



Sun Cluster 3.x With Sun StorEdge 3900 or 6900 Series System Manual

Sun Microsystems, Inc.
4150 Network Circle
Santa Clara, CA 95054
U.S.A.

Part No: 817-0179-10
July 2003, Revision A

Copyright 2003 Sun Microsystems, Inc. 4150 Network Circle, Santa Clara, CA 95054 U.S.A. All rights reserved.

This product or document is protected by copyright and 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, AnswerBook, AnswerBook2, and Solaris 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 2003 Sun Microsystems, Inc. 4150 Network Circle, Santa Clara, CA 95054 U.S.A. Tous droits réservés.

Ce produit ou document est protégé par un copyright et 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 du système 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, AnswerBook, AnswerBook2, et Solaris 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.

CETTE PUBLICATION EST FOURNIE "EN L'ETAT" ET AUCUNE GARANTIE, EXPRESSE OU IMPLICITE, N'EST ACCORDEE, Y COMPRIS DES GARANTIES CONCERNANT LA VALEUR MARCHANDE, L'APTITUDE DE LA PUBLICATION A REPOUDRE A UNE UTILISATION PARTICULIERE, OU LE FAIT QU'ELLE NE SOIT PAS CONTREFAISANTE DE PRODUIT DE TIERS. CE DENI DE GARANTIE NE S'APPLIQUERAIT PAS, DANS LA MESURE OU IL SERAIT TENU JURIDIQUEMENT NUL ET NON AVENU.



030807 @ 5943



Contents

Preface 5

1	Installing and Maintaining a Sun StorEdge 3900 or 6900 Series System	11
	Installing Storage Systems	12
	▼ How to Install a Storage System	12
	Configuring Storage Systems	15
	How to Configure a LUN	15
	How to Configure VLUNs on the Virtualization Engines (StorEdge 6900 Series Only)	16
	Maintaining Storage Systems	17
	▼ How to Add a Storage System	19
	▼ How to Remove a Storage System	24
	▼ How to Replace a Virtualization Engine (StorEdge 6900 Series Only)	26
	▼ How to Replace a Node-to-Switch Component	27
	▼ How to Remove a Storage Array	27
	How to Upgrade Storage Array Firmware	30
	▼ How to Replace a Disk Drive	32
	▼ How to Replace a Node's Host Adapter	32

Index 35

Preface

The *Sun Cluster 3.x With Sun StorEdge 3900 or 6900 Series System Manual* provides procedures specific to Sun StorEdge™ 3900 and 6900 storage devices that are placed in a Sun™ Cluster environment.

Use this manual with any version of Sun Cluster 3.x software. Unless otherwise noted, procedures are the same for all Sun Cluster 3.x versions. See the “Revision History” on page 6 for a list of changes to this manual.

Who Should Use This Book

This book is for Sun representatives who are performing the initial installation of a Sun Cluster configuration and for system administrators who are responsible for maintaining the system.

This document is intended for experienced system administrators with extensive knowledge of Sun software and hardware. Do not use this document as a planning or presales guide. You should have already determined your system requirements and purchased the appropriate equipment and software before reading this document.

How This Book Is Organized

This book contains one chapter that consists of three major sections.

Section 1 discusses how to install Sun StorEdge 3900 or 6900 storage systems

Section 2 discusses how to configure logical units on Sun StorEdge 3900 or 6900 storage systems

Section 3 describes how to maintain Sun StorEdge 3900 or 6900 storage systems in a running cluster.

Revision History

The following table lists the information that has been revised or added since the initial release of this documentation. The table also lists the revision date for these changes.

TABLE P-1 Sun Cluster 3.x With Sun StorEdge 3900 or 6900 Series System Manual

Revision Date	New Information
July 2003	Updated documentation with editorial changes.
April 2003	Corrected references in “How to Upgrade Storage Array Firmware” on page 30.
	Updated documentation to state that Virtualization Engines require no cluster-specific procedures.
	Provided a list of FRUs that require no cluster-specific procedures.

Related Documentation

The following books provide conceptual information or procedures to administer hardware and applications. If you plan to use this documentation in a hard-copy format, ensure that you have the following books available for your reference.

The following Sun Cluster books support the Sun Cluster 3.1 release. If you are maintaining a different version of Sun Cluster software, refer to the appropriate documentation. All Sun Cluster documentation is available on <http://docs.sun.com>.

Documentation that is not available on <http://docs.sun.com> is listed with the appropriate URL.

TABLE P-2 Hardware Documentation

Title	Part Number
<i>Sun StorEdge Traffic Manager Installation and Configuration Guide</i> Available on http://www.sun.com/products-n-solutions/hardware/docs	816-1420
<i>Sun StorEdge T3 and T3+ Array Administrator's Guide</i> Available on http://www.sun.com/products-n-solutions/hardware/docs	816-0776
<i>Sun StorEdge T3 and T3+ Array Configuration Guide</i> Available on http://www.sun.com/products-n-solutions/hardware/docs	816-4771
<i>Sun StorEdge T3 and T3+ Array Installation, Operation, and Service Manual</i> Available on http://www.sun.com/products-n-solutions/hardware/docs	816-0773
<i>Sun StorEdge T3 Disk Tray Release Notes</i> Available on http://www.sun.com/products-n-solutions/hardware/docs	806-1497
<i>Sun StorEdge 3900 and 6900 Series Cabinet Installation and Service Manual</i>	816-3244
<i>Sun StorEdge 3900 and 6900 Series Reference Manual</i>	816-3245

