvquit
```

8. Verify the contents of the `nvrामrc` script you created in [Step 6](#), as shown in the following example.

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

```
{0} ok printenv nvrामrc  
nvrामrc =  
    probe-all  
    cd /sbus@6,0/QLGC,isp@2,10000  
    6 " scsi-initiator-id" integer-property  
    device-end  
    cd /sbus@2,0/QLGC,isp@2,10000  
    6 " scsi-initiator-id" integer-property  
    device-end  
    install-console  
    banner
```

9. Set the parameter to instruct the OpenBoot PROM Monitor to use the `nvrामrc` script:

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

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. Select each host adapter's device tree node, then display the node's properties to confirm that the `scsi-initiator-id` for each host adapter is set to 7:

```
{0} ok show-disks
...
b) /sbus@6,0/QLGC,isp@2,10000/sd
...
d) /sbus@2,0/QLGC,isp@2,10000/sd
...
{0} ok cd /sbus@6,0/QLGC,isp@2,10000
{0} ok .properties
scsi-initiator-id      00000007
```

11. Install the Solaris operating environment, then apply any required Solaris patches.

For the procedure on installing the Solaris operating environment, see the *Sun Cluster 3.0 U1 Installation Guide*. For the location of patches and installation instructions, see the *Sun Cluster 3.0 U1 Release Notes*.

12. Read the following two conditions carefully to determine whether you must reboot the cluster nodes now:

- If you are using a version of RAID Manager later than 6.22 or you are using a version of the Solaris operating environment earlier than Solaris 8 Update 4, go to [Step 13](#).
- If you are using RAID Manager 6.22 *and* the Solaris 8 Update 4 or later operating environment, reboot both cluster nodes now.

```
# reboot
```

13. Install the RAID Manager software.

For the procedure on installing the RAID Manager software, see the *Sun StorEdge RAID Manager Installation and Support Guide*.

Note – RAID Manager 6.22 or a compatible version is required for clustering with Sun Cluster 3.0.

Note – For the most current list of software, firmware, and patches that are required for the StorEdge A3x00/A3500FC controller module, refer to EarlyNotifier 20029, “A1000/A3x00/A3500FC Software/Firmware Configuration Matrix.” This document is available online to Sun service providers and to customers with SunSolve service contracts at the SunSolve site: <http://sunsolve.sun.com>.

14. Install any StorEdge A3500/A3500FC controller module or RAID Manager patches.

For more information, see the *Sun StorEdge RAID Manager Release Notes*.

15. Check the StorEdge A3500/A3500FC controller module NVSRAM file revision, and if necessary, install the most recent revision.

For the NVSRAM file revision number and boot level, see the *Sun StorEdge RAID Manager Release Notes*. For the procedure on upgrading the NVSRAM file, see the *Sun StorEdge RAID Manager User’s Guide*.

16. Check the StorEdge A3500/A3500FC controller module firmware revision, and if necessary, install the most recent revision.

For the firmware revision number and boot level, see the *Sun StorEdge RAID Manager Release Notes*. For the procedure on upgrading the firmware, see the *Sun StorEdge RAID Manager User’s Guide*.

17. Set the Rdac parameters in the `/etc/osa/rmparams` file:

```
Rdac_RetryCount=1
Rdac_NoAltOffline=TRUE
```

18. Verify that the controller module is set to active/active mode (if it is not, set it to active/active).

For more information on controller modes, see the *Sun StorEdge RAID Manager Installation and Support Guide* and the *Sun StorEdge RAID Manager User’s Guide*.

19. **Set up the StorEdge A3500/A3500FC controller module with logical unit numbers (LUNs) and hot spares.**

For the procedure on setting up the StorEdge A3500/A3500FC controller module with LUNs and hot spares, see the *Sun StorEdge RAID Manager User's Guide*.

Note – Use the `format` command to verify Solaris logical device names.

20. **Copy the `/etc/raid/rdac_address` file from the node on which you created the LUNs to the other node to ensure consistency across both nodes.**
21. **Ensure that the new logical name for the LUN you created in [Step 19](#) appears in the `/dev/rdisk` directory on both nodes by running the `hot_add` command on both nodes:**

```
# /etc/raid/bin/hot_add
```

Where to Go From Here

To continue with Sun Cluster software and data services installation tasks, see the *Sun Cluster 3.0 U1 Installation Guide* and the *Sun Cluster 3.0 U1 Data Services Installation and Configuration Guide*.

Configuring a Sun StorEdge A3500/A3500FC System

This section describes the procedures for configuring a StorEdge A3500/A3500FC system *after* installing Sun Cluster software. [TABLE 7-1](#) lists these procedures.

Configuring a StorEdge A3500/A3500FC system *before* installing Sun Cluster software is the same as doing so in a non-cluster environment. For procedures on configuring StorEdge A3500/A3500FC systems before installing Sun Cluster, see the *Sun StorEdge RAID Manager User's Guide*.

TABLE 7-1 Task Map: Configuring StorEdge A3500/A3500FC Disk Drives

Task	For Instructions, Go To
Create a logical unit number (LUN).	"How to Create a LUN" on page 133
Remove a LUN.	"How to Delete a LUN" on page 136
Reset the StorEdge A3500/A3500FC configuration.	"How to Reset StorEdge A3500/A3500FC LUN Configuration" on page 139
Rebalance running LUNs. Follow the same procedure that is used in a non-cluster environment.	<i>Sun StorEdge RAID Manager User's Guide</i> <i>Sun StorEdge RAID Manager Release Notes</i>
Create a hot spare. Follow the same procedure that is used in a non-cluster environment.	<i>Sun StorEdge RAID Manager User's Guide</i> <i>Sun StorEdge RAID Manager Release Notes</i>
Delete a hot spare. Follow the same procedure that is used in a non-cluster environment.	<i>Sun StorEdge RAID Manager User's Guide</i> <i>Sun StorEdge RAID Manager Release Notes</i>
Increase the size of a drive group. Follow the same procedure that is used in a non-cluster environment.	<i>Sun StorEdge RAID Manager User's Guide</i> <i>Sun StorEdge RAID Manager Release Notes</i>

▼ How to Create a LUN

Use this procedure to create a logical unit number (LUN) from unassigned disk drives or remaining capacity. See the *Sun StorEdge RAID Manager Release Notes* for the latest information about LUN administration.

This product supports the use of hardware RAID and host-based software RAID. For host-based software RAID, this product supports RAID levels 0+1 and 1+0.

Note – You must use hardware RAID for Oracle Parallel Server (OPS) data stored on the StorEdge A3500/A3500FC arrays. Do not place OPS data under volume management control. You must place all non-OPS data that is stored on the arrays under volume management control. Use either hardware RAID, host-based software RAID, or both types of RAID to manage your non-OPS data.

Hardware RAID uses the StorEdge A3500/A3500FC system's hardware redundancy to ensure that independent hardware failures do not impact data availability. By mirroring across separate arrays, host-based software RAID ensures that independent hardware failures do not impact data availability when an entire array is offline. Although you can use hardware RAID and host-based software RAID concurrently, you need only one RAID solution to maintain a high degree of data availability.

Note – When you use host-based software RAID with hardware RAID, the hardware RAID levels you use affect the hardware maintenance procedures because they affect volume management administration.

If you use hardware RAID level 1, 3, or 5, you can perform most maintenance procedures in [“Maintaining a StorEdge A3500/A3500FC System” on page 144](#) without volume management disruptions.

If you use hardware RAID level 0, some maintenance procedures in [“Maintaining a StorEdge A3500/A3500FC System” on page 144](#) require additional volume management administration because the availability of the LUNs is impacted.

1. With all cluster nodes booted and attached to the StorEdge A3500/A3500FC system, create the LUN on one node.

Shortly after the LUN formatting completes, a logical name for the new LUN appears in `/dev/rdisk` on all cluster nodes that are attached to the StorEdge A3500/A3500FC system.

For the procedure on creating a LUN, see the *Sun StorEdge RAID Manager User's Guide*.

If the following warning message is displayed, ignore it and continue with the next step:

```
scsi: WARNING:
/sbus@40,0/SUNW,socal@0,0/sf@1,0/ssd@w200200a0b80740db,4 (ssd0):
corrupt label - wrong magic number
```

Note – Use the `format(1M)` command to verify Solaris logical device names.

2. Copy the `/etc/raid/rdac_address` file from the node on which you created the LUN to the other node to ensure consistency across both nodes.
3. Ensure that the new logical name for the LUN you created in [Step 1](#) appears in the `/dev/rdisk` directory on both nodes by running the `hot_add` command on both nodes:

```
# /etc/raid/bin/hot_add
```

4. On one node, update the global device namespace:

```
# scgdevs
```

5. Use the `scdidadm` command to verify that the DID numbers for the LUNs are the same on both nodes. In the sample output that follows, the DID numbers are different:

```
# scdidadm -L
...
33          e07a:/dev/rdsk/c1t4d2          /dev/did/rdsk/d33
33          e07c:/dev/rdsk/c0t4d2          /dev/did/rdsk/d33
```

6. Are the DID numbers you received from running the `scdidadm` command in [Step 5](#) the same for both your nodes?
 - If the DID numbers are the same, go to [Step 7](#).
 - If the DID numbers are different, perform the procedure in “[How to Correct Mismatched DID Numbers](#)” on [page 142](#) before you continue with [Step 7](#) of this procedure.
7. If you want a volume manager to manage the new LUN you created in [Step 1](#), run the appropriate Solstice DiskSuite or 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.

Note – Do not configure StorEdge A3500/A3500FC LUNs as quorum devices; the use of StorEdge A3500/A3500FC LUNs as quorum devices is not supported.

▼ How to Delete a LUN

Use this procedure to delete a LUN(s). See the *Sun StorEdge RAID Manager Release Notes* for the latest information about LUN administration.



Caution – This procedure removes all data on the LUN you delete.



Caution – Do not delete LUN 0.

1. From one node that is connected to the StorEdge A3500/A3500FC system, use the `format` command to determine the paths to the LUN you are deleting (sample output follows):

```
# format
AVAILABLE DISK SELECTIONS:
0. c0t5d0 <SYMBIOS-StorEdgeA3500FCr-0301 cyl13 alt2 hd64 sec64>
/pseudo/rdnexus@0/rdriver@5,0
1. c0t5d1 <SYMBIOS-StorEdgeA3500FCr-0301 cyl12025 alt2 hd64 sec64>
/pseudo/rdnexus@0/rdriver@5,1
```

2. Does a volume manager manage the LUN you are deleting?

- If not, go to [Step 3](#).
- If a volume manager does manage the LUN, 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. See the following paragraph for additional VERITAS Volume Manager commands that are required.

LUNs that were managed by VERITAS Volume Manager must be completely removed from VERITAS Volume Manager control before you can delete them. To remove the LUNs, after you delete the LUN from any disk group, use the following commands:

```
# vxdisk offline cNtXdY
# vxdisk rm cNtXdY
```

3. From one node, delete the LUN.

For the procedure on deleting a LUN, see the *Sun StorEdge RAID Manager User's Guide*.

4. Remove the paths to the LUN(s) you are deleting:

```
# rm /dev/rdisk/cNtXdY*
# rm /dev/dsk/cNtXdY*

# rm /dev/osa/dev/dsk/cNtXdY*
# rm /dev/osa/dev/rdisk/cNtXdY*
```

5. Use the `lad` command to determine the *alternate* paths to the LUN(s) you are deleting.

The RAID Manager software creates two paths to the LUN in the `/dev/osa/dev/rdisk` directory. Substitute the `cNtXdY` number from the other controller module in the disk array to determine the alternate path.

For example, with this configuration:

```
# lad
c0t5d0 1T93600714 LUNS: 0 1
c1t4d0 1T93500595 LUNS: 2
```

The alternate paths would be:

```
/dev/osa/dev/dsk/c1t4d1*
/dev/osa/dev/rdisk/c1t4d1*
```

6. Remove the *alternate* paths to the LUN(s) you are deleting:

```
# rm /dev/osa/dev/dsk/cNtXdY*
# rm /dev/osa/dev/rdisk/cNtXdY*
```

7. On both nodes, remove all obsolete device IDs (DID)s:

