

Sun Cluster 3.1 - 3.2 With Sun StorEdge 3510 or 3511 FC RAID Array Manual for Solaris OS



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

Part No: 819-3016-12
January 2009, Revision A

Copyright 2009 Sun Microsystems, Inc. 4150 Network Circle, Santa Clara, CA 95054 U.S.A. All rights reserved.

Sun Microsystems, Inc. has intellectual property rights relating to technology embodied in the product that is described in this document. In particular, and without limitation, these intellectual property rights may include one or more U.S. patents or pending patent applications in the U.S. and in other countries.

U.S. Government Rights – Commercial software. Government users are subject to the Sun Microsystems, Inc. standard license agreement and applicable provisions of the FAR and its supplements.

This distribution may include materials developed by third parties.

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, the Solaris logo, the Java Coffee Cup logo, docs.sun.com, Sun StorEdge, Netra, OpenBoot, Solstice DiskSuite, UNIX, Java, and Solaris are trademarks or registered trademarks of Sun Microsystems, Inc. or its subsidiaries 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.

Products covered by and information contained in this publication are controlled by U.S. Export Control laws and may be subject to the export or import laws in other countries. Nuclear, missile, chemical or biological weapons or nuclear maritime end uses or end users, whether direct or indirect, are strictly prohibited. Export or reexport to countries subject to U.S. embargo or to entities identified on U.S. export exclusion lists, including, but not limited to, the denied persons and specially designated nationals lists is strictly prohibited.

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 2009 Sun Microsystems, Inc. 4150 Network Circle, Santa Clara, CA 95054 U.S.A. Tous droits réservés.

Sun Microsystems, Inc. détient les droits de propriété intellectuelle relatifs à la technologie incorporée dans le produit qui est décrit dans ce document. En particulier, et ce sans limitation, ces droits de propriété intellectuelle peuvent inclure un ou plusieurs brevets américains ou des applications de brevet en attente aux Etats-Unis et dans d'autres pays.

Cette distribution peut comprendre des composants développés par des tierces personnes.

Certains composants de ce produit peuvent être dérivées du logiciel Berkeley BSD, licenciés par l'Université de Californie. UNIX est une marque déposée aux Etats-Unis et dans d'autres pays; elle est licenciée exclusivement par X/Open Company, Ltd.

Sun, Sun Microsystems, le logo Sun, le logo Solaris, le logo Java Coffee Cup, docs.sun.com, Sun StorEdge, Netra, OpenBoot, Solstice DiskSuite, UNIX, Java et Solaris sont des marques de fabrique ou des marques déposées de Sun Microsystems, Inc., ou ses filiales, 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.

Les produits qui font l'objet de cette publication et les informations qu'il contient sont régis par la législation américaine en matière de contrôle des exportations et peuvent être soumis au droit d'autres pays dans le domaine des exportations et importations. Les utilisations finales, ou utilisateurs finaux, pour des armes nucléaires, des missiles, des armes chimiques ou biologiques ou pour le nucléaire maritime, directement ou indirectement, sont strictement interdites. Les exportations ou réexportations vers des pays sous embargo des Etats-Unis, ou vers des entités figurant sur les listes d'exclusion d'exportation américaines, y compris, mais de manière non exclusive, la liste de personnes qui font objet d'un ordre de ne pas participer, d'une façon directe ou indirecte, aux exportations des produits ou des services qui sont régis par la législation américaine en matière de contrôle des exportations et la liste de ressortissants spécifiquement désignés, sont rigoureusement interdites.

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 CONTREFACON.