TABLE P-3 Sun Cluster Documentation

Application	Title	Part Number
Concepts	<i>Sun Cluster 3.1 Concepts Guide</i>	816-3383
Hardware	<i>Sun Cluster 3.x Hardware Administration Manual</i> Sun Cluster 3.x Hardware Administration Collection at http://docs.sun.com/db/coll/1024.1/	817-0168
Software Installation	<i>Sun Cluster 3.1 Software Installation Guide</i>	816-3388
Data Services	<i>Sun Cluster 3.1 Data Service Planning and Administration Guide</i> Sun Cluster 3.1 Data Service Collection at http://docs.sun.com/db/coll/573.10/	817-1526
API Development	<i>Sun Cluster 3.1 Data Services Developer's Guide</i>	816-3385
Administration	<i>Sun Cluster 3.1 System Administration Guide</i>	816-3384

TABLE P-3 Sun Cluster Documentation (Continued)

Application	Title	Part Number
Error Messages	<i>Sun Cluster 3.1 Error Messages Guide</i>	816-3382
Man Pages	<i>Sun Cluster 3.1 Man Page Reference Manual</i>	816-5251
Release Notes	<i>Sun Cluster 3.1 Release Notes</i>	816-5317
	<i>Sun Cluster 3.1 Release Notes Supplement</i>	816-3381

Using UNIX Commands

This document contains information on commands used to install, configure, or upgrade a Sun Cluster configuration. This document might not contain complete 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 sources for this information.

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

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 number and serial number 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>/usr/cluster/bin/scinstall -pv</code>	Displays Sun Cluster release and package version information

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

Accessing Sun Documentation Online

The docs.sun.comSM Web site enables you to access Sun technical documentation online. You can browse the docs.sun.com archive or search for a specific book title or subject. The URL is `http://docs.sun.com`.

Typographic Conventions

The following table describes the typographic changes used in this book.

TABLE P-4 Typographic Conventions

Typeface or Symbol	Meaning	Example
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. <code>machine_name%</code> you have mail.
AaBbCc123	What you type, contrasted with on-screen computer output	<code>machine_name%</code> su Password:

TABLE P-4 Typographic Conventions (Continued)

Typeface or Symbol	Meaning	Example
<i>AaBbCc123</i>	Command-line placeholder: replace with a real name or value	To delete a file, type rm <i>filename</i> .
<i>AaBbCc123</i>	Book titles, new words, or terms, or words to be emphasized.	Read Chapter 6 in <i>User's Guide</i> . These are called <i>class</i> options. You must be <i>root</i> to do this.

Shell Prompts in Command Examples

The following table shows the default system prompt and superuser prompt for the C shell, Bourne shell, and Korn shell.

TABLE P-5 Shell Prompts

Shell	Prompt
C shell prompt	machine_name%
C shell superuser prompt	machine_name#
Bourne shell and Korn shell prompt	\$
Bourne shell and Korn shell superuser prompt	#

Installing and Maintaining a Sun StorEdge 3900 or 6900 Series System

This chapter contains the procedures about how to install, configure, and maintain Sun StorEdge™ 3900 Series and Sun StorEdge 6900 Series systems.

This chapter contains the following procedures.

- “How to Install a Storage System” on page 12
- “How to Configure a LUN” on page 15
- “How to Add a Storage System” on page 19
- “How to Remove a Storage System” on page 24
- “How to Replace a Node-to-Switch Component” on page 27
- “How to Replace a Virtualization Engine (StorEdge 6900 Series Only)” on page 26
- “How to Remove a Storage Array” on page 27
- “How to Upgrade Storage Array Firmware” on page 30
- “How to Replace a Disk Drive” on page 32
- “How to Replace a Node’s Host Adapter” on page 32

The storage system configuration utilities can be run from a menu-driven interface or a command line interface. This chapter describes the menu-driven interface.

For detailed information about storage system architecture, features, and configuration utilities, see the following manuals.

- *Sun StorEdge 3900 and 6900 Series Reference Manual*
- *Sun StorEdge 3900 and 6900 Series Installation and Service Manual*

Installing Storage Systems

Note – This section contains the procedure for an *initial installation* of storage systems in a new Sun™ Cluster that is not running. If you are adding systems to an existing cluster, use the procedure in “How to Add a Storage System” on page 19.

▼ How to Install a Storage System

Perform the steps in this procedure in conjunction with the procedures in the Sun Cluster software installation documentation and your server hardware manual.

1. Install host adapters in the nodes that to be connected to the system.

For the procedure about how to install host adapters, see the documentation that shipped with your host adapters and nodes.

2. Unpack, place, and level the system cabinet.

For instructions, see the *Sun StorEdge 3900 and 6900 Series Installation and Service Manual*.

3. Install cables in the following order.

a. Install the system power cord.

b. Install the system grounding strap.

c. Install the cables from the FC switches to the nodes as shown in Figure 1–1.

d. Install the Ethernet cable to the local area network (LAN).

For cabling instructions, see the *Sun StorEdge 3900 and 6900 Series Installation and Service Manual*.




FIGURE 1-1 Installing a Sun StorEdge 3900 or 6900 Series System

4. Power on the storage system and the nodes.

For the procedure about how to power on the system, see the *Sun StorEdge 3900 and 6900 Series Installation and Service Manual*. For the procedure about how to power on a node, see the documentation that shipped with your node hardware.