```
# scdidadm -C
```

8. Switch resources and device groups off the node:

```
# scswitch -Sh nodename
```

9. Shut down the node:

```
# shutdown -y -g0 -i0
```

10. Boot the node and wait for it to rejoin the cluster:

```
# boot -r
```

11. Repeat [Step 3](#) through [Step 10](#) on the other node that is attached to the StorEdge A3500/A3500FC system.

▼ How to Reset StorEdge A3500/A3500FC LUN Configuration

Use this procedure to reset StorEdge A3500/A3500FC LUN configuration.



Caution – Resetting LUN configuration results in a new DID number being assigned to LUN 0. This is because the software assigns a new worldwide number (WWN) to the new LUN.

1. From one node that is connected to the StorEdge A3500/A3500FC system, use the `format` command to determine the paths to the LUN(s) you are resetting, as shown in the following example (sample output shown below):

```
# format
AVAILABLE DISK SELECTIONS:
0. c0t5d0 <SYMBIOS-StorEdgeA3500FCr-0301 cyl13 alt2 hd64 sec64>
  /pseudo/rdnexus@0/rdriver@5,0
1. c0t5d1 <SYMBIOS-StorEdgeA3500FCr-0301 cyl12025 alt2 hd64 sec64>
  /pseudo/rdnexus@0/rdriver@5,1
```

2. Does a volume manager manage the LUN(s) on the controller module you are resetting?

- If not, go to [Step 3](#).
- If a volume manager does manage the LUN, 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. See the following paragraph for additional VERITAS Volume Manager commands that are required.

LUNs that were managed by VERITAS Volume Manager must be completely removed from VERITAS Volume Manager control before you can delete them. To remove the LUNs, after you delete the LUN from any disk group, use the following commands:

```
# vxdisk offline cNtXdY
# vxdisk rm cNtXdY
```

3. On one node, reset the LUN configuration.

For the procedure for resetting StorEdge A3500/A3500FC LUN configuration, see the *Sun StorEdge RAID Manager User's Guide*.

Note – Use the `format` command to verify Solaris logical device names.

4. Set the controller module back to active/active mode (it was set to active/passive when reset).

For more information on controller modes, see the *Sun StorEdge RAID Manager Installation and Support Guide* and the *Sun StorEdge RAID Manager User's Guide*.

5. By using the `format` command, label the new LUN 0.

6. Remove the paths to the old LUN(s) you reset:

```
# rm /dev/rdisk/cNtXdY*
# rm /dev/dsk/cNtXdY*

# rm /dev/osa/dev/dsk/cNtXdY*
# rm /dev/osa/dev/rdisk/cNtXdY*
```

7. Use the `lad` command to determine the alternate paths to the old LUN(s) you reset.

The RAID Manager software creates two paths to the LUN in the `/dev/osa/dev/rdisk` directory. Substitute the `cNtXdY` number from the other controller module in the disk array to determine the alternate path.

For example, with this configuration:

```
# lad
c0t5d0 1T93600714 LUNS: 0 1
c1t4d0 1T93500595 LUNS: 2
```

The alternate paths would be as follows.

```
/dev/osa/dev/dsk/c1t4d1*
/dev/osa/dev/rdisk/c1t4d1*
```

8. Remove the *alternate* paths to the old LUN(s) you reset:

```
# rm /dev/osa/dev/dsk/cNtXdY*
# rm /dev/osa/dev/rdisk/cNtXdY*
```

9. Update device namespaces on both nodes:

```
devfsadm -C
```

10. Remove all obsolete DIDs on both nodes:

```
# scdidadm -C
```

11. Switch resources and device groups off the node:

```
# scswitch -Sh nodename
```

12. Shut down the node:

```
# shutdown -y -g0 -i0
```

13. Boot the node and wait for it to rejoin the cluster:

```
# boot -r
```

If the following error message appears, ignore it and continue with the next step. The DID will be updated when the procedure is complete.

```
device id for '/dev/rdisk/c0t5d0' does not match physical disk's id.
```

14. After the node has rebooted and joined the cluster, repeat [Step 6](#) through [Step 13](#) on the other node that is attached to the StorEdge A3500/A3500FC system.

The DID number for the original LUN 0 is removed and a new DID is assigned to LUN 0.

▼ How to Correct Mismatched DID Numbers

Use this section to correct mismatched device ID (DID) numbers that might appear during the creation of A3500/A3500FC LUNs. You correct the mismatch by deleting Solaris and Sun Cluster paths to the LUNs that have DID numbers that are different. After rebooting, the paths are corrected.

Note – Use this procedure only if you are directed to do so from [“How to Create a LUN” on page 133](#).

1. From one node that is connected to the StorEdge A3500/A3500FC system, use the `format` command to determine the paths to the LUN(s) that have different DID numbers:

```
# format
```

2. Remove the paths to the LUN(s) that have different DID numbers:

```
# rm /dev/rdisk/cNtXdY*
# rm /dev/dsk/cNtXdY*

# rm /dev/osa/dev/dsk/cNtXdY*
# rm /dev/osa/dev/rdisk/cNtXdY*
```

3. Use the `lad` command to determine the *alternate* paths to the LUN(s) that have different DID numbers.

The RAID Manager software creates two paths to the LUN in the `/dev/osa/dev/rdisk` directory. Substitute the `cNtXdY` number from the other controller module in the disk array to determine the alternate path.

For example, with this configuration:

```
# lad
c0t5d0 1T93600714 LUNS: 0 1
c1t4d0 1T93500595 LUNS: 2
```

The alternate paths would be as follows.

```
/dev/osa/dev/dsk/c1t4d1*
/dev/osa/dev/rdisk/c1t4d1*
```

4. Remove the *alternate* paths to the LUN(s) that have different DID numbers:

```
# rm /dev/osa/dev/dsk/cNtXdY*
# rm /dev/osa/dev/rdisk/cNtXdY*
```

5. On both nodes, remove all obsolete DIDs:

```
# scdidadm -C
```

6. Switch resources and device groups off the node:

```
# scswitch -Sh nodename
```

7. Shut down the node:

```
# shutdown -y -g0 -i0
```

8. Boot the node and wait for it to rejoin the cluster:

```
# boot -r
```

9. Repeat [Step 1](#) through [Step 8](#) on the other node that is attached to the StorEdge A3500/A3500FC system.
10. Return to [“How to Create a LUN” on page 133](#).

Maintaining a StorEdge A3500/A3500FC System

This section contains the procedures for maintaining a StorEdge A3500/A3500FC system in a Sun Cluster environment. Some maintenance tasks listed in [TABLE 7-2](#) are performed the same as in a non-cluster environment, so the task's procedures are referenced rather than contained in this section. [TABLE 7-2](#) lists the procedures for maintaining the StorEdge A3500/A3500FC system.

TABLE 7-2 Tasks: Maintaining a StorEdge A3500/A3500FC System

Task	For Instructions, Go To
A3500/A3500FC system/controller module procedures:	
Add a StorEdge A3500/A3500FC system to a running cluster.	"How to Add a StorEdge A3500/A3500FC System to a Running Cluster" on page 148
Remove a StorEdge A3500/A3500FC system from a running cluster.	"How to Remove a StorEdge A3500/A3500FC System From a Running Cluster" on page 158
Replace a failed StorEdge A3500/A3500FC controller module or restore an offline controller module.	"How to Replace a Failed Controller or Restore an Offline Controller" on page 161
Upgrade StorEdge A3500/A3500FC controller module firmware and NVSRAM file.	"How to Upgrade Controller Module Firmware in a Running Cluster" on page 163
Replace a power cord to a StorEdge A3500/A3500FC controller module. Shut down the cluster, then follow the same procedure used in a non-cluster environment.	Sun Cluster 3.0 U1 System Administration Guide for procedures on shutting down a cluster Sun StorEdge A3500/A3500FC Controller Module Guide for replacement procedures
Replace a StorEdge A3500/A3500FC controller module fan canister. Follow the same procedure that is used in a non-cluster environment.	Sun StorEdge A3500/A3500FC Controller Module Guide

TABLE 7-2 Tasks: Maintaining a StorEdge A3500/A3500FC System (*Continued*)

Task	For Instructions, Go To
<p>Replace the StorEdge A3500/A3500FC controller module card cage.</p> <p>Shut down the cluster, then follow the same procedure that is used in a non-cluster environment.</p>	<p><i>Sun Cluster 3.0 U1 System Administration Guide</i> for procedures on shutting down a cluster</p> <p><i>Sun StorEdge A3500/A3500FC Controller Module Guide</i> for replacement procedures</p>
<p>Replace the entire StorEdge A3500/A3500FC controller module assembly.</p> <p>Shut down the cluster, then follow the same procedure that is used in a non-cluster environment.</p>	<p><i>Sun Cluster 3.0 U1 System Administration Guide</i> for procedures on shutting down a cluster</p> <p><i>Sun StorEdge A3500/A3500FC Controller Module Guide</i> for replacement procedures</p>
Cable/hub/connector procedures:	
<p>Replace a SCSI cable from the controller module to the disk array.</p> <p>Follow the same procedure that is used in a non-cluster environment.</p>	<p><i>Sun StorEdge A3500/A3500FC Controller Module Guide</i></p> <p><i>Sun StorEdge RAID Manager User's Guide</i> <i>Sun StorEdge RAID Manager Release Notes</i></p>
<p>Replace a StorEdge A3500-to-host SCSI cable.</p> <p>Follow the same procedure that is used in a non-cluster environment.</p>	<p><i>Sun StorEdge A3500/A3500FC Controller Module Guide</i></p> <p><i>Sun StorEdge RAID Manager User's Guide</i> <i>Sun StorEdge RAID Manager Release Notes</i></p>
<p>Replace a StorEdge A3500FC-to-host (or hub) fiber-optic cable.</p> <p>Follow the same procedure that is used in a non-cluster environment.</p>	<p><i>Sun StorEdge A3500/A3500FC Controller Module Guide</i></p> <p><i>Sun StorEdge RAID Manager User's Guide</i> <i>Sun StorEdge RAID Manager Release Notes</i></p>
<p>Replace a StorEdge FC100 hub.</p> <p>Follow the same procedure that is used in a non-cluster environment.</p>	<p><i>Sun StorEdge A3500/A3500FC Controller Module Guide</i></p> <p><i>Sun StorEdge FC-100 Hub Installation and Service Manual</i></p>

TABLE 7-2 Tasks: Maintaining a StorEdge A3500/A3500FC System (Continued)

Task	For Instructions, Go To
<p>Replace a StorEdge FC100 hub gigabit interface converter (GBIC) that connects cables to the host or hub.</p> <p>Follow the same procedure that is used in a non-cluster environment.</p>	<p><i>Sun StorEdge A3500/A3500FC Controller Module Guide</i> <i>Sun StorEdge FC-100 Hub Installation and Service Manual</i></p>
<p>Replace a GBIC on a node.</p> <p>Follow the same procedure that is used in a non-cluster environment.</p>	<p><i>Sun StorEdge A3500/A3500FC Controller Module Guide</i></p>
<p>Cabinet/power subassembly procedures:</p>	
<p>Replace the power supply fan canister.</p> <p>Follow the same procedure that is used in a non-cluster environment.</p>	<p><i>Sun StorEdge A3500/A3500FC Controller Module Guide</i></p>
<p>Replace a DC-power or battery harness.</p> <p>Shut down the cluster, then follow the same procedure that is used in a non-cluster environment.</p>	<p><i>Sun Cluster 3.0 U1 System Administration Guide</i> for procedures on shutting down a cluster <i>Sun StorEdge A3500/A3500FC Controller Module Guide</i> for replacement procedures</p>
<p>Replace the battery unit.</p> <p>Shut down the cluster, then follow the same procedure that is used in a non-cluster environment.</p>	<p><i>Sun Cluster 3.0 U1 System Administration Guide</i> for procedures on shutting down a cluster <i>Sun StorEdge A3500/A3500FC Controller Module Guide</i> for replacement procedures</p>
<p>Replace the power supply housing.</p> <p>Shut down the cluster, then follow the same procedure that is used in a non-cluster environment.</p>	<p><i>Sun Cluster 3.0 U1 System Administration Guide</i> for procedures on shutting down a cluster <i>Sun StorEdge A3500/A3500FC Controller Module Guide</i></p>
<p>Replace a power cord to the cabinet power distribution unit.</p> <p>Shut down the cluster, then follow the same procedure that is used in a non-cluster environment.</p>	<p><i>Sun Cluster 3.0 U1 System Administration Guide</i> for procedures on shutting down a cluster <i>Sun StorEdge Expansion Cabinet Installation and Service Manual</i> for replacement procedures.</p>

TABLE 7-2 Tasks: Maintaining a StorEdge A3500/A3500FC System (*Continued*)

Task	For Instructions, Go To
D1000 Disk Array Procedures:	
Add a disk drive.	“How to Add a Disk Drive in a Running Cluster” on page 165
Replace a disk drive.	“How to Replace a Failed Disk Drive in a Running Cluster” on page 166
Remove a disk drive.	“How to Remove a Disk Drive From a Running Cluster” on page 167
Upgrade disk drive firmware.	“How to Upgrade Disk Drive Firmware in a Running Cluster” on page 167
<p>Replace a power cord to a StorEdge D1000 disk array.</p> <p>Shut down the cluster, then follow the same procedure that is used in a non-cluster environment.</p>	<p><i>Sun Cluster 3.0 U1 System Administration Guide</i> for procedures on shutting down a cluster</p> <p><i>Sun StorEdge A1000 and D1000 Installation, Operations, and Service Manual</i> for replacement procedures</p>
Node/host adapter procedures:	
<p>Replace a host adapter in a node.</p> <p>Follow the procedure for your type of controller module (StorEdge A3500 or StorEdge A3500FC).</p>	<p>“How to Replace a Host Adapter in a Node (Connected to a StorEdge A3500 System)” on page 168</p> <p>“How to Replace a Host Adapter in a Node (Connected to a StorEdge A3500FC System)” on page 170</p>

▼ How to Add a StorEdge A3500/A3500FC System to a Running Cluster

Use this procedure to add a StorEdge A3500/A3500FC system to a running cluster.

1. Install the RAID Manager software.

For the procedure on installing RAID Manager software, see the *Sun StorEdge RAID Manager Installation and Support Guide*.

Note – RAID Manager 6.22 or a compatible version is required for clustering with Sun Cluster 3.0.

Note – For the most current list of software, firmware, and patches that are required for the StorEdge A3x00/A3500FC controller module, refer to EarlyNotifier 20029, “A1000/A3x00/A3500FC Software/Firmware Configuration Matrix.” This document is available online to Sun service providers and to customers with SunSolve service contracts at the SunSolve site: <http://sunsolve.sun.com>.

2. Install any StorEdge A3500/A3500FC system patches.

For the location of patches and installation instructions, see the *Sun Cluster 3.0 U1 Release Notes*.

3. Set the Rdac parameters in the `/etc/osa/rmparams` file:

```
Rdac_RetryCount=1
Rdac_NoAltOffline=TRUE
```

4. Power on the StorEdge A3500/A3500FC system.

For the procedure on powering on the StorEdge A3500/A3500FC system, see the *Sun StorEdge A3500/A3500FC Controller Module Guide*.

5. Depending on which type of system you are adding:

- If you are adding a StorEdge A3500 system, go to [Step 6](#).
- If you are adding a StorEdge A3500FC system, set the loop ID of the controller module by installing jumpers to the appropriate pins on the rear of the controller module.

For diagrams and information about setting FC-AL ID settings, see the *Sun StorEdge A3500/A3500FC Controller Module Guide*.

6. Are you installing new host adapters to your nodes for connection to the StorEdge A3500/A3500FC system?

- If not, go to [Step 8](#).
- If you are installing new host adapters, shut down and power off the *first* node.

```
# scswitch -S -h nodename  
# shutdown -y -g0 -i0
```

For the full procedure on shutting down and powering off a node, see the *Sun Cluster 3.0 U1 System Administration Guide*.

7. Install the host adapters in the first node.

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

8. Cable the StorEdge A3500/A3500FC system to the first node. Depending on which type of system you are adding:

- If you are adding a StorEdge A3500 system, connect the differential SCSI cable between the node and the controller module as shown in [FIGURE 7-3](#). 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.
- If you are installing a StorEdge A3500FC system, see [FIGURE 7-4](#) for a sample StorEdge A3500FC cabling connection. The example shows the first node that is connected to a StorEdge A3500FC controller module.

For more sample configurations, see the *Sun StorEdge A3500/A3500FC Hardware Configuration Guide*.

For the procedure on installing the cables, see the *Sun StorEdge A3500/A3500FC Controller Module Guide*.




FIGURE 7-3 Sample StorEdge A3500 Cabling

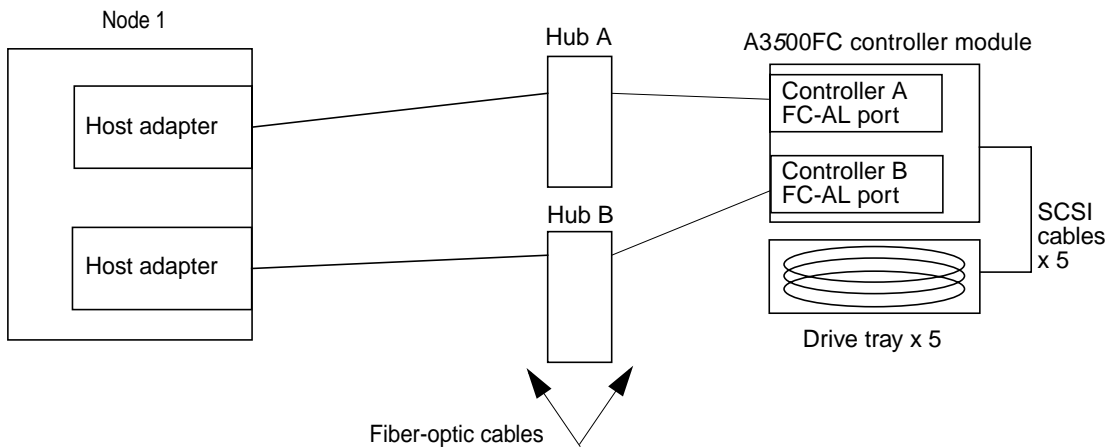


FIGURE 7-4 Sample StorEdge A3500FC Cabling (1st Node Attached)

9. Did you power off the first node to install a host adapter?

- If not, go to [Step 10](#).
- If you did power off the first node, power it and the StorEdge A3500 system on, *but do not allow the node to boot*. If necessary, halt the system to continue with OpenBoot PROM (OBP) Monitor tasks.

10. Depending on which type of controller module you are adding, do the following:

- If you are installing a StorEdge A3500FC controller module, go to [Step 15](#).
- If you are adding a StorEdge A3500 controller module, find the paths to the SCSI host adapters.

```
{0} ok show-disks
...
b) /sbus@6,0/QLGC,isp@2,10000/sd
...
d) /sbus@2,0/QLGC,isp@2,10000/sd
...
```

Identify and record the two controllers that are to be connected to the disk arrays, and record these paths. Use this information to change the SCSI addresses of these controllers in the `nvrामrc` script in [Step 11](#). Do not include the `sd` directories in the device paths.

11. Edit the `nvrामrc` script to change the `scsi-initiator-id` for the host adapters on the *first* node.

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 that has SCSI address 7 as the host adapter on the “second node.”

To avoid conflicts, 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 that has an available SCSI address as the host adapter on the “first node.”

For a partial list of `nvrामrc` editor and `nvedit` keystroke commands, see [Appendix B](#) of this guide. For a full list of commands, see the *OpenBoot 3.x Command Reference Manual*.

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

Note – Insert exactly one space after the quotation mark and before `scsi-initiator-id`.

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

12. Store the changes.

The changes you make through the `nvedit` command are recorded on a temporary copy of the `nvrामrc` 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
```

13. Verify the contents of the `nvrामrc` script you created in [Step 11](#), as shown in the following example.

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

```
{0} ok printenv nvrामrc
nvrामrc =                probe-all
                        cd /sbus@6,0/QLGC,isp@2,10000
                        6 "scsi-initiator-id" integer-property
                        device-end
                        cd /sbus@2,0/QLGC,isp@2,10000
                        6 "scsi-initiator-id" integer-property
                        device-end
                        install-console
                        banner
{0} ok
```

14. Instruct the OpenBoot PROM Monitor to use the `nvrामrc` script:

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

15. Did you power off the first node to install a host adapter?

- If not, go to [Step 21](#).
- If you powered off the first node, boot it now and wait for it to join the cluster.

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

For more information on booting nodes, see the *Sun Cluster 3.0 U1 System Administration Guide*.

16. Are you installing new host adapters to the second node for connection to the StorEdge A3500/A3500FC system?

- If not, go to [Step 21](#).
- If you are installing new host adapters, shut down and power off the *second* node.

```
# scswitch -S -h nodename  
# shutdown -y -g0 -i0
```

For the procedure on shutting down and powering off a node, see the *Sun Cluster 3.0 U1 System Administration Guide*.

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. Cable the StorEdge A3500/A3500FC system to your node. Depending on which type of controller module you are adding, do the following:

- If you are adding a StorEdge A3500 controller module, connect the differential SCSI cable between the node and the controller module as shown in [FIGURE 7-3](#). 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.
- If you are installing a StorEdge A3500FC controller module, see [FIGURE 7-5](#) for a sample StorEdge A3500FC cabling connection. The example shows two nodes that are connected to a StorEdge A3500FC controller module.

For more sample configurations, see the *Sun StorEdge A3500/A3500FC Hardware Configuration Guide*.

For the procedure on installing the cables, see the *Sun StorEdge A3500/A3500FC Controller Module Guide*.




FIGURE 7-5 Sample StorEdge A3500FC Cabling (2nd Node Attached)

19. Did you power off the second node to install a host adapter?

- If not, go to [Step 21](#).
- If you did power off the second node, power it and the StorEdge A3500/A3500FC system on, *but do not allow the node to boot*. If necessary, halt the system to continue with OpenBoot PROM (OBP) Monitor tasks.

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

If the node does not recognize the new hardware, check all hardware connections and repeat installation steps you performed in [Step 17](#).

```
{0} ok show-disks
...
b) /sbus@6,0/QLGC,isp@2,10000/sd
...
d) /sbus@2,0/QLGC,isp@2,10000/sd
...
{0} ok
```

21. Depending on which type of controller module you are adding, do the following:

- If you are installing a StorEdge A3500FC controller module, go to [Step 26](#).
- If you are adding a StorEdge A3500 controller module, 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 that are 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@6,0/QLGC,isp@2,10000
{0} ok .properties
scsi-initiator-id      00000007
...
```

22. Did you power off the second node to install a host adapter?

- If not, go to [Step 26](#).
- If you powered off the second node, boot it now and wait for it to join the cluster.

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

For more information, see the *Sun Cluster 3.0 U1 System Administration Guide*.

23. Check the StorEdge A3500/A3500FC controller module NVSRAM file revision, and if necessary, install the most recent revision.

For the NVSRAM file revision number and boot level, see the *Sun StorEdge RAID Manager Release Notes*. For the procedure on upgrading the NVSRAM file, see the *Sun StorEdge RAID Manager User's Guide*.

24. Check the StorEdge A3500/A3500FC controller module firmware revision, and, if necessary, install the most recent firmware revision.

For the revision number and boot level of the StorEdge A3500/A3500FC controller module firmware, see the *Sun StorEdge RAID Manager Release Notes*. For the procedure on upgrading the StorEdge A3500/A3500FC controller firmware, see ["How to Upgrade Controller Module Firmware in a Running Cluster" on page 163](#).

25. One at a time, boot each node into cluster mode.