Contents

Preface	5
1 Installing and Maintaining Sun StorEdge 3510 and 3511 Fibre Channel RAID Arrays	11
Installing Storage Arrays	11
Storage Array Cabling Configurations	11
▼ How to Install a Storage Array	15
Adding a Storage Array to a Running Cluster	17
Configuring Storage Arrays in a Running Cluster	22
▼ How to Create and Map a LUN	23
▼ How to Unmap and Remove a LUN	24
Maintaining Storage Arrays	26
StorEdge 3510 and 3511 FC RAID Array FRUs	28
▼ How to Remove a Storage Array From a Running Cluster	28
▼ How to Upgrade Storage Array Firmware	31
▼ How to Replace a Disk Drive	32
▼ How to Replace a Host Adapter	33
Replacing a Node-to-Switch Component	36
▼ How to Replace a Chassis in a Running Cluster	38
Index	41

Preface

The *Sun Cluster 3.1 - 3.2 With Sun StorEdge 3510 or 3511 FC RAID Array Manual for Solaris OS* provides procedures specific to Sun StorEdge™ 3510 and 3511 FC RAID storage arrays placed in a Sun™ Cluster environment.

Use this manual with any version of Sun Cluster 3.1 or 3.2 software on SPARC® based clusters and with any supported version of Sun Cluster 3.1 or 3.2 on x86 based clusters. Unless otherwise noted, procedures are the same for all Sun Cluster 3.1 or 3.2 versions. See the “[Revision History](#)” on page 6 for a list of changes to this manual.

See the *Sun Cluster 3.0-3.1 Release Notes Supplement* for current information on bugs or restrictions.

Note – This Sun Cluster release supports systems that use the SPARC and x86 families of processor architectures: UltraSPARC, SPARC64, and AMD64. In this document, the label x86 refers to systems that use the AMD64 family of processor architectures. The information in this document pertains to both platforms unless otherwise specified in a special chapter, section, note, bulleted item, figure, table, or example.

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 a pre-sales 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

Chapter 1, “Installing and Maintaining Sun StorEdge 3510 and 3511 Fibre Channel RAID Arrays,” discusses how to install, configure, and maintain the RAID version of the StorEdge 3510 and 3511 storage arrays.

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 3.1 - 3.2 With Sun StorEdge 3510 or 3511 FC RAID Array Manual for Solaris OS

Revision Date	New Information
March 2008	Replaced outdated information about SunSolve with information about Sun Connection Update Manager .
January 2009	Updated table in Preface in to include different versions of Sun Cluster documentation.

Related Documentation

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

The following Sun Cluster books support the Sun Cluster 3.1 and 3.2 releases. If you are maintaining a different version of Sun Cluster software, refer to the appropriate documentation. All Sun Cluster documentation is available at <http://docs.sun.com>. Documentation that is not available at <http://docs.sun.com> is listed with the appropriate URL.

TABLE P-2 Hardware Documentation

Title	Part Number
<i>Sun StorEdge 3000 Family Installation, Operation, and Service Manual, Sun StorEdge 3510 FC Array</i>	816-7300
<i>Sun StorEdge 3000 Family RAID Firmware 4.1x User's Guide</i>	817-3711
<i>Sun StorEdge 3000 Family Best Practices Manual</i>	816-7325

TABLE P-2 Hardware Documentation (Continued)

Title	Part Number
<i>Sun StorEdge 3510 FC Array Release Notes</i>	816-7301
<i>Sun StorEdge 3511 FC Array with SATA Release Notes</i>	817-6597
<i>Sun StorEdge 3000 Family FRU Installation Guide</i>	816-7326
<i>Sun StorEdge 3000 Family Rack Installation Guide for 2U Arrays</i>	817-3629

TABLE P-3 Sun Cluster Documentation

Documentation
Solaris Cluster 3.2
Sun Cluster 3.1

Using UNIX Commands

This document contains information about commands that are used to install, configure, or upgrade a Sun Cluster configuration. This document might not contain complete information about 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™ Operating System (Solaris OS)
- Other software documentation that you received with your system
- Solaris Operating System 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 10)
- The release number of Sun Cluster (for example, Sun Cluster 3.2)

Use the following commands to gather information about 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/clnode show-rev</code> <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.