5. Set the host name, IP address, date, and time zone for the system's storage service processor.

For detailed instructions, see the initial field installation instructions in the *Sun StorEdge 3900 and 6900 Series Reference Manual*.

6. For StorEdge 3900 Series systems only: Reconfigure the StorEdge 3900 Series Fibre Channel switches. Ensure that every node can access each storage array

that is intended for that cluster.

7. **On all nodes, install the Solaris operating environment. Apply the required Solaris patches for Sun Cluster software and storage array support.**

For the procedure about how to install the Solaris operating environment, see your Sun Cluster software installation documentation.

See the *Sun Cluster release notes documentation* for information about accessing Sun's Info Docs on <http://sunsolve.sun.com>. The Info Docs list information about required patches and firmware that are available for download. For the procedure about how to apply host adapter firmware, see the firmware patch README file.

8. **Install any required patches or software for Sun StorEdge Traffic Manager software support on all nodes.**

For the procedure on how to install the Sun StorEdge Traffic Manager software, see the *Sun StorEdge Traffic Manager Installation and Configuration Guide* at <http://www.sun.com/products-n-solutions/hardware/docs/>. For a list of required patches for Sun StorEdge Traffic Manager software, see the *Sun StorEdge Traffic Manager Software Release Notes* on <http://www.sun.com/storage/san>.

9. **Edit the `/kernel/drv/scsi_vhci.conf` file to activate the Sun StorEdge Traffic Manager software functionality. Set the `mpxio-disable` parameter to `no`.**

```
mpxio-disable="no"
```

10. **Shutdown the entire cluster.**

For the procedure about how to shut down a cluster, see your Sun Cluster system administration documentation.

11. **Perform a reconfiguration boot on all nodes to create the new Solaris device files and links.**

12. **On all nodes, update the `/devices` and `/dev` entries.**

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

13. **On all nodes, update the paths to the DID instances.**

```
# scdidadm -C
# scdidadm -r
.
```

14. **For StorEdge 6900 Series systems only: On any node, view the virtualization engine (VE) controller status and enable the VE controllers.**

```
# cfgadm -al
# cfgadm -c configure c::controller id
```

See the `cfgadm(1M)` man page for more information about the command and the command's options.

15. **On all nodes, confirm that all storage arrays that you installed are now visible.**

```
# luxadm probe
```

Where to Go From Here

To continue with Sun Cluster software installation tasks, see your Sun Cluster software installation documentation.

Configuring Storage Systems

You can customize the configuration of the logical unit numbers (LUNs) on the storage arrays in your storage system. To customize the configuration of the LUNs, use the system's menu-driven or command-line interface. The *Sun StorEdge 3900 and 6900 Series Reference Manual* describes the factory default settings for the storage arrays in the storage systems.

For StorEdge 6900 Series systems only, you can also customize the configuration of the virtual LUNs (VLUNs) on the virtualization engines in the system. To customize the configuration of the VLUNs, use the system's menu-driven or command-line interface. The *Sun StorEdge 3900 and 6900 Series Reference Manual* describes the factory default settings for the virtualization engines in the StorEdge 6900 Series systems.

Note – When you upgrade firmware on a storage device or on an enclosure, redefine the stripe size of a LUN, or perform other LUN operations, a device ID (DID) might change unexpectedly. When you perform a check of the DID configuration by running the `scdidadm -c` command, the following error message appears on your console if the DID changed unexpectedly.

```
device id for nodename:/dev/rdsd/cXtYdZsN does not match physical
device's id for ddecimalnumber, device may have been replaced.
```

To fix DIDs that report this error, run the `scdidadm -R` command for each affected device.

How to Configure a LUN

Use the following summary of the steps to configure LUNs, or logical volumes, on the storage arrays in your storage system.

- Menu-driven Interface

1. On the storage service processor (SSP), start the Configuration Utilities.
2. Select T3+ Configuration Utility from the main menu.

3. Type your storage array password when you are prompted.
 4. Step through the submenus to configure the storage arrays. See the *Sun StorEdge 3900 and 6900 Series Reference Manual* for more information about the submenus.
- **Command-line Interface**
 1. On the SSP, type your storage array password if prompted.
 2. Use the SSP storage array commands to configure the storage arrays in your system. See the *Sun StorEdge 3900 and 6900 Series Reference Manual* and the man pages for more information.

How to Configure VLUNs on the Virtualization Engines (StorEdge 6900 Series Only)

The storage system's virtualization engines enable you to divide storage array LUNs into small virtual LUNs (VLUNs). You might want to customize storage usage in a storage area network (SAN).

Use the following summary of the steps to configure the VLUNs on the virtualization engines in your storage system.

- **Menu-driven Interface**
 1. On the SSP, start the Configuration Utilities.
 2. Select VE Configuration Utility from the main menu.
 3. Step through the submenus to configure the virtualization engines. See the *Sun StorEdge 3900 and 6900 Series Reference Manual* for more information about the submenus.
 4. After you configure the VLUNs, from any node update the global device IDs.


```
# scgdevs
```
 5. On one node connected to the partner-group, label the new VLUNs.


```
# format
```
- **Command-line Interface**
 1. Use the SSP virtualization engine commands to configure the virtualization engines in your system. See the *Sun StorEdge 3900 and 6900 Series Reference Manual* and the man pages for more information.
 2. After you configure the VLUNs, from any node update the global device IDs.


```
# scgdevs
```
 3. On one node connected to the partner-group, run the `format` command to label the new VLUNs.


```
# format
```

Maintaining Storage Systems

This section contains the procedures about how to maintain storage systems. Table 1–1 lists these procedures. This section does not include procedures about how to add or remove disk drives. Storage arrays in your storage system operate only when fully configured with disk drives.