```
# reboot
```

26. On one node, verify that the DIDs have been assigned to the StorEdge A3500/A3500FC LUNs for all nodes that are attached to the StorEdge A3500/A3500FC system:

```
# scdidadm -L
```

27. Verify that the controller module is set to active/active mode (if it is not, set it to active/active).

For more information on controller modes, see the *Sun StorEdge RAID Manager Installation and Support Guide* and the *Sun StorEdge RAID Manager User's Guide*.

Where to Go From Here

To create a LUN from disk drives that are unassigned, see [“How to Create a LUN” on page 133](#).

To upgrade StorEdge A3500/A3500FC controller module firmware, see [“How to Upgrade Controller Module Firmware in a Running Cluster” on page 163](#).

▼ How to Remove a StorEdge A3500/A3500FC System From a Running Cluster

Use this procedure to remove a StorEdge A3500/A3500FC system from a running cluster.



Caution – This procedure removes all data that is on the StorEdge A3500/A3500FC system you remove.

1. **Migrate any Oracle Parallel Server (OPS) tables, data services, or volumes off the StorEdge A3500/A3500FC system.**
2. **Halt all activity to the StorEdge A3500/A3500FC controller module.**
See the *Sun StorEdge RAID Manager User's Guide* and your operating system documentation for instructions.
3. **Does a volume manager manage any of the LUNs on the StorEdge A3500/A3500FC controller module you are removing?**

- If not, go to [Step 4](#).
- If a volume manager does manage the LUN, 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. See the following paragraph for additional VERITAS Volume Manager commands that are required.

LUNs that were managed by VERITAS Volume Manager must be completely removed from VERITAS Volume Manager control before you can delete them. To remove the LUNs, after you delete the LUN from any disk group, use the following commands:

```
# vxdisk offline cNtXdY
# vxdisk rm cNtXdY
```

4. **Disconnect all cables from the StorEdge A3500/A3500FC system and remove the hardware from your cluster.**
5. **From one node, delete the LUN.**

For the procedure on deleting a LUN, see the *Sun StorEdge RAID Manager User's Guide*.

6. Remove the paths to the LUN(s) you are deleting:

```
# rm /dev/rdisk/cNtXdY*
# rm /dev/dsk/cNtXdY*

# rm /dev/osa/dev/dsk/cNtXdY*
# rm /dev/osa/dev/rdisk/cNtXdY*
```

7. Use the `lad` command to determine the *alternate* paths to the LUN(s) you are deleting.

The RAID Manager software creates two paths to the LUN in the `/dev/osa/dev/rdisk` directory. Substitute the `cNtXdY` number from the other controller module in the disk array to determine the alternate path.

For example, with this configuration:

```
# lad
c0t5d0 1T93600714 LUNS: 0 1
c1t4d0 1T93500595 LUNS: 2
```

The alternate paths would be the following.

```
/dev/osa/dev/dsk/c1t4d1*
/dev/osa/dev/rdisk/c1t4d1*
```

8. Remove the *alternate* paths to the LUN(s) you are deleting:

```
# rm /dev/osa/dev/dsk/cNtXdY*
# rm /dev/osa/dev/rdisk/cNtXdY*
```

9. On all cluster nodes, remove references to the StorEdge A3500/A3500FC system:

```
# scdidadm -C
```

10. Are you are removing the *last* StorEdge A3500FC controller module from a hub in your cluster?

- If not, go to [Step 11](#).
- If you are removing the last StorEdge A3500FC controller module from a particular hub, remove the hub hardware and cables from your cluster.

11. Remove any unused host adapter from nodes that were attached to the StorEdge A3500/A3500FC system:

a. Shut down and power off the first node from which you are removing a host adapter:

```
# scswitch -S -h nodename  
# shutdown -y -g0 -i0
```

For the procedure on shutting down and powering off a node, see the *Sun Cluster 3.0 U1 System Administration Guide*.

b. Remove the host adapter from the first node.

See the documentation that came with your node hardware for removal instructions.

c. Boot the node and wait for it to rejoin the cluster.

```
# boot -r
```

d. Repeat [Step a](#) through [Step c](#) for the second node that was attached to the StorEdge A3500/A3500FC system.

12. Switch the cluster back online.

```
# scswitch -z
```

13. Are you removing the *last* StorEdge A3500/A3500FC system from your cluster?

- If not, you are finished with this procedure.
- If you are removing the last StorEdge A3500/A3500FC system from your cluster, remove StorEdge A3500/A3500FC software packages. For the procedure on removing software packages, see the documentation that shipped with your StorEdge A3500/A3500FC system.

▼ How to Replace a Failed Controller or Restore an Offline Controller

Use this procedure to replace a StorEdge A3500/A3500FC controller, or to restore an offline controller.

For conceptual information on SCSI reservations and failure fencing, see the *Sun Cluster 3.0 U1 Concepts*.

1. **On both nodes, to prevent LUNs from automatic assignment to the controller that is being brought online, 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 5](#) 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 the *Sun StorEdge RAID Manager Installation and Support Guide*.

2. **Restart the RAID Manager daemon:**

```
# /etc/init.d/amdemon stop
# /etc/init.d/amdemon start
```

3. **Do you have a failed controller?**

- If your controller module is offline, but does not have a failed controller, go to [Step 4](#).
- If you have a failed controller, replace the failed controller with a new controller, *but do not bring the controller online*.

For the procedure on replacing StorEdge A3500/A3500FC controllers, see the *Sun StorEdge A3500/A3500FC Controller Module Guide* and the *Sun StorEdge RAID Manager Installation and Support Guide* for additional considerations.

4. **On one node, use the RAID Manager GUI's Recovery application to restore the controller online.**

Note – You must use the RAID Manager GUI'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 the *Sun StorEdge RAID Manager User's Guide*. If you have problems with bringing the controller online, see the *Sun StorEdge RAID Manager Installation and Support Guide*.

- 5. On one node that is connected to the StorEdge A3500/A3500FC system, verify that the controller has the correct SCSI reservation state.**

Run the `sccidadm(1M)` repair option (`-R`) on LUN 0 of the controller you want to bring online:

```
# sccidadm -R /dev/dsk/cNtXdY
```

- 6. Set the controller to active/active mode and assign LUNs to it.**

For more information on controller modes, see the *Sun StorEdge RAID Manager Installation and Support Guide* and the *Sun StorEdge RAID Manager User's Guide*.

- 7. Reset the `System_LunReDistribution` parameter in the `/etc/raid/rmparams` file to `true`.**

For the procedure on changing the `rmparams` file, see the *Sun StorEdge RAID Manager Installation and Support Guide*.

- 8. Restart the RAID Manager daemon:**

```
# /etc/init.d/amdemon stop  
# /etc/init.d/amdemon start
```


How to Upgrade Controller Module Firmware in a Running Cluster

Use this procedure to upgrade firmware in a StorEdge A3500/A3500FC controller module, in a running cluster. Depending on which firmware you are upgrading, you must use either the *online* or *offline* method, as described in the *Sun StorEdge RAID Manager User's Guide*.

1. Are you upgrading the NVSRAM firmware file?

- If you are not upgrading the NVSRAM file, you can use the online method.
Upgrade the firmware by using the online method, as described in the *Sun StorEdge RAID Manager User's Guide*. No special steps are required for a cluster environment.
- If you are upgrading the NVSRAM file, you must use the offline method, using one of the following two procedures:
 - If the data on your StorEdge A3500/A3500FC controller module is mirrored on another controller module, use the procedure that is described in [Step 2](#).
 - If the data on your StorEdge A3500/A3500FC controller module is *not* mirrored on another controller module, use the procedure that is described in [Step 3](#).

2. Use this step if you are upgrading the NVSRAM and other firmware files on a controller module that has its data mirrored:

- a. **Halt all activity to the StorEdge A3500/A3500FC controller module.**
For more information, see your Solstice DiskSuite or VERITAS Volume Manager documentation.
- b. **Update the firmware files by using the offline method, as described in the *Sun StorEdge RAID Manager User's Guide*.**
- c. **Restore all activity to the StorEdge A3500/A3500FC controller module.**
For more information, see your Solstice DiskSuite or VERITAS Volume Manager documentation.
This step completes the firmware upgrade.

3. Use this step if you are upgrading the NVSRAM and other firmware files on a controller module that does *not* have its data mirrored:

a. Shut down the entire cluster.

```
# scshutdowm -y -g0 -i0
```

For the full procedure on shutting down a cluster, see the *Sun Cluster 3.0 U1 System Administration Guide*.

b. Boot one node that is attached to the controller module into non-cluster mode.

```
ok boot -x
```

For the full procedure on booting a node into non-cluster mode, see the *Sun Cluster 3.0 U1 System Administration Guide*.

c. Update the firmware files using the offline method, as described in the *Sun StorEdge RAID Manager User's Guide*.

d. Reboot both nodes into cluster mode.

```
ok boot
```

For the full procedure on booting nodes into the cluster, see the *Sun Cluster 3.0 U1 System Administration Guide*.

This step completes the firmware upgrade.

▼ How to Add a Disk Drive in a Running Cluster

Use this procedure to add a disk drive to a StorEdge A3500/A3500FC controlled disk array that is in a running cluster.



Caution – If the disk drive you are adding was previously owned by another controller module, preformat the disk drive to wipe clean the old DacStore information before adding it to this disk array.

1. Install the new disk drive to the disk array.

For the procedure on installing a disk drive, see the *Sun StorEdge D1000 Storage Guide*.

2. Allow the disk drive to spin up (approximately 30 seconds).

3. Run Health Check to ensure that the new disk drive is not defective.

For instructions on running Recovery Guru and Health Check, see the *Sun StorEdge RAID Manager User's Guide*.

4. Fail the new drive, then revive the drive to update DacStore on the drive.

For instructions on failing drives and manual recovery procedures, see the *Sun StorEdge RAID Manager User's Guide*.

5. Repeat [Step 1](#) through [Step 4](#) for each disk drive you are adding.

Where to Go From Here

To create LUNs for the new drives, see [“How to Create a LUN” on page 133](#) for more information.

▼ How to Replace a Failed Disk Drive in a Running Cluster

Use this procedure to replace a failed disk drive in a running cluster.

1. Does replacing the disk drive affects any LUN's availability?

- If not, go to [Step 2](#).
- If the replacement does affect LUN availability, remove the LUN(s) from volume management control. For more information, see your Solstice DiskSuite or VERITAS Volume Manager documentation.

2. Replace the disk drive in the disk array.

For the procedure on replacing a disk drive, see the *Sun StorEdge D1000 Storage Guide*.

3. Run Health Check to ensure that the new disk drive is not defective.

For instructions on running Recovery Guru and Health Check, see the *Sun StorEdge RAID Manager User's Guide*.

4. Does the failed drive belong to a drive group?

- If the drive does *not* belong to a device group, go to [Step 5](#).
- If the drive is part of a device group, reconstruction is started automatically. If reconstruction does not start automatically for any reason, then select Reconstruct from the Manual Recovery application. Do not select Revive. When reconstruction is complete, go to [Step 6](#).

5. Fail the new drive, then revive the drive to update DacStore on the drive.

For instructions on failing drives and manual recovery procedures, see the *Sun StorEdge RAID Manager User's Guide*.

6. If you removed LUNs from volume management control in [Step 1](#), return the LUN(s) to volume management control.

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

▼ How to Remove a Disk Drive From a Running Cluster

Use this procedure to remove a disk drive from a running cluster.

1. **Remove the logical unit number (LUN) that is associated with the disk drive you are removing.**

For the procedure on removing a LUN, see [“How to Delete a LUN” on page 136](#).

2. **Remove the disk drive from the disk array.**

For the procedure on removing a disk drive, see the *Sun StorEdge D1000 Storage Guide*.



Caution – After you remove the disk drive, install a dummy drive to maintain proper cooling.

How to Upgrade Disk Drive Firmware in a Running Cluster

Note – Only qualified service personnel should perform disk drive firmware updates. If you need to upgrade drive firmware, contact your local Sun solution center or Sun service provider.

▼ How to Replace a Host Adapter in a Node (Connected to a StorEdge A3500 System)

Note – This section describes the procedure for replacing a failed host adapter in a running node that is attached to a StorEdge A3500 *SCSI-based* system. For the same procedure for a cluster node that is attached to an A3500FC system, see [“How to Replace a Host Adapter in a Node \(Connected to a StorEdge A3500FC System\)”](#) on page 170.

In the following procedure, node 1’s host adapter on SCSI bus A needs replacement but node 2 remains in service.

Note – Several steps in this procedure require that you halt I/O activity. To halt I/O activity, take the controller module offline by using the RAID Manager GUI’s manual recovery procedure in the *Sun StorEdge RAID Manager User’s Guide*.

1. Without powering off the node, shut down node 1.