Documentation, Support, and Training

The Sun web site provides information about the following additional resources:

- [Documentation \(http://www.sun.com/documentation/\)](http://www.sun.com/documentation/)
- [Support \(http://www.sun.com/support/\)](http://www.sun.com/support/)
- [Training \(http://www.sun.com/training/\)](http://www.sun.com/training/)

Sun Welcomes Your Comments

Sun is interested in improving its documentation and welcomes your comments and suggestions. To share your comments, go to <http://docs.sun.com> and click Feedback.

Typographic Conventions

The following table describes the typographic conventions that are used in this book.

TABLE P-4 Typographic Conventions

Typeface	Meaning	Example
AaBbCc123	The names of commands, files, and directories, and onscreen computer output	Edit your <code>.login</code> file. Use <code>ls -a</code> to list all files. <code>machine_name% you have mail.</code>

TABLE P-4 Typographic Conventions (Continued)

Typeface	Meaning	Example
AaBbCc123	What you type, contrasted with onscreen computer output	machine_name% su Password:
<i>aabbcc123</i>	Placeholder: replace with a real name or value	The command to remove a file is <i>rm filename</i> .
<i>AaBbCc123</i>	Book titles, new terms, and terms to be emphasized	Read Chapter 6 in the <i>User's Guide</i> . A <i>cache</i> is a copy that is stored locally. Do <i>not</i> save the file. Note: Some emphasized items appear bold online.

Shell Prompts in Command Examples

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

TABLE P-5 Shell Prompts

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

Installing and Maintaining Sun StorEdge 3510 and 3511 Fibre Channel RAID Arrays

This chapter describes the procedures for installing, configuring, and maintaining the Sun™ StorEdge™ 3510 FC RAID array and the Sun StorEdge 3511 FC RAID array with SATA in a Sun Cluster environment. This chapter contains the following main topics:

- “Installing Storage Arrays” on page 11
- “Configuring Storage Arrays in a Running Cluster” on page 22
- “Maintaining Storage Arrays” on page 26

Before you perform any of the tasks in this chapter, read the entire procedure. If you are not reading an online version of this document, have the books listed in the [Preface](#) available.

For conceptual information on multihost disks, see the Sun Cluster concepts documentation.

Installing Storage Arrays

This section contains the procedures listed in [Table 1-1](#)

TABLE 1-1 Task Map: Installing Storage Arrays

Task	Information
Install a storage array in a new cluster, before the OS and Sun Cluster software are installed.	“How to Install a Storage Array” on page 15
Add a storage array to an existing cluster.	“Adding a Storage Array to a Running Cluster” on page 17

Storage Array Cabling Configurations

You can install the StorEdge 3510 and 3511 FC RAID arrays in several different configurations. Use the *Sun StorEdge 3000 Family Best Practices Manual* to help evaluate your needs and

determine which configuration is best for your situation. See your Sun service provider for currently supported Sun Cluster configurations.

The following figures provide examples of configurations with multipathing solutions. With direct attach storage (DAS) configurations with multipathing, you map each LUN to each host channel. All nodes can see all 256 LUNs.