Caution – If you remove any field replaceable unit (FRU) from the storage arrays for an extended period of time, thermal complications might result. To prevent these complications, the storage array is designed so that an orderly shutdown occurs. This shutdown occurs when you remove a component for longer than 30 minutes. Therefore, a replacement part must be immediately available before you start an FRU replacement procedure. You must replace an FRU within 30 minutes. If you do not, the storage array, and all attached storage arrays, shut down and power off.

Note – When you upgrade firmware on a storage device or on an enclosure, redefine the stripe size of a LUN, or perform other LUN operations, a device ID (DID) might change unexpectedly. When you perform a check of the DID configuration by running the `scdidadm -c` command, the following error message appears on your console if the DID changed unexpectedly.

```
device id for nodename:/dev/rdisk/cXtYdZsN does not match physical
device's id for ddecimalnumber, device may have been replaced.
```

To fix DIDs that report this error, run the `scdidadm -R` command for each affected device.

TABLE 1–1 Task Map: Maintaining a Storage System

Task	Information
Add a storage system.	"How to Add a Storage System" on page 19
Remove a storage system.	"How to Remove a Storage System" on page 24
Replace a virtualization engine. This procedure applies to a StorEdge 6900 Series only.	"How to Replace a Virtualization Engine (StorEdge 6900 Series Only)" on page 26

TABLE 1–1 Task Map: Maintaining a Storage System (Continued)

Task	Information
Replace a node-to-switch fiber-optic cable.	"How to Replace a Node-to-Switch Component" on page 27
Replace a gigabit interface converter (GBIC) or Small Form Factor Pluggable (SFP) on a node's host adapter.	"How to Replace a Node-to-Switch Component" on page 27
Replace a GBIC or an SFP on an FC switch, connecting to a node.	"How to Replace a Node-to-Switch Component" on page 27
Remove a storage array.	"How to Remove a Storage Array" on page 27
Upgrade storage array firmware.	"How to Upgrade Storage Array Firmware" on page 30
Replace a disk drive.	"How to Replace a Disk Drive" on page 32
Replace a node's host adapter.	"How to Replace a Node's Host Adapter" on page 32
Add a node to the storage array.	Sun Cluster system administration documentation
Remove a node from the storage array.	Sun Cluster system administration documentation

The following is a list of administrative tasks that require no cluster-specific procedures.

See the *Sun StorEdge T3 and T3+ Array Installation, Operation, and Service Manual* for the following procedures.

- Replacing an Ethernet cable..
- Replacing a power cable on the storage array.
- Replacing a unit interconnect card (UIC).
- Replacing a Power and Cooling Unit (PCU).

See the *Sun StorEdge 3900 and 6900 Series Reference Manual* for the following procedures.

- Replacing an Ethernet hub.
- Replacing a storage service processor.
- Replacing a GBIC or an SFP on an FC switch that connects to a storage array.
- Replacing a storage array-to-switch fiber-optic cable.
- Replacing an FC switch.
- Removing a Virtualization Engine.
- Replacing a Virtualization Engine.
- Upgrading a Virtualization Engine firmware.

▼ How to Add a Storage System

Note – Use this procedure to add new storage systems to a running cluster. To install systems to a new Sun Cluster that is not running, use the procedure in “How to Install a Storage System” on page 12.

This procedure defines Node A as the node with which you begin working. Node B is the remaining node.

1. Unpack, place, and level the system cabinet.

For instructions, see the *Sun StorEdge 3900 and 6900 Series Installation and Service Manual*.

2. Install cables in the following order.

a. Install the system power cord.

b. Install the system grounding strap.

c. Install the cables from the FC switches to the nodes as shown in Figure 1–1.

d. Install the Ethernet cable to the LAN.

For cabling instructions, see the *Sun StorEdge 3900 and 6900 Series Installation and Service Manual*.

3. Power on the new system.

Note – The storage arrays in your system might require several minutes to boot.

For instructions, see the *Sun StorEdge 3900 and 6900 Series Installation and Service Manual*.

4. Set the host name, IP address, date, and time zone for the system’s Storage Service Processor.

For instructions, see the initial field installation instructions in the *Sun StorEdge 3900 and 6900 Series Reference Manual*.

5. For StorEdge 3900 Series systems only: Reconfigure the StorEdge 3900 Series Fibre Channel switches. Ensure that every node can access each storage array that is intended for that cluster.