```
# scswitch -s -h nodename
# shutdown -y -g0 -i0
```

For the procedure on shutting down a node, see the *Sun Cluster 3.0 U1 System Administration Guide*.

2. From node 2, halt I/O activity to SCSI bus A.

See the *Sun StorEdge RAID Manager User’s Guide* for instructions.

3. From the controller module end of the SCSI cable, disconnect the SCSI bus A cable that connects the StorEdge A3500 controller module to node 1, then replace this cable with a differential SCSI terminator.

4. Restart I/O activity on SCSI bus A.

See the *Sun StorEdge RAID Manager User’s Guide* for instructions.

5. Does servicing the failed host adapter affect SCSI bus B?

- If SCSI bus B is not affected, go to [Step 9](#).
- If SCSI bus B is affected, continue with [Step 6](#).

6. From node 2, halt I/O activity to the StorEdge A3500 controller module on SCSI bus B.

See the *Sun StorEdge RAID Manager User’s Guide* for instructions.

7. **From the controller module end of the SCSI cable, disconnect the SCSI bus B cable that connects the StorEdge A3500 controller module to node 1 and replace this cable with a differential SCSI terminator.**
8. **Restart I/O activity on SCSI bus B.**
See the *Sun StorEdge RAID Manager User's Guide* for instructions.
9. **Power off node 1.**
10. **Replace node 1's host adapter.**
See the documentation that came with your node hardware for instructions.
11. **Power on node 1, but do not allow it to boot. If necessary, halt the system.**
12. **From node 2, halt I/O activity to the StorEdge A3500 controller module on SCSI bus A.**
See the *Sun StorEdge RAID Manager User's Guide* for instructions.
13. **Remove the differential SCSI terminator from SCSI bus A, then reinstall the SCSI cable to connect the StorEdge A3500 controller module to node 1.**
14. **Restart I/O activity on SCSI bus A.**
See the *Sun StorEdge RAID Manager User's Guide* for instructions.
15. **Did you install a differential SCSI terminator to SCSI bus B in [Step 7](#)?**
 - If not, skip to [Step 18](#).
 - If you did install a SCSI terminator to SCSI bus B, halt I/O activity on SCSI bus B, then continue with [Step 16](#).
16. **Remove the differential SCSI terminator from SCSI bus B, then reinstall the SCSI cable to connect the StorEdge A3500 controller module to node 1.**
17. **Restart I/O activity on SCSI bus B.**
See the *Sun StorEdge RAID Manager User's Guide* for instructions.
18. **Bring the StorEdge A3500 controller module back online.**
See the *Sun StorEdge RAID Manager User's Guide* for instructions.
19. **Rebalance all logical unit numbers (LUNs).**
See the *Sun StorEdge RAID Manager User's Guide* for instructions.
20. **Boot node 1 into cluster mode.**

```
{0} ok boot
```

▼ How to Replace a Host Adapter in a Node (Connected to a StorEdge A3500FC System)

Note – This section describes the procedure for replacing a failed host adapter in a node that is attached to a StorEdge A3500FC *fiber-optic-based* system. For the same procedure for a cluster node that is attached to a StorEdge A3500 system, see [“How to Replace a Host Adapter in a Node \(Connected to a StorEdge A3500 System\)”](#) on page 168.

Note – Several steps in this procedure require that you halt I/O activity. To halt I/O activity, take the controller module offline by using the RAID Manager GUI’s manual recovery procedure in the *Sun StorEdge RAID Manager User’s Guide*.

1. Move all Sun Cluster data services off of the node in which you are replacing a host adapter.

See the *Sun Cluster 3.0 U1 Data Services Installation and Configuration Guide* for instructions.

2. Halt all I/O activity on the affected StorEdge A3500FC controller module.

See the *Sun StorEdge RAID Manager User’s Guide* for instructions.

3. Shut down and power off the node in which you are replacing a host adapter.

```
# scswitch -S -h nodename  
# shutdown -y -g0 -i0
```

For the procedure on shutting down and powering off a node, see the *Sun Cluster 3.0 U1 System Administration Guide*.

4. Disconnect the fiber-optic cable from the host adapter that you are replacing.

5. Replace the host adapter in the node.

See the documentation that came with your node hardware for instructions.

6. Connect the fiber optic cable to the new host adapter that you just installed.

7. Boot the node into cluster mode.

```
{0} ok boot
```


8. Restart I/O activity on the affected StorEdge A3500FC controller module.

See the *Sun StorEdge RAID Manager User's Guide* and your operating system documentation for instructions.

9. Check the status of the affected StorEdge A3500FC controller module.

Observe the front panel LEDs and interpret them by using the *Sun StorEdge A3500/A3500FC Controller Module Guide*.

10. Rebalance LUNs that are running on the affected StorEdge A3500FC controller module.

For more information, see the *Sun StorEdge RAID Manager User's Guide*.

11. Move the Sun Cluster data services back to the node in which you replaced the host adapter.

See the *Sun Cluster 3.0 U1 Data Services Installation and Configuration Guide* for instructions.

Installing and Maintaining a Sun StorEdge T3 Disk Tray Single-Controller Configuration

This chapter provides the procedures for installing, configuring, and maintaining Sun StorEdge™ T3 disk trays in a single-controller (non-interconnected) configuration.

This chapter contains the following procedures:

- [“How to Install a StorEdge T3 Disk Tray” on page 174](#)
- [“How to Create a Sun StorEdge T3 Disk Tray Logical Volume” on page 177](#)
- [“How to Remove a Sun StorEdge T3 Disk Tray Logical Volume” on page 179](#)
- [“How to Upgrade StorEdge T3 Disk Tray Firmware” on page 184](#)
- [“How to Replace a Disk Drive” on page 185](#)
- [“How to Add a StorEdge T3 Disk Tray” on page 186](#)
- [“How to Remove a StorEdge T3 Disk Tray” on page 195](#)
- [“How to Replace a Host-to-Hub Component” on page 198](#)
- [“How to Replace a Sun StorEdge FC-100 Hub and Hub-to-Disk Tray Component” on page 199](#)
- [“How to Replace a StorEdge T3 Disk Tray Controller” on page 200](#)
- [“How to Replace a StorEdge T3 Disk Tray Chassis” on page 201](#)
- [“How to Replace a Host Adapter” on page 202](#)

For conceptual information on multihost disks, see the *Sun Cluster 3.0 U1 Concepts* document.

Installing a StorEdge T3 Disk Tray

This section provides the procedure for an initial installation of a new StorEdge T3 disk tray.

▼ How to Install a StorEdge T3 Disk Tray

Use this procedure to install and configure a new StorEdge T3 disk tray in a cluster that is not running. Perform the steps in this procedure in conjunction with the procedures in the *Sun Cluster 3.0 U1 Installation Guide* and your server hardware manual.

1. Install the host adapters in the nodes that are to be connected to the StorEdge T3 disk trays.

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

2. Install the Sun StorEdge FC-100 hubs.

For the procedure on installing a Sun StorEdge FC-100 hub, see the *FC-100 Hub Installation and Service Manual*.

3. Set up a Reverse Address Resolution Protocol (RARP) server on the network you want the new StorEdge T3 disk tray to reside on.

This RARP server enables you to assign an IP address to the new StorEdge T3 disk tray by using the StorEdge T3 disk tray's unique MAC address.

For the procedure on setting up a RARP server, see the *Sun StorEdge T3 Installation, Operation, and Service Manual*.

4. Cable and power on the StorEdge T3 disk trays as shown in [FIGURE 8-1](#).

Note – No restrictions are placed on the hub port assignments. You can connect your StorEdge T3 disk tray and node to any hub port.

For the procedure on installing fiber-optic cables, see the *Sun StorEdge T3 Configuration Guide*. For the procedure on powering on the StorEdge T3 disk tray, see the *Sun StorEdge T3 Installation, Operation, and Service Manual*.




FIGURE 8-1 Cabling a StorEdge T3 Disk Tray in a Single-Controller Configuration

Note – Although [FIGURE 8-1](#) shows a single-controller configuration, two disk trays are shown to illustrate how two non-interconnected disk trays are typically cabled in a cluster to allow data sharing and host-based mirroring.

5. (Optional) Configure the StorEdge T3 disk tray with logical volumes.

For the procedure on configuring the StorEdge T3 disk tray with logical volumes, see the *Sun StorEdge T3 Disk Tray Administrator's Guide*.

6. Telnet to the StorEdge T3 disk tray you are adding, and install the necessary StorEdge T3 disk tray controller firmware.

Revision 1.16a firmware is required for the StorEdge T3 disk tray controller. For the procedure on upgrading firmware, see the firmware patch README.

7. Ensure that this new StorEdge T3 disk tray has a unique target address.

For the procedure on verifying and assigning a target address, see the *Sun StorEdge T3 Configuration Guide*.

8. Reset the StorEdge T3 disk tray.

For the procedure on rebooting or resetting a StorEdge T3 disk tray, see the *Sun StorEdge T3 Installation, Operation, and Service Manual*.

9. Install the Solaris operating environment, and apply the required Solaris patches for Sun Cluster software and StorEdge T3 disk tray support.

For the procedure on installing the Solaris operating environment, see the *Sun Cluster 3.0 U1 Installation Guide*. For the location of required Solaris patches and installation instructions for Sun Cluster software support, see the *Sun Cluster 3.0 U1 Release Notes*. For a list of required Solaris patches for StorEdge T3 disk tray support, see the *Sun StorEdge T3 Disk Tray Release Notes*.

Where to Go From Here

To continue with Sun Cluster software installation tasks, see the *Sun Cluster 3.0 U1 Installation Guide*.

Configuring a StorEdge T3 Disk Tray

This section provides the procedures for configuring a StorEdge T3 disk tray in a running cluster. The following table lists these procedures.

TABLE 8-1 Task Map: Configuring a StorEdge T3 Disk Tray

Task	For Instructions, Go To
Create a disk tray logical volume	“How to Create a Sun StorEdge T3 Disk Tray Logical Volume” on page 177
Remove a disk tray logical volume	“How to Remove a Sun StorEdge T3 Disk Tray Logical Volume” on page 179

▼ How to Create a Sun StorEdge T3 Disk Tray Logical Volume

Use this procedure to create a logical volume. This procedure assumes all cluster nodes are booted and attached to the StorEdge T3 disk tray that is to host the logical volume you are creating.

1. **Telnet to the StorEdge T3 disk tray that is to host the logical volume you are creating.**
2. **Create the logical volume.**

The creation of a logical volume involves adding, mounting, and initializing the logical volume.

For the procedure on creating and initializing a logical volume, see the *Sun StorEdge T3 Disk Tray Administrator's Guide*. For the procedure on mounting a logical volume, see the *Sun StorEdge T3 Installation, Operation, and Service Manual*.

3. **On all cluster nodes, update the `/devices` and `/dev` entries.**

```
# devfsadm
```

After this process, a Solaris logical device name for the new logical volume appears in the `/dev/rdisk` and `/dev/dsk` directories on all cluster nodes that are attached to the StorEdge T3 disk tray.

4. If you are running VERITAS Volume Manager, update VERITAS Volume Manager's device tables on all cluster nodes that are attached to the logical volume you created in [Step 2](#). Otherwise, proceed to [Step 5](#).
5. If necessary, partition the logical volume.
6. From any node in the cluster, update the global device namespace.

```
# scgdevs
```

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

Where to Go From Here

To create a new resource or reconfigure a running resource to use the new StorEdge T3 disk tray logical volume, see the *Sun Cluster 3.0 U1 Data Services Installation and Configuration Guide*.

To configure a logical volume as a quorum device, see the *Sun Cluster 3.0 U1 System Administration Guide* for the procedure on adding a quorum device.

▼ How to Remove a Sun StorEdge T3 Disk Tray Logical Volume

Use this procedure to remove a logical volume. This procedure assumes all cluster nodes are booted and attached to the StorEdge T3 disk tray that hosts the logical volume you are removing.

This procedure defines Node A as the node you begin working with, and Node B as the remaining node.



Caution – This procedure removes all data on the logical volume you are removing.

1. If necessary, migrate all data and volumes off the logical volume you are removing. Otherwise, proceed to [Step 2](#).
2. Is the logical volume you are removing a quorum device?