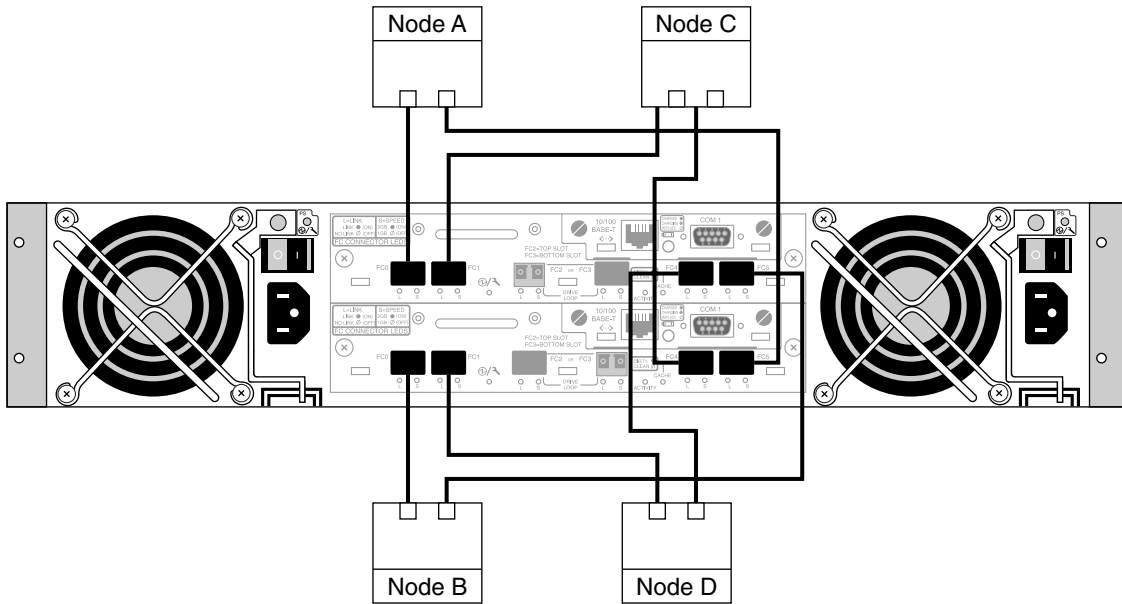


FIGURE 1-1 Sun StorEdge 3510 DAS Configuration With Multipathing and Two Controllers

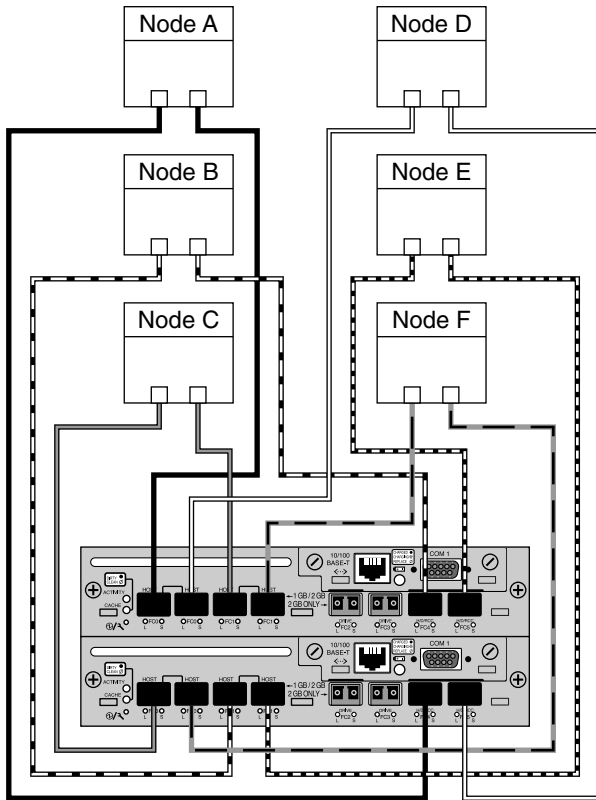


FIGURE 1-2 Sun StorEdge 3511 DAS Configuration With Multipathing and Two Controllers

The two-controller SAN configurations allow 32 LUNs to be mapped to each pair of host channels. Since these configurations use multipathing, each node sees a total of 64 LUNs.