6. Determine the resource groups and device groups running on all nodes.

Record this information because you use this information in Step 47 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 from-node
```

8. Do you need to install host adapters in Node A?

- If no, skip to Step 14.
- If yes, proceed to Step 9.

9. Is the host adapter that you are installing the first host adapter on Node A?

- If no, skip to Step 11.
- If yes, determine whether the required support packages for the host adapter are already installed on this node. The following packages are required.

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

10. Are the required support packages already installed?

- If yes, skip to Step 11.
- If no, install the required support packages that are missing.

The support packages are located in the `Product` directory of the Solaris CD-ROM.

11. Shut down and power off Node A.

For the procedure about how to shut down and power off a node, see your Sun Cluster system administration documentation.

12. Install the host adapters in Node A.

For the procedure about how to install host adapters, see the documentation that shipped with your host adapters and nodes.

13. Power on and boot Node A into noncluster mode.

For the procedure about how to boot a node in noncluster mode, see your Sun Cluster system administration documentation.

14. If necessary, upgrade the host adapter firmware on Node A.

See the *Sun Cluster release notes documentation* for information about accessing Sun's Info Docs on <http://sunsolve.sun.com>. The Info Docs list information about required patches and firmware that are available for download. For the procedure about how to apply host adapter firmware, see the firmware patch README file.

15. If necessary, install GBICs or SFPs in the FC switches.

For the procedure about how to install a GBIC or an SFP to an FC switch, see the documentation that shipped with your FC switch hardware.

16. Connect fiber-optic cables between Node A and the FC switches in your storage system as shown in Figure 1–1.

17. If necessary, install the required Solaris patches for storage array support on Node A.

See the *Sun Cluster release notes documentation* for information about accessing Sun's Info Docs on <http://sunsolve.sun.com>. The Info Docs list information about required patches and firmware that are available for download. For the procedure about how to apply host adapter firmware, see the firmware patch README file.

18. Install any required patches or software for Sun StorEdge Traffic Manager software support on Node A.

For the procedure on how to install the Sun StorEdge Traffic Manager software, see the *Sun StorEdge Traffic Manager Installation and Configuration Guide* at <http://www.sun.com/products-n-solutions/hardware/docs/>. For a list of required patches for Sun StorEdge Traffic Manager software, see the *Sun StorEdge Traffic Manager Software Release Notes* on <http://www.sun.com/storage/san>.

19. Edit the `/kernel/drv/scsi_vhci.conf` file to activate the Sun StorEdge Traffic Manager software functionality. Set the `mpxio-disable` parameter to `no`.

```
mpxio-disable="no"
```

20. Shut down Node A.

For the procedure about how to shut down and power off a node, see your Sun Cluster system administration documentation.

21. Perform a reconfiguration boot on Node A to create the new Solaris device files and links.

22. On all nodes, update the `/devices` and `/dev` entries.

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

23. On all nodes, update the paths to the DID instances

```
# scdidadm -C
# scdidadm -r
.
```

24. Are you adding a StorEdge 3900 Series or StorEdge 6900 Series system?

- If you are adding a StorEdge 3900 Series system, skip to Step 26.
- If you are adding StorEdge 6900 Series systems: On Node A, view the virtualization engine (VE) controller status and enable the VE controllers.

```
# cfgadm -al
# cfgadm -c configure c::controller id
```

25. (Optional) Configure VLUNs on the VEs in the new StorEdge 6900 Series system.

For the procedure about how to configure VLUNs in a cluster, see “How to Configure VLUNs on the Virtualization Engines (StorEdge 6900 Series Only)” on page 16.

26. (Optional) On Node A, verify that the device IDs (DIDs) were assigned to the new storage array.

```
# scdidadm -l
```

27. Do you need to install host adapters in Node B?

- If no, skip to Step 33.
- If yes, proceed to Step 28.

28. Is the host adapter that you are installing the first host adapter on Node B?

- If no, skip to Step 30.
- If yes, determine whether the required support packages for the host adapter are already installed on this node. The following packages are required.

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

29. Are the required support packages already installed?

- If yes, proceed to Step 30.
- If no, install the missing support packages.

The support packages are located in the Product directory of the Solaris CD-ROM.

30. Shut down and power off Node B.

For the procedure about how to shut down and power off a node, see your Sun Cluster system administration documentation.

31. Install the host adapters in Node B.

For the procedure about how to install host adapters, see the documentation that shipped with your host adapters and nodes.

32. Power on and boot Node B into noncluster mode.

For the procedure about how to shut down and power off a node, see your Sun Cluster system administration documentation.

33. If necessary, upgrade the host adapter firmware on Node B.

See the *Sun Cluster release notes documentation* for information about accessing Sun's Info Docs on <http://sunsolve.sun.com>. The Info Docs list information about required patches and firmware that are available for download. For the procedure about how to apply host adapter firmware, see the firmware patch README file.

34. If necessary, install GBICs or SFPs in the FC switches.

For the procedure about how to install a GBIC or an SFP to an FC switch, see the documentation that shipped with your FC switch hardware.

35. Connect fiber-optic cables between the FC switches in your storage system and Node B as shown in Figure 1-1.

36. If necessary, install the required Solaris patches for storage array support on Node B.

For a list of required Solaris patches for storage array support, see the *Sun StorEdge T3 Disk Tray Release Notes*.

37. Install any required patches or software for Sun StorEdge Traffic Manager software support on Node B.

For the procedure on how to install the Sun StorEdge Traffic Manager software, see the *Sun StorEdge Traffic Manager Installation and Configuration Guide* at <http://www.sun.com/products-n-solutions/hardware/docs/>. For a list of required patches for Sun StorEdge Traffic Manager software, see the *Sun StorEdge Traffic Manager Software Release Notes* on <http://www.sun.com/storage/san>.

38. Edit the `/kernel/drv/scsi_vhci.conf` file to activate the Sun StorEdge Traffic Manager software functionality. Set the `mpxio-disable` parameter to `no`.