```
# scstat -q
```

- If yes, remove the quorum device before you proceed.
- If no, proceed to [Step 3](#).

For the procedure on removing a quorum device, see the *Sun Cluster 3.0 U1 System Administration Guide*.

3. If you are running VERITAS Volume Manager, update VERITAS Volume Manager's device tables on all cluster nodes that are attached to the logical volume you are removing. Otherwise, proceed to [Step 4](#).
4. Run the appropriate Solstice DiskSuite or VERITAS Volume Manager commands to remove the reference to the logical unit number (LUN) from any diskset or disk group.
For more information, see your Solstice DiskSuite or VERITAS Volume Manager documentation.

5. Remove the logical volume.

For the procedure on deleting a logical volume, see the *Sun StorEdge T3 Disk Tray Administrator's Guide*.

6. Determine the resource groups and device groups that are running on Node A and Node B.

Record this information because you will use it in [Step 15](#) of this procedure to return resource groups and device groups to these nodes.

```
# scstat
```

7. Move all resource groups and device groups off Node A.

```
# scswitch -S -h nodename
```

8. Shut down Node A.

```
# shutdown -y -g0 -i0
```

9. Boot Node A into cluster mode.

```
{0} ok boot
```

For more information, see the *Sun Cluster 3.0 U1 System Administration Guide*.

10. On Node A, remove the obsolete device IDs (DIDs).

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

11. Move all resource groups and device groups off Node B.

```
# scswitch -S -h nodename
```

12. Shut down Node B.

```
# shutdown -y -g0 -i0
```

13. Boot Node B into cluster mode.

```
{0} ok boot
```

For more information, see the *Sun Cluster 3.0 U1 System Administration Guide*.

14. On Node B, remove the obsolete DIDs.

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

15. Return the resource groups and device groups you identified in [Step 6](#) to Node A and Node B.

```
# scswitch -z -g resource-group -h nodename  
# scswitch -z -D device-group-name -h nodename
```

For more information, see the *Sun Cluster 3.0 U1 System Administration Guide*.

Where to Go From Here

To create a logical volume, see [“How to Create a Sun StorEdge T3 Disk Tray Logical Volume” on page 177](#).

Maintaining a StorEdge T3 Disk Tray

This section provides the procedures for maintaining a StorEdge T3 disk tray. The following table lists these procedures. This section does not include a procedure for adding a disk drive and a procedure for removing a disk drive because a StorEdge T3 disk tray only operates when fully configured.



Caution – If you remove any field-replaceable unit (FRU) for an extended period of time, thermal complications might result. To prevent this complication, the StorEdge T3 disk tray is designed so an orderly shutdown occurs when you remove a component for longer than 30 minutes. A replacement part must be immediately available before starting a FRU replacement procedure. You must replace a FRU within 30 minutes or the StorEdge T3 disk tray, and all attached StorEdge T3 disk trays, will shut down and power off.

TABLE 8-2 Task Map: Maintaining a StorEdge T3 Disk Tray

Task	For Instructions, Go To
Upgrade StorEdge T3 disk tray firmware.	“How to Upgrade StorEdge T3 Disk Tray Firmware” on page 184
Replace a disk drive.	“How to Replace a Disk Drive” on page 185
Add a StorEdge T3 disk tray.	“How to Add a StorEdge T3 Disk Tray” on page 186
Remove a StorEdge T3 disk tray.	“How to Remove a StorEdge T3 Disk Tray” on page 195
Replace a host-to-hub fiber-optic cable.	“How to Replace a Host-to-Hub Component” on page 198
Replace an FC-100/S host adapter GBIC.	“How to Replace a Host-to-Hub Component” on page 198
Replace an FC-100 hub GBIC that connects a FC-100 hub to a host.	“How to Replace a Host-to-Hub Component” on page 198
Replace a hub-to-disk tray fiber optic cable.	“How to Replace a Sun StorEdge FC-100 Hub and Hub-to-Disk Tray Component” on page 199
Replace an FC-100 hub GBIC that connects the FC-100 hub to a StorEdge T3 disk tray.	“How to Replace a Sun StorEdge FC-100 Hub and Hub-to-Disk Tray Component” on page 199

TABLE 8-2 Task Map: Maintaining a StorEdge T3 Disk Tray (*Continued*)

Task	For Instructions, Go To
Replace a Sun StorEdge FC-100 hub.	“How to Replace a Sun StorEdge FC-100 Hub and Hub-to-Disk Tray Component” on page 199
Replace a Sun StorEdge FC-100 hub power cord.	“How to Replace a Sun StorEdge FC-100 Hub and Hub-to-Disk Tray Component” on page 199
Replace a media interface adapter (MIA).	“How to Replace a Sun StorEdge FC-100 Hub and Hub-to-Disk Tray Component” on page 199
Replace a StorEdge T3 disk tray controller.	“How to Replace a StorEdge T3 Disk Tray Controller” on page 200
Replace a StorEdge T3 disk tray chassis.	“How to Replace a StorEdge T3 Disk Tray Chassis” on page 201
Replace a host adapter.	“How to Replace a Host Adapter” on page 202
Replace a Power and Cooling Unit (PCU). Follow the same procedure used in a non-cluster environment.	<i>Sun StorEdge T3 Installation, Operation, and Service Manual</i>
Replace a unit interconnect card (UIC). Follow the same procedure used in a non-cluster environment.	<i>Sun StorEdge T3 Installation, Operation, and Service Manual</i>
Replace a StorEdge T3 disk tray power cable. Follow the same procedure used in a non-cluster environment.	<i>Sun StorEdge T3 Installation, Operation, and Service Manual</i>
Replace an Ethernet cable. Follow the same procedure used in a non-cluster environment.	<i>Sun StorEdge T3 Installation, Operation, and Service Manual</i>

▼ How to Upgrade StorEdge T3 Disk Tray Firmware

Use this procedure to upgrade StorEdge T3 disk tray firmware in a running cluster. StorEdge T3 disk tray firmware includes controller firmware, unit interconnect card (UIC) firmware, and disk drive firmware.



Caution – Perform this procedure on one StorEdge T3 disk tray at a time. This procedure requires that you reset the StorEdge T3 disk tray you want to upgrade. If you reset more than one StorEdge T3 disk tray, your cluster will lose access to data if the StorEdge T3 disk trays are submirrors of each other.

1. **On one node attached to the StorEdge T3 disk tray you are upgrading, detach that StorEdge T3 disk tray's submirrors.**

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

2. **Apply the controller, disk drive, and UIC firmware patches.**

For the list of required StorEdge T3 disk tray patches, see the *Sun StorEdge T3 Disk Tray Release Notes*. For the procedure on applying firmware patches, see the firmware patch README. For the procedure on verifying the firmware level, see the *Sun StorEdge T3 Installation, Operation, and Service Manual*.

3. **Reset the StorEdge T3 disk tray, if you have not already done so.**

For the procedure on rebooting a StorEdge T3 disk tray, see the *Sun StorEdge T3 Installation, Operation, and Service Manual*.

4. **Reattach the submirrors to resynchronize them.**

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

▼ How to Replace a Disk Drive

Use this procedure to replace one failed disk drive in a StorEdge T3 disk tray in a running cluster.



Caution – If you remove any field replaceable unit (FRU) for an extended period of time, thermal complications might result. To prevent this complication, the StorEdge T3 disk tray is designed so an orderly shutdown occurs when you remove a component for longer than 30 minutes. A replacement part must be immediately available before starting a FRU replacement procedure. You must replace a FRU within 30 minutes or the StorEdge T3 disk tray, and all attached StorEdge T3 disk trays, will shut down and power off.

1. **If the failed disk drive impacted the logical volume’s availability, remove the logical volume from volume management control. Otherwise, proceed to [Step 2](#).**

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

2. **Replace the disk drive.**

For the procedure on replacing a disk drive, see the *Sun StorEdge T3 Installation, Operation, and Service Manual*.

3. **If you removed a LUN from volume management control in [Step 1](#), return the LUN(s) to volume management control. Otherwise, [Step 2](#) completes this procedure.**

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

▼ How to Add a StorEdge T3 Disk Tray

Use this procedure to add a new StorEdge T3 disk tray to a running cluster.

This procedure defines Node A as the node you begin working with, and Node B as the remaining node.

1. **Set up a Reverse Address Resolution Protocol (RARP) server on the network the new StorEdge T3 disk tray is to reside on, and then assign an IP address to the new StorEdge T3 disk tray.**

This RARP server enables you to assign an IP address to the new StorEdge T3 disk tray by using the StorEdge T3 disk tray's unique MAC address.

For the procedure on setting up a RARP server, see the *Sun StorEdge T3 Installation, Operation, and Service Manual*.

2. **Install the media interface adapter (MIA) in the StorEdge T3 disk tray you want to add as shown in [FIGURE 8-2](#).**

For the procedure on installing a media interface adapter (MIA), see the *Sun StorEdge T3 Configuration Guide*.

3. **If necessary, install a gigabit interface converter (GBIC) in the Sun StorEdge FC-100 hub as shown in [FIGURE 8-2](#).**

This GBIC enables you to connect the Sun StorEdge FC-100 hub to the StorEdge T3 disk tray you want to add.

Note – No restrictions are placed on the hub port assignments. You can connect your StorEdge T3 disk tray and node to any hub port.

For the procedure on installing an FC-100 hub GBIC, see the *FC-100 Hub Installation and Service Manual*.

4. **Install the 10Base-T Ethernet cable between the StorEdge T3 disk tray and the Local Area Network (LAN), as shown in [FIGURE 8-2](#).**
5. **Power on the StorEdge T3 disk tray.**

Note – The StorEdge T3 disk tray might require a few minutes to boot.

For the procedure on powering on a StorEdge T3 disk tray, see the *Sun StorEdge T3 Installation, Operation, and Service Manual*.

6. **Telnet to the StorEdge T3 disk tray you are adding, and, if necessary, install the required StorEdge T3 disk tray controller firmware.**

Revision 1.16a firmware is required for the StorEdge T3 disk tray controller. For the procedure on upgrading firmware, see the firmware patch [README](#).

7. Does this new StorEdge T3 disk tray have a unique target address?

- If yes, proceed to [Step 8](#).
- If no, change the target address for this new StorEdge T3 disk tray.

For the procedure on verifying and assigning a target address, see the *Sun StorEdge T3 Configuration Guide*.

8. Install a fiber-optic cable between the Sun StorEdge FC-100 hub and the StorEdge T3 disk tray as shown in [FIGURE 8-2](#).

Note – No restrictions are placed on the hub port assignments. You can connect your StorEdge T3 disk tray and node to any hub port.

For the procedure on installing a fiber-optic cable, see the *Sun StorEdge T3 Configuration Guide*.

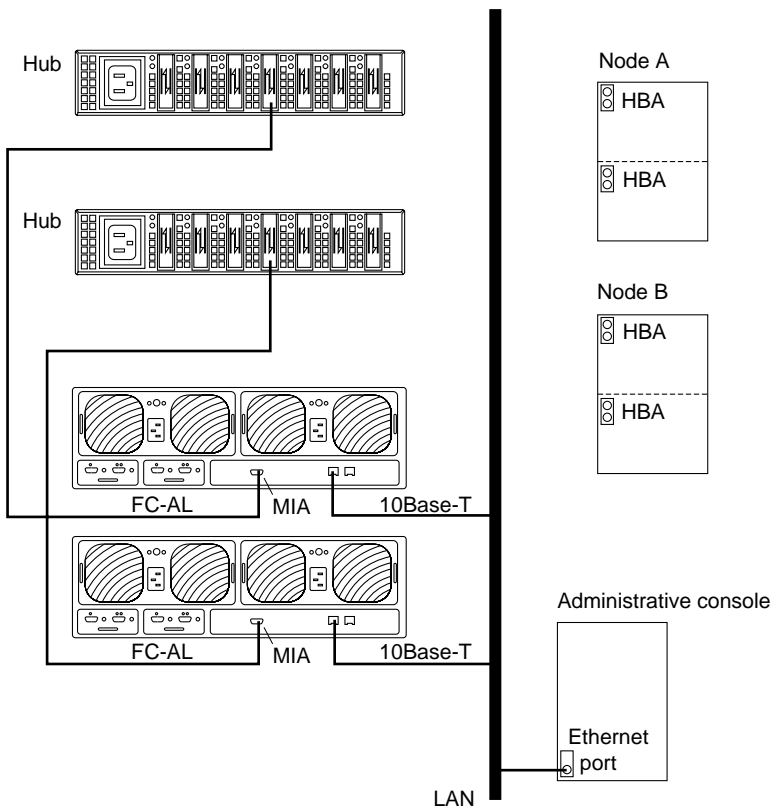


FIGURE 8-2 Adding a StorEdge T3 Disk Tray in a Single-Controller Configuration

Note – Although [FIGURE 8-2](#) shows a single-controller configuration, two disk trays are shown to illustrate how two non-interconnected disk trays are typically cabled in a cluster to allow data sharing and host-based mirroring.

9. Configure the new StorEdge T3 disk tray.

For the procedure on creating a logical volume, see the *Sun StorEdge T3 Disk Tray Administrator's Guide*.

10. Determine the resource groups and device groups that are running on Node A and Node B.

Record this information because you will use it in [Step 42](#) of this procedure to return resource groups and device groups to these nodes.