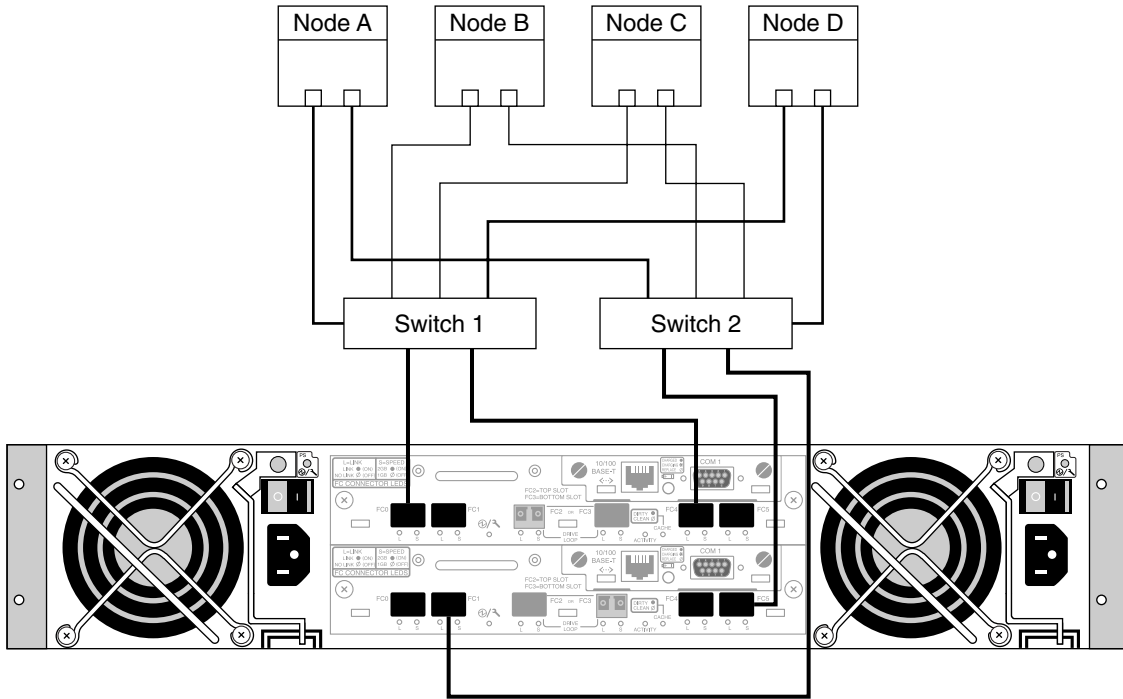


FIGURE 1-3 Sun StorEdge 3510 SAN Configuration With Multipathing and Two Controllers

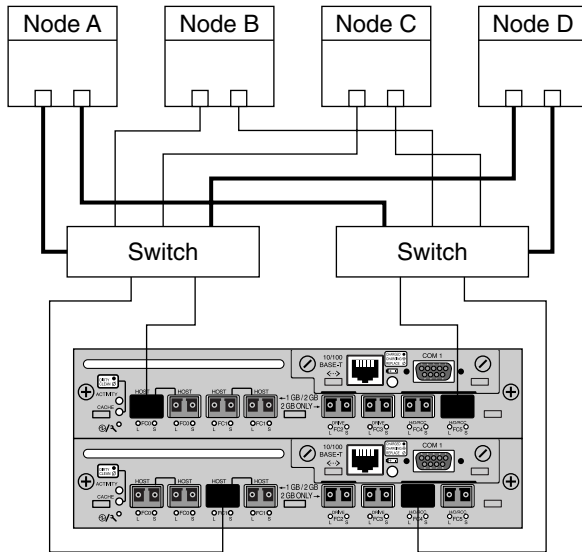


FIGURE 1-4 Sun StorEdge 3511 SAN Configuration With Multipathing and Two Controllers

▼ How to Install a Storage Array

Before installing or configuring your cluster, see “Known Problems” in the *Sun Cluster 3.0-3.1 Release Notes Supplement* for important information about the StorEdge 3510 and 3511 FC storage arrays.

Use this procedure to install and configure storage arrays *before* installing the Solaris Operating System and Sun Cluster software on your cluster nodes. If you need to add a storage array to an existing cluster, use the procedure in “Adding a Storage Array to a Running Cluster” on page 17

Before You Begin This procedure assumes that the hardware is not connected.

Note – If you plan to attach a StorEdge 3510 or 3511 FC expansion storage array to a StorEdge 3510 or 3511 FC RAID storage array, attach the expansion storage array before connecting the RAID storage array to the cluster nodes. See the *Sun StorEdge 3000 Family Installation, Operation, and Service Manual* for more information.