```
mpxio-disable="no"
```

39. Shut down Node B.

For the procedure about how to shut down and power off a node, see your Sun Cluster system administration documentation.

40. Perform a reconfiguration boot to create the new Solaris device files and links on Node B.

41. On all nodes, update the `/devices` and `/dev` entries.

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

42. On all nodes, update the paths to the DID instances.

```
# scdidadm -C
# scdidadm -r
.
```

43. Are you adding a StorEdge 3900 Series or StorEdge 6900 Series system?

- If you are adding a StorEdge 3900 Series system, skip to Step 45.
- If you are adding StorEdge 6900 Series systems: On Node A, view the virtualization engine (VE) controller status and enable the VE controllers.

```
# cfgadm -al
# cfgadm -c configure c::controller id
```

44. (Optional) Configure VLUNs on the VEs in the new StorEdge 6900 Series system.

For the procedure about how to configure VLUNs in a cluster, see “How to Configure VLUNs on the Virtualization Engines (StorEdge 6900 Series Only)” on page 16.

45. (Optional) On Node B, verify that the DIDs are assigned to the new storage arrays.

```
# scdidadm -l
```

46. On one node attached to the new storage arrays, reset the SCSI reservation state.

```
# scdidadm -R n
```

Where *n* is the DID instance of a storage array LUN that you are adding to the cluster.

Note – Repeat this command on the same node for each storage array LUN that you are adding to the cluster.

47. Return the resource groups and device groups that you identified in Step 6 to all nodes.

```
# scswitch -z -g resource-group -h nodename
# scswitch -z -D device-group-name -h nodename
```

For more information, see your Sun Cluster system administration documentation.

48. 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 Storage System

Use this procedure to permanently remove a storage system and the associated submirrors from a running cluster.

This procedure defines Node A as the node with which you begin working. Node B is the remaining node.



Caution – During this procedure, you lose access to the data that resides on each storage array partner-group. This partner-group is in the storage system that you are removing.

1. **If necessary, back up all database tables, data services, and volumes associated with each storage array partner-group. This partner-group is in the storage system that you are removing.**
2. **If necessary, run the appropriate Solstice DiskSuite commands or VERITAS Volume Manager commands to detach the submirrors from each storage array partner-group. This partner-group is in the storage system that you are removing. Detach the submirrors to stop all I/O activity to the partner-groups.**
For more information, see your Solstice DiskSuite or VERITAS Volume Manager documentation.
3. **Run the appropriate Solstice DiskSuite commands or VERITAS Volume Manager commands to remove references to each LUN or VLUN. Each LUN or VLUN is in the storage system that you are removing.**
For more information, see your Solstice DiskSuite or VERITAS Volume Manager documentation.
4. **Determine the resource groups and device groups running on all nodes.**
Record this information because you use this information in Step 15 of this procedure to return resource groups and device groups to these nodes.

```
# scstat
```
5. **Shut down Node A.**
For the procedure about how to shut down and power off a node, see your Sun Cluster system administration documentation.
6. **Disconnect the cables that connect Node A to the FC switches in your storage system.**
7. **Without enabling the node to boot, power on Node A.**
For more information, see your Sun Cluster system administration documentation.
8. **Boot Node A into cluster mode.**
For more information about how to boot nodes, see your Sun Cluster system administration documentation.
9. **Shut down Node B.**

For the procedure about how to shut down and power off a node, see your Sun Cluster system administration documentation.

10. Disconnect the cables that connect Node B to the FC switches.

11. Without enabling the node to boot, power on Node B.

For more information, see your Sun Cluster system administration documentation.

12. Boot Node B into cluster mode.

For more information on booting nodes, see your Sun Cluster system administration documentation.

13. On all nodes, update the `/devices` and `/dev` entries.

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

14. On all nodes, update the paths to the DID instances.

```
# scdidadm -C
# scdidadm -r
.
```

15. Return the resource groups and device groups that you identified in Step 4 to all nodes.

```
# scswitch -z -g resource-group -h nodename
# scswitch -z -D device-group-name -h nodename
```

▼ How to Replace a Virtualization Engine (StorEdge 6900 Series Only)

Use this procedure to replace a virtualization engine (VE) in a storage system in a running cluster.

1. Replace the VE hardware.

Follow the instructions in the *Sun StorEdge 3900 and 6900 Series Reference Manual*.

2. On any node, view the virtualization engine (VE) controller status and enable the VE controllers.

```
# cfgadm -al
# cfgadm -c configure c::controller id
```

▼ How to Replace a Node-to-Switch Component

Use this procedure to replace the following node-to-switch components in a running cluster. This procedure defines Node A as the node connected to the node-to-switch component you are replacing. Node B is the remaining node node-to-switch connected to the component you are replacing.

- Node-to-switch fiber-optic cable
- GBIC or an SFP on an FC switch, connecting to a node

1. Determine the resource groups and device groups running on Node A.

Record this information because you use this information in Step 4 of this procedure to return resource groups and device groups to these nodes.

```
# scstat
```

2. Move all resource groups and device groups to Node B.

```
# scswitch -S -h from-node
```

3. Replace the node-to-switch component.

- For the procedure about how to replace a fiber-optic cable between a node and an FC switch, see the documentation that shipped with your switch hardware.
- For the procedure about how to replace a GBIC or an SFP on an FC switch, see the documentation that shipped with your FC switch hardware.

4. Return the resource groups and device groups that you identified in Step 1 to Node A.