```
# scstat
```

11. Move all resource groups and device groups off Node A.

```
# scswitch -S -h nodename
```

12. Do you need to install a host adapter in Node A?

- If yes, proceed to [Step 13](#).
- If no, skip to [Step 20](#).

13. Is the host adapter you are installing the first FC-100/S host adapter on Node A?

- If no, skip to [Step 15](#).
- If yes, determine whether the Fibre Channel support packages are already installed on these nodes. This product requires the following packages.

```
# 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 social Device Driver
system    SUNWluxlx  Sun Enterprise Network Array social Device Driver (64-bit)
system    SUNWluxop  Sun Enterprise Network Array firmware and utilities
```

14. Are the Fibre Channel support packages installed?

- If yes, proceed to [Step 15](#).

- If no, install them.

The StorEdge T3 disk tray packages are located in the `Product` directory of the Solaris CD-ROM. Use the `pkgadd` command to add any necessary packages.

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

15. Stop the Sun Cluster software on Node A and shut down Node A.

```
# shutdown -y -g0 -i0
```

For the procedure on shutting down a node, see the *Sun Cluster 3.0 U1 System Administration Guide*.

16. Power off Node A.

17. Install the host adapter in Node A.

For the procedure on installing a host adapter, see the documentation that shipped with your host adapter and node.

18. If necessary, power on and boot Node A.

```
{0} ok boot -x
```

For more information, see the *Sun Cluster 3.0 U1 System Administration Guide*.

19. If necessary, upgrade the host adapter firmware on Node A.

For the required host adapter firmware patch, see the *Sun StorEdge T3 Disk Tray Release Notes*. For the procedure on applying the host adapter firmware patch, see the firmware patch README.

20. If necessary, install gigabit interface converters (GBIC), as shown in [FIGURE 8-3](#).

For the procedure on installing an FC-100 hub GBIC, see the *FC-100 Hub Installation and Service Manual*.

21. If necessary, connect a fiber-optic cable between the Sun StorEdge FC-100 hub and Node A as shown in [FIGURE 8-3](#).

Note – No restrictions are placed on hub port assignments. You can connect your StorEdge T3 disk tray and node to any hub port.

For the procedure on installing an FC-100/S host adapter GBIC, see your host adapter documentation. For the procedure on installing a fiber-optic cable, see the *Sun StorEdge T3 Configuration Guide*.




FIGURE 8-3 Adding a StorEdge T3 Disk Tray in a Single-Controller Configuration

22. If necessary, install the required Solaris patches for StorEdge T3 disk tray support on Node A.

For a list of required Solaris patches for StorEdge T3 disk tray support, see the *Sun StorEdge T3 Disk Tray Release Notes*.

23. Shut down Node A.

```
# shutdown -y -g0 -i0
```

24. Perform a reconfiguration boot to create the new Solaris device files and links on Node A.

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

25. Label the new logical volume.

For the procedure on labeling a logical volume, see the *Sun StorEdge T3 Disk Tray Administrator's Guide*.

26. (Optional) On Node A, verify that the device IDs (DIDs) are assigned to the new StorEdge T3 disk tray.

```
# scdidadm -l
```

27. Do you need to install a host adapter in Node B?

- If yes, proceed to [Step 28](#).
- If no, skip to [Step 36](#).

28. Is the host adapter you want to install the first FC-100/S host adapter on Node B?

- If no, skip to [Step 30](#).
- If yes, determine whether the Fibre Channel support packages are already installed on these nodes. This product requires the following packages.

```
# 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 social Device Driver
system    SUNWluxlx   Sun Enterprise Network Array social Device Driver (64-bit)
system    SUNWluxop   Sun Enterprise Network Array firmware and utilities
```

29. Are the Fibre Channel support packages installed?

- If yes, proceed to [Step 30](#).
- If no, install them.

The StorEdge T3 disk tray packages are located in the `Product` directory of the Solaris CD-ROM. Use the `pkgadd` command to add any necessary packages.

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

30. Move all resource groups and device groups off Node B.

```
# scswitch -s -h nodename
```

31. Stop the Sun Cluster software on Node B, and shut down the node.

```
# shutdown -y -g0 -i0
```

For the procedure on shutting down a node, see the *Sun Cluster 3.0 U1 System Administration Guide*.

32. Power off Node B.

For more information, see the *Sun Cluster 3.0 U1 System Administration Guide*.

33. Install the host adapter in Node B.

For the procedure on installing a host adapter, see the documentation that shipped with your host adapter and node.

34. If necessary, power on and boot Node B.

```
{0} ok boot -x
```

For more information, see the *Sun Cluster 3.0 U1 System Administration Guide*.

35. If necessary, upgrade the host adapter firmware on Node B.

For the required host adapter firmware patch, see the *Sun StorEdge T3 Disk Tray Release Notes*. For the procedure on applying the host adapter firmware patch, see the firmware patch README.

36. If necessary, install gigabit interface converters (GBIC) as shown in [FIGURE 8-4](#).

For the procedure on installing an FC-100 hub GBIC, see the *FC-100 Hub Installation and Service Manual*.

37. If necessary, connect a fiber-optic cable between the Sun StorEdge FC-100 hub and Node B as shown in [FIGURE 8-4](#).

For the procedure on installing a FC-100/S host adapter GBIC, see your host adapter documentation. For the procedure on installing a fiber-optic cable, see the *Sun StorEdge T3 Configuration Guide*.




FIGURE 8-4 Adding a StorEdge T3 Disk Tray in a Single-Controller Configuration

38. If necessary, install the required Solaris patches for StorEdge T3 disk tray support on Node B.

For a list of required Solaris patches for StorEdge T3 disk tray support, see the *Sun StorEdge T3 Disk Tray Release Notes*.

39. Shut down Node B.

```
# shutdown -y -g0 -i0
```

40. Perform a reconfiguration boot to create the new Solaris device files and links on Node B.

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

41. (Optional) On Node B, verify that the device IDs (DIDs) are assigned to the new StorEdge T3 disk tray.

```
# scdidadm -l
```

42. Return the resource groups and device groups you identified in [Step 10](#) to Node A and Node B.

```
# scswitch -z -g resource-group -h nodename  
# scswitch -z -D device-group-name -h nodename
```

For more information, see the *Sun Cluster 3.0 U1 System Administration Guide*.

43. Perform volume management administration to incorporate the new logical volumes into the cluster.

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

▼ How to Remove a StorEdge T3 Disk Tray

Use this procedure to permanently remove a StorEdge T3 disk tray and its submirrors from a running cluster. This procedure provides the flexibility to remove the host adapters from the nodes for the StorEdge T3 disk tray you are removing.

This procedure defines Node A as the node you want to begin working with, and Node B as the remaining node.



Caution – During this procedure, you will lose access to the data that resides on the StorEdge T3 disk tray you are removing.

1. **If necessary, back up all database tables, data services, and volumes that are associated with the StorEdge T3 disk tray that you are removing.**
2. **Detach the submirrors from the StorEdge T3 disk tray you are removing in order to stop all I/O activity to the StorEdge T3 disk tray.**

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

3. **Run the appropriate Solstice DiskSuite or VERITAS Volume Manager commands to remove the references to the LUN(s) from any diskset or disk group.**

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

4. **Determine the resource groups and device groups that are running on Node B.**

```
# scstat
```

5. **Move all resource groups and device groups off Node A.**

```
# scswitch -s -h nodename
```

6. **Stop the Sun Cluster software on Node A, and shut down Node A.**

```
# shutdown -y -g0 -i0
```

For the procedure on shutting down a node, see the *Sun Cluster 3.0 U1 System Administration Guide*.

7. Is the StorEdge T3 disk tray you are removing the last StorEdge T3 disk tray that is connected to Node A?

- If yes, disconnect the fiber-optic cable between Node A and the Sun StorEdge FC-100 hub that is connected to this StorEdge T3 disk tray, then disconnect the fiber-optic cable between the Sun StorEdge FC-100 hub and this StorEdge T3 disk tray.
- If no, proceed to [Step 8](#).

For the procedure on removing a fiber-optic cable, see the *Sun StorEdge T3 Configuration Guide*.

8. Do you want to remove the host adapter from Node A?

- If yes, power off Node A.
- If no, skip to [Step 10](#).

9. Remove the host adapter from Node A.

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

10. Without allowing the node to boot, power on Node A.

For more information, see the *Sun Cluster 3.0 U1 System Administration Guide*.

11. Boot Node A into cluster mode.

```
{0} ok boot
```

12. Move all resource groups and device groups off Node B.

```
# scswitch -s -h nodename
```

13. Stop the Sun Cluster software on Node B, and shut down Node B.

```
# shutdown -y -g0 -i0
```

14. Is the StorEdge T3 disk tray you are removing the last StorEdge T3 disk tray that is connected to the Sun StorEdge FC-100 hub.

- If yes, disconnect the fiber-optic cable that connects this Sun StorEdge FC-100 hub and Node B.
- If no, proceed to [Step 15](#).

For the procedure on removing a fiber-optic cable, see the *Sun StorEdge T3 Configuration Guide*.

15. Do you want to remove the host adapter from Node B?

- If yes, power off Node B.
- If no, skip to [Step 18](#).

16. Remove the host adapter from Node B.

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

17. Without allowing the node to boot, power on Node B.

For more information, see the *Sun Cluster 3.0 U1 System Administration Guide*.

18. Boot Node B into cluster mode.

```
{0} ok boot
```

For more information, see the *Sun Cluster 3.0 U1 System Administration Guide*.

19. On all cluster nodes, update the /devices and /dev entries.

```
# devfsadm -C  
# sctidadm -C
```

20. Return the resource groups and device groups you identified in [Step 4](#) to Node A and Node B.