1 Install host adapters in the nodes that connect to the storage array.

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

2 If necessary, install the Fibre Channel (FC) switches.

For the procedure on installing an FC switch, see the documentation that shipped with your switch hardware.

Note – You must use FC switches when installing storage arrays in a SAN configuration.

3 If necessary, install gigabit interface converters (GBICs) or Small Form-Factor Pluggables (SFPs) in the FC switches.

For the procedures on installing a GBIC or an SFP to an FC switch, see the documentation that shipped with your FC switch hardware.

4 Cable the storage array.

For the procedures on connecting your FC storage array, see *Sun StorEdge 3000 Family Installation, Operation, and Service Manual*.

- If you plan to create a storage area network (SAN), connect the storage array to the FC switches using fiber-optic cables.
- If you plan to have a DAS configuration, connect the storage array to the nodes.

5 Power on the storage arrays.

Verify that all components are powered on and functional.

For the procedure on powering up the storage arrays and checking LEDs, see the *Sun StorEdge 3000 Family Installation, Operation, and Service Manual*.

6 Set up and configure the storage array.

For procedures on setting up logical drives and LUNs, see the *Sun StorEdge 3000 Family Installation, Operation, and Service Manual* or the *Sun StorEdge 3000 Family RAID Firmware 3.27 User's Guide*.

For the procedure on configuring the storage array, see *Sun StorEdge 3000 Family Installation, Operation, and Service Manual*.

7 On all nodes, install the Solaris operating system and 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 “[How to Install Solaris Software](#)” in *Sun Cluster Software Installation Guide for Solaris OS*.

8 Install any required storage array controller firmware.

Sun Cluster software requires patch version 113723–03 or later for each Sun StorEdge 3510 array in the cluster.

See the Sun Cluster release notes documentation for information about accessing Sun's EarlyNotifier web pages. The EarlyNotifier web pages list information about any required patches or firmware levels that are available for download.

9 Install any required patches or software for Solaris I/O multipathing software support to nodes and enable multipathing.

When using these arrays, Sun Cluster software requires Sun StorEdge SAN Foundation software:

- **SPARC: For the Sun StorEdge 3510 storage array, at least Sun StorEdge SAN Foundation software version 4.2.**
- **SPARC: For the Sun StorEdge 3511 storage array, at least Sun StorEdge SAN Foundation software version 4.4.**
- **x86: For x86 based clusters, at least the Sun StorEdge SAN Foundation software that is bundled with Solaris 10.**

For the procedure about how to install the Solaris I/O multipathing software, see [“How to Install Sun Multipathing Software”](#) in *Sun Cluster Software Installation Guide for Solaris OS*.

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

```
# devfsadm -C
```

11 On all nodes, confirm that the storage arrays that you installed are visible.

```
# luxadm probe
```

12 If necessary, label the LUNs.

```
# format
```

13 Install the Sun Cluster software and volume management software.

For software installation procedures, see the Sun Cluster software installation documentation.

See Also To continue with Sun Cluster software installation tasks, see the Sun Cluster software installation documentation.

Adding a Storage Array to a Running Cluster

Use this procedure to add new storage array to a running cluster. To install to a new Sun Cluster that is not running, use the procedure in [“How to Install a Storage Array”](#) on page 15.

If you need to add a storage array to more than two nodes, repeat the steps for each additional node that connects to the storage array.

Note – This procedure assumes that your nodes are not configured with dynamic reconfiguration functionality.

If your nodes are configured for dynamic reconfiguration, see the Sun Cluster system administration documentation and skip steps that instruct you to shut down the node.

▼ **How to Perform Initial Configuration Tasks on the Storage Array**

1 Power on the storage array.

2 Set up and configure the storage array.

For the procedures on configuring the storage array, see the *Sun StorEdge 3000 Family Installation, Operation, and Service Manual*.

3 If necessary, upgrade the storage array's controller firmware.

Sun Cluster software requires patch version 113723-03 or later for each Sun StorEdge 3510 array in the cluster.