```
# scswitch -z -g resource-group -h nodename  
# scswitch -z -D device-group-name -h nodename
```

▼ How to Remove a Storage Array

Use this procedure to permanently remove storage array partner groups and their submirrors from a storage system. This storage system is in a running cluster.

This procedure defines Node A as the node with which you begin working. Node B is the remaining node.



Caution – During this procedure, you lose access to the data that resides on each storage array partner-group you are removing.

1. If necessary, back up all database tables, data services, and volumes associated with each partner-group you are removing.

2. **If necessary, run the appropriate Solstice DiskSuite commands or VERITAS Volume Manager commands to detach the submirrors from each storage array or partner-group that you are removing. Detach the submirrors to stop all I/O activity to the storage array or partner-group.**

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

3. **Run the appropriate Solstice DiskSuite commands or VERITAS Volume Manager commands to remove references to each LUN. Each LUN belongs to the storage array or partner-group that you are removing.**

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

4. **Determine the resource groups and device groups running on all nodes.**

Record this information because you use this information in Step 20 of this procedure to return resource groups and device groups to these nodes.

```
# scstat
```

5. **Shut down Node A.**

For the procedure about how to shut down and power off a node, see your Sun Cluster system administration documentation.

6. **Disconnect the fiber-optic cables that connect the storage arrays and the FC switches. Then disconnect the Ethernet cables.**

7. **Is any storage array that you are removing the last storage array connected to an FC switch that is on Node A?**

- If no, skip to Step 11.
- If yes, disconnect the fiber-optic cable between Node A and the FC switch that was connected to this storage array.

For the procedure about how to remove a fiber-optic cable, see the *Sun StorEdge T3 and T3+ Array Configuration Guide*.

8. **Do you want to remove the host adapters from Node A?**

- If no, skip to Step 11.
- If yes, power off Node A.

9. **Remove the host adapters from Node A.**

For the procedure about how to remove host adapters, see the documentation that shipped with your host adapter and nodes.

10. **Without enabling the node to boot, power on Node A.**

For more information, see your Sun Cluster system administration documentation.

11. **Boot Node A into cluster mode.**

For more information on booting nodes, see your Sun Cluster system administration documentation.

12. Shut down Node B.

For the procedure about how to shut down and power off a node, see your Sun Cluster system administration documentation.

13. Is any storage array that you are removing the last storage array connected to an FC switch that is on Node B?

- If no, skip to Step 17.
- If yes, disconnect the fiber-optic cable that connects this FC switch to Node B.

For the procedure about how to remove a fiber-optic cable, see the *Sun StorEdge T3 and T3+ Array Configuration Guide*.

14. Do you want to remove the host adapters from Node B?

- If no, skip to Step 17.
- If yes, power off Node B.

15. Remove the host adapters from Node B.

For the procedure about how to remove host adapters, see the documentation that shipped with your nodes.

16. Without enabling the node to boot, power on Node B.

For more information, see your Sun Cluster system administration documentation.

17. Boot Node B into cluster mode.

For more information on booting nodes, see your Sun Cluster system administration documentation.

18. On all nodes, update the `/devices` and `/dev` entries.

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

19. On all nodes, update the paths to the DID instances.

```
# scdidadm -C
# scdidadm -r
.
```

20. Return the resource groups and device groups that you identified in Step 4 to all nodes.

```
# scswitch -z -g resource-group -h nodename
# scswitch -z -D device-group-name -h nodename
```

How to Upgrade Storage Array Firmware

Use one of the following procedures to upgrade the firmware on the storage arrays in your storage system. The procedure that you choose depends on whether your storage array partner-group has been configured to support submirrors of a node's volumes. Storage array firmware includes controller firmware, unit interconnect card (UIC) firmware, EPROM firmware, and disk drive firmware.

- "Upgrading Firmware on Storage Arrays That Support Submirrored Data" on page 30
- "Upgrading Firmware on Storage Arrays That Do Not Support Submirrored Data" on page 31

Note – For all firmware, always read any README files that accompany the firmware for the latest information and special notes.

▼ Upgrading Firmware on Storage Arrays That Support Submirrored Data

This procedure defines Node A as the node on which you are upgrading firmware. Node B is the remaining node.



Caution – Perform this procedure on one storage array at a time. This procedure requires you to reset the storage arrays that you are upgrading. If you reset more than one storage array at a time, your cluster loses access to data.

1. **On the node that currently owns the disk group or diskset to which the submirror belongs, detach the submirrors of the storage array. This storage array is the storage array on which you are upgrading firmware.**

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

2. **Disconnect both storage array-to-switch fiber-optic cables from the two storage arrays of the partner-group.**

3. **Apply the controller, disk drive, and UIC firmware patches.**

For the list of required storage array patches, see the *Sun StorEdge 3900 and 6900 Series Reference Manual*. For the procedure about how to apply firmware patches, see the firmware patch README file. For the procedure about how to verify the firmware level, see the *Sun StorEdge 3900 and 6900 Series Reference Manual*.

4. **If you have not already done so, reset the storage arrays.**

For the procedure about how to reset a storage array, see the *Sun StorEdge T3 and T3+ Array Installation, Operation, and Service Manual*.

5. **Disable the storage array controller that is attached to Node B. Disable the storage array controller so that all logical volumes are managed by the remaining controller.**

For more information, see the *Sun StorEdge T3 and T3+ Array Administrator's Guide*.

6. **Reconnect both storage array-to-switch fiber-optic cables to the two storage arrays of the partner-group.**
7. **On one node connected to the partner-group, verify that the storage array controllers are visible to the node.**