```
# scswitch -z -g resource-group -h nodename  
# scswitch -z -D device-group-name -h nodename
```

▼ How to Replace a Host-to-Hub Component

Use this procedure to replace the following host-to-hub components.

- Host-to-hub fiber-optic cable
- FC-100/S host adapter GBIC
- FC-100 hub GBIC that connects a Sun StorEdge FC-100 hub to a node

1. **On the node that is connected to the host-to-hub connection you are replacing, determine the resource groups and device groups that are running on this node.**

```
# scstat
```

2. **Move all resource groups and device groups to another node.**

```
# scswitch -s -h nodename
```

3. **Replace the host-to-hub component.**

For the procedure on replacing a fiber-optic cable, see the *Sun StorEdge T3 Configuration Guide*. For the procedure on replacing an FC-100 hub GBIC, see the *FC-100 Hub Installation and Service Manual*. For the procedure on replacing a FC-100/S host adapter GBIC, see your host adapter documentation.

4. **Return the resource groups and device groups you identified in [Step 1](#) to the node that is connected to the host-to-hub connection you replaced.**

```
# scswitch -z -g resource-group -h nodename  
# scswitch -z -D device-group-name -h nodename
```

▼ How to Replace a Sun StorEdge FC-100 Hub and Hub-to-Disk Tray Component

Use this procedure to replace the following hub-to-disk tray components.

- Fiber-optic cable that connects a Sun StorEdge FC-100 hub to a StorEdge T3 disk tray
- FC-100 hub GBIC that connects a Sun StorEdge FC-100 hub to a StorEdge T3 disk tray
- Sun StorEdge FC-100 hub
- Sun StorEdge FC-100 hub power cord
- Media interface adapter (MIA)

- 1. Detach the submirrors on the StorEdge T3 disk tray that is connected to the hub-to-disk tray fiber-optic cable you are replacing in order to stop all I/O activity to this StorEdge T3 disk tray.**

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

- 2. Replace the hub-to-disk tray component.**

For the procedure on replacing a fiber-optic cable, see the *Sun StorEdge T3 Configuration Guide*. For the procedure on replacing an FC-100 hub GBIC, a Sun StorEdge FC-100 hub, or a Sun StorEdge FC-100 hub power cord, see the *FC-100 Hub Installation and Service Manual*. For the procedure on replacing an MIA, see the *Sun StorEdge T3 Installation, Operation, and Service Manual*.

- 3. Reattach the submirrors to resynchronize them.**

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

▼ How to Replace a StorEdge T3 Disk Tray Controller

Use this procedure to replace a StorEdge T3 disk tray controller.

- 1. Detach the submirrors on the StorEdge T3 disk tray that is connected to the controller you are replacing in order to stop all I/O activity to this StorEdge T3 disk tray.**

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

- 2. Replace the controller.**

For the procedure on replacing a StorEdge T3 controller, see the *Sun StorEdge T3 Installation, Operation, and Service Manual*.

- 3. Reattach the submirrors to resynchronize them.**

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

▼ How to Replace a StorEdge T3 Disk Tray Chassis

Use this procedure to replace a StorEdge T3 disk tray chassis. This procedure assumes that you are retaining all FRUs other than the chassis and the backplane. To replace the chassis, you must replace both the chassis and the backplane because these components are manufactured as one part.

Note – Only trained, qualified service providers should use this procedure to replace a StorEdge T3 disk tray chassis.

1. **Detach the submirrors on the StorEdge T3 disk tray that is connected to the chassis you are replacing in order to stop all I/O activity to this StorEdge T3 disk tray.**

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

2. **Replace the chassis/backplane.**

For the procedure on replacing a StorEdge T3 chassis, see the *Sun StorEdge T3 Field Service Manual*.

3. **Reattach the submirrors to resynchronize them.**

Note – Account for the change in the World Wide Number (WWN).

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

▼ How to Replace a Host Adapter

Use this procedure to replace a failed host adapter in a running cluster. Node A in this procedure refers to the node with the failed host adapter you are replacing. Node B is a backup node.

1. Determine the resource groups and device groups that are running on Node A and Node B.

Record this information because you will use it in [Step 8](#) of this procedure to return resource groups and device groups to these nodes.

```
# scstat
```

2. Move all resource groups and device groups off Node A.

```
# scswitch -s -h nodename
```

3. Shut down Node A.

```
# shutdown -y -g0 -i0
```

4. Power off Node A.

For more information, see the *Sun Cluster 3.0 U1 System Administration Guide*.

5. Replace the failed host adapter.

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

6. Power on Node A.

For more information, see the *Sun Cluster 3.0 U1 System Administration Guide*.

7. Boot Node A into cluster mode.

```
{0} ok boot
```

For more information, see the *Sun Cluster 3.0 U1 System Administration Guide*.

8. Return the resource groups and device groups you identified in [Step 1](#) to Node A and Node B.

```
# scswitch -z -g resource-group -h nodename  
# scswitch -z -D device-group-name -h nodename
```

For more information, see the *Sun Cluster 3.0 U1 System Administration Guide*.

Verifying Sun Cluster Hardware Redundancy

This appendix describes 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 206](#)
- [“How to Test Cluster Interconnects” on page 207](#)
- [“How to Test Network Adapter Failover Groups” on page 209](#)

If your cluster passes these tests, your hardware has adequate redundancy. This redundancy means that your 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 206](#) and [“How to Test Cluster Interconnects” on page 207](#), 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 has 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 the *Sun Cluster 3.0 U1 Concepts* document.

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 that is 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 consoles of the existing nodes.

2. On another node, run the `scstat` command to verify that the secondary took ownership of the device group that is mastered by the primary.

Look for the output that shows the device group ownership.

```
# scstat
```

3. Power on the initial primary and boot it 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](#) because the system boot process moves ownership of the device group back to the initial primary. Otherwise, proceed to [Step 4](#) 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 primary 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 a `faulted` status to the cluster transport path you disconnected. This fault does not result in a failover.

2. **Disconnect the remaining cluster transport cable from the primary node you identified in [Step 1](#).**

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 a `faulted` status to the cluster transport path you disconnected. This action causes the primary node to go down, resulting in a partitioned cluster.

For conceptual information on failure fencing or *split brain*, see the *Sun Cluster 3.0 U1 Concepts* document.

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](#), into cluster mode.**

```
{0} ok boot
```

6. Verify that the Sun Cluster software assigned a path online status to each cluster transport path you reconnected in [Step 4](#).

```
# scstat
```

If you have the device group failback option enabled, skip [Step 7](#) because the system boot process moves ownership of the device group back to the initial primary. Otherwise, go to [Step 7](#) 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.

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 that are 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 `nvrामrc` editor and `nvedit` keystroke commands.

An `nvrामrc` 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 `nvrामrc` script contains data, add the entries to the end of the script. To edit an `nvrामrc` script or merge new lines in an `nvrामrc` script, use `nvrामrc` editor and `nvedit` keystroke commands.

[TABLE B-1](#) and [TABLE B-2](#) list useful `nvrामrc` editor and `nvedit` keystroke commands. For an entire list of `nvrामrc` editor and `nvedit` keystroke commands, see the *OpenBoot 3.x Command Reference Manual*.

TABLE B-1 NVRAMRC Editor Commands

Command	Description
<code>nvedit</code>	Enter the <code>nvrामrc</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>nvrामrc</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>nvrामrc</code> and discard the contents of the temporary buffer.
<code>nvquit</code>	Discard the contents of the temporary buffer, without writing it to <code>nvrामrc</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>nvrampc</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>nvrampc</code> editing buffer.
^R	Replace the current line.
Delete	Delete the previous character.
Return	Insert a new line at the cursor position and advance to the next line.

Recabling Disk Devices

This appendix contains the procedures for recabling disk devices.

This appendix provides the following procedures:

- [“How to Move a Disk Cable to a New Host Adapter” on page 214](#)
- [“How to Move a Disk Cable From One Node to Another” on page 216](#)
- [“How to Update Sun Cluster Software to Reflect Proper Device Configuration” on page 218](#)

Moving a Disk Cable

Although you can move a disk cable to a different host adapter on the same bus because of a failed host adapter, it is better to replace the failed host adapter rather than recable to a different host adapter. However, you might want to move a disk cable to a different host adapter on the same bus to improve performance.

This section provides the following two procedures for moving a disk cable:

- [“How to Move a Disk Cable to a New Host Adapter” on page 214](#)
- [“How to Move a Disk Cable From One Node to Another” on page 216](#)

Use one of these two procedures to prevent interference with normal operation of your cluster when you want to move a disk cable to a different host adapter on the same bus. If you do not follow these procedures correctly, you might see an error the next time you run the `scdidadm -r` command or the `scgdevs` command. If you see an error message that says `did reconfiguration discovered invalid diskpath`, go to [“How to Update Sun Cluster Software to Reflect Proper Device Configuration” on page 218](#).

▼ How to Move a Disk Cable to a New Host Adapter

Use this procedure to move a disk cable to a new host adapter within a node.

Caution – Failure to follow this cabling procedure might introduce invalid device IDs (DIDs) and render the devices inaccessible.

1. Stop all I/O to the affected disk(s).

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

2. Unplug the cable from the old host adapter.

3. From the local node, unconfigure all drives that are affected by the recabling.

```
# cfgadm
```

Or reboot the local node.

```
# reboot -- -r
```

4. From the local node, update the Solaris device link.

```
# devfsadm -C
```

5. From the local node, update the DID device path.

```
# scdidadm -C
```

6. Connect the cable to the new host adapter.

7. From the local node, configure the drives in the new location.

```
# cfgadm
```

Or reboot the local node.

```
# reboot -- -r
```

8. Add the new DID device path.

```
# scgdevs
```

Where to Go From Here

If you did not follow this procedure correctly, you might see an error the next time you run the `scdidadm -r` command or the `scgdevs` command. If you see an error message that says `did reconfiguration discovered invalid diskpath`, go to [“How to Update Sun Cluster Software to Reflect Proper Device Configuration” on page 218](#).

▼ How to Move a Disk Cable From One Node to Another

Use this procedure to move a disk cable from one node to another node.

Caution – Failure to follow this cabling procedure might introduce invalid device IDs (DIDs) and render the devices inaccessible.

1. **Delete all references to the DID device path you want to remove from all volume manager and data service configurations.**

For more information, see the *Sun Cluster 3.0 U1 Data Services Installation and Configuration Guide* and your Solstice DiskSuite or VERITAS Volume Manager documentation.

2. **Stop all I/O to the affected disk(s).**

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

3. **Unplug the cable from the old node.**

4. **From the old node, unconfigure all drives that are affected by the recabling.**

```
# cfgadm
```

Or reboot the old node.

```
# reboot -- -r
```

5. **From the old node, update the Solaris device link.**

```
# devfsadm -C
```

6. **From the old node, update the DID device path.**

```
# scdidadm -C
```

7. **Connect the cable to the new node.**

8. From the new node, configure the drives in the new location.

```
# cfgadm
```

Or reboot the new node.

```
# reboot -- -r
```

9. From the new node, create the new Solaris device links.

```
# devfsadm
```

10. From the new node, add the new DID device path.

```
# scgdevs
```

11. Add the DID device path on the new node to any volume manager and data service configurations that are required.

When you configure data services, check that your node failover preferences are set to reflect the new configuration.

For more information, see the *Sun Cluster 3.0 U1 Data Services Installation and Configuration Guide* and your Solstice DiskSuite or VERITAS Volume Manager documentation.

Where to Go From Here

If you did not follow this procedure correctly, you might see an error the next time you run the `scdidadm -r` command or the `scgdevs` command. If you see an error message that says `did reconfiguration discovered invalid diskpath`, go to [“How to Update Sun Cluster Software to Reflect Proper Device Configuration”](#) on page 218.

▼ How to Update Sun Cluster Software to Reflect Proper Device Configuration

If you see the following error when you run the `scdidadm -r` command or the `scgdevs` command, the Sun Cluster software does not reflect the proper device configuration because of improper device recabling.

```
did reconfiguration discovered invalid diskpath.  
This path must be removed before a new path  
can be added. Please run did cleanup (-C)  
then re-run did reconfiguration (-r).
```

Use this procedure to ensure that the Sun Cluster software becomes aware of the new configuration and to guarantee device availability.

1. Ensure that your cluster meets the following conditions.

- The cable configuration is correct.
- The cable you are removing is detached from the old node.
- The old node is removed from any volume manager or data service configurations that are required.

For more information, see the *Sun Cluster 3.0 U1 Data Services Installation and Configuration Guide* and your Solstice DiskSuite or VERITAS Volume Manager documentation.

2. From all nodes, one node at a time, unconfigure all drives.

```
# cfgadm
```

Or reboot all nodes, one node at a time.

```
# reboot -- -r
```

3. From all nodes, one node at a time, update the Solaris device link.

```
# devfsadm -C
```

4. From all nodes, one node at a time, update the DID device path.

```
# scdidadm -C
```


5. From all nodes, one node at a time, configure all drives.

```
# cfgadm
```

Or reboot all nodes, one node at a time.

```
# reboot -- -r
```

6. From any node, add the new DID device path.

```
# scgdevs
```

7. From all nodes that are affected by the recabling, verify that SCSI reservations are in the correct state.

```
# scdidadm -R device
```