```
# format
```

8. **Enable the storage array controller that you disabled in Step 5.**
9. **Reattach the submirrors that you detached in Step 1 to resynchronize the submirrors.**

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

▼ Upgrading Firmware on Storage Arrays That Do Not Support Submirrored Data

In a partner-pair configuration, you can have nonmirrored data. However, this configuration requires you to shut down the cluster when upgrading firmware.

1. **Shutdown the entire cluster.**

For the procedure about how to shut down a cluster, see your Sun Cluster system administration documentation.

2. **Apply the controller, disk drive, and UIC firmware patches.**

For the list of required storage array patches, see the *Sun StorEdge 3900 and 6900 Series Reference Manual*. For the procedure about how to apply firmware patches, see the firmware patch README file. For the procedure about how to verify the firmware level, see the *Sun StorEdge 3900 and 6900 Series Reference Manual*.

3. **If you have not already done so, reset the storage arrays.**

For the procedure about how to reset a storage array, see the *Sun StorEdge T3 and T3+ Array Installation, Operation, and Service Manual*.

4. **Boot all nodes back into the cluster.**

For more information on booting nodes, see your Sun Cluster system administration documentation.

5. **On one node connected to the partner-group, verify that the storage array controllers are visible to the node.**

```
# format
```

▼ How to Replace a Disk Drive

Use this procedure to replace a failed disk drive in a storage array in your storage system. This storage system is 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 storage array is designed so that an orderly shutdown occurs. This shutdown occurs when you remove a component for longer than 30 minutes. Therefore, a replacement part must be immediately available before starting an FRU replacement procedure. You must replace an FRU within 30 minutes. If you do not, the storage array, and all attached storage arrays, shut down and power off.

1. Did the failed disk drive impact the storage array LUN's availability?

- If no, proceed to Step 2.
- If yes, remove the LUN from volume management control.

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

2. Replace the disk drive in the storage array.

For the procedure about how to replace a disk drive, see the *Sun StorEdge T3 and T3+ Array Installation, Operation, and Service Manual*.

3. Did you remove a LUN from volume management control in Step 1?

- If no, this step completes this procedure.
- If yes, return the LUN to volume management control now.

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

▼ How to Replace a Node's Host Adapter

Use this procedure to replace a failed host adapter in a running cluster. This procedure defines Node A as the node with the failed host adapter you are replacing. Node B is the remaining node.

1. Determine the resource groups and device groups running on all nodes.

Record this information because you use this information in Step 6 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 from-node
```

Shut down and then power off Node A.

For the procedure about how to shut down and power off a node, see your Sun Cluster system administration documentation.

3. Replace the failed host adapter.

For the procedure about how to remove and add host adapters, see the documentation that shipped with your nodes.

4. Power on Node A.

For more information, see your Sun Cluster system administration documentation.

5. Boot Node A into cluster mode.

For more information on booting nodes, see your Sun Cluster system administration documentation.

6. Return the resource groups and device groups that you identified in Step 1 to all nodes.

```
# scswitch -z -g resource-group -h nodename
```

```
# scswitch -z -D device-group-name -h nodename
```

For more information, see your Sun Cluster system administration documentation.

Index

A

adding
 See also installing
 nodes, 18
 systems, 19
arrays, *See* storage arrays

C

configuring
 LUNs, 15
 VLUNs, 16
cooling units, replacing, 18

D

disk drives, replacing, 32

E

Ethernet cables, replacing, 18
Ethernet hubs, replacing, 18

G

GBICs, replacing, 18

H

help, 8
host adapters, replacing, 32

I

installing
 See also adding
 systems, 12

L

logical unit numbers, *See* LUNs
logical units, LUNs, 15
logical volumes, *See* LUNs
LUNs, configuring, 15

N

node-to-switch components, replacing, 27
nodes
 adding, 18
 removing, 18

P

power cables, replacing, 18
power units, replacing, 18

R

removing

- nodes, 18
- storage arrays, 27
- systems, 24
- Virtualization Engines, 18

replacing

- cooling units, 18
- disk drives, 32
- Ethernet cables, 18
- Ethernet hubs, 18
- GBICs, 18
- host adapters, 32
- node-to-switch components, 27
- power cables, 18
- power units, 18
- SFPs, 18
- storage array-to-switch fiber-optic cables, 18
- storage service processor, 18
- switches, 18
- unit interconnect cards, 18
- Virtualization Engines, 18
- virtualization engines, 26

S

SFPs, replacing, 18

storage array firmware, upgrading, 30

storage array-to-switch fiber-optic cables,
replacing, 18

storage arrays

- See* systems
- removing, 27

storage service processor, replacing, 18

switches, replacing, 18

systems

- adding, 19
- installing, 12
- removing, 24

T

technical support, 8

U

unit interconnect cards, replacing, 18

upgrading, storage array firmware, 30

upgrading firmware, Virtualization Engines, 18

V

Virtualization Engines

removing, 18

replacing, 18

virtualization engines, replacing, 26

Virtualization Engines

upgrading firmware, 18

VLUNs, configuring, 16