See the Sun Cluster release notes documentation for information about accessing Sun's EarlyNotifier web pages. The EarlyNotifier web pages list information about any required patches or firmware levels that are available for download. For the procedure on applying any host adapter's firmware patch, see the firmware patch README file.

4 Configure the new storage array. Map the LUNs to the host channels.

For the procedures on setting up logical drives and LUNs, see the *Sun StorEdge 3000 Family Installation, Operation, and Service Manual* or *Sun StorEdge 3000 Family RAID Firmware 3.27 User's Guide*.

5 To continue adding the storage array, proceed to [“How to Connect the Storage Array to FC Switches” on page 18](#).

▼ **How to Connect the Storage Array to FC Switches**

Use this procedure if you plan to add a storage array to a SAN environment. If you do not plan to add the storage array to a SAN environment, go to [“How to Connect the Node to the FC Switches or the Storage Array” on page 19](#).

1 Install the SFPs in the storage array that you plan to add.

For the procedure on installing an SFP, see the *Sun StorEdge 3000 Family Installation, Operation, and Service Manual*.

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

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

3 Install a fiber-optic cable between the new storage array and each FC switch.

For the procedure on installing a fiber-optic cable, see the *Sun StorEdge 3000 Family Installation, Operation, and Service Manual*.

4 To finish adding your storage array, see “How to Connect the Node to the FC Switches or the Storage Array” on page 19.

▼ How to Connect the Node to the FC Switches or the Storage Array

Use this procedure when you add a storage array to a SAN or DAS configuration. In SAN configurations, you connect the node to the FC switches. In DAS configurations, you connect the node directly to the storage array.

Before You Begin This procedure provides the long forms of the Sun Cluster commands. Most commands also have short forms. Except for the forms of the command names, the commands are identical. For a list of the commands and their short forms, see [Appendix A, “Sun Cluster Object-Oriented Commands,”](#) in *Sun Cluster 3.1 - 3.2 Hardware Administration Manual for Solaris OS*.

To perform this procedure, become superuser or assume a role that provides `solaris.cluster.read` and `solaris.cluster.modify` role-based access control (RBAC) authorization.

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

Record this information because you will use it in [Step 12](#) and [Step 13](#) of this procedure to return resource groups and device groups to these nodes.

■ **If you are using Sun Cluster 3.2, use the following commands:**

```
# clresourcegroup status +
# cldevicegroup status +
```

■ **If you are using Sun Cluster 3.1, use the following command:**

```
# scstat
```

2 Move all resource groups and device groups off the node that you plan to connect.

■ **If you are using Sun Cluster 3.2, use the following command:**

```
# clnode evacuate nodename
```

- **If you are using Sun Cluster 3.1, use the following command:**

```
# scswitch -S -h nodename
```

- 3 If you need to install host adapters in the node, see the documentation that shipped with your host adapters and install the adapters.**

- 4 If necessary, install GBICs or SFPs to the FC switches or the storage array.**

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

For the procedure on installing a GBIC or an SFP to a storage array, see the *Sun StorEdge 3000 Family Installation, Operation, and Service Manual*.

- 5 Connect fiber-optic cables between the node and the FC switches or the storage array.**

For the procedure on installing a fiber-optic cable, see the *Sun StorEdge 3000 Family Installation, Operation, and Service Manual*.

- 6 If necessary, install the required Solaris patches for storage array support on the node.**

See the Sun Cluster release notes documentation for information about accessing Sun's EarlyNotifier web pages. The EarlyNotifier web pages list information about any required patches or firmware levels that are available for download. For the procedure on applying any host adapter's firmware patch, see the firmware patch README file.

- 7 On the node, update the /devices and /dev entries.**

```
# devfsadm -C
```

- 8 On the node, update the paths to the device ID instances.**

- **If you are using Sun Cluster 3.2, use the following command:**

```
# cldevice populate
```

- **If you are using Sun Cluster 3.1, use the following command:**

```
# scgdevs
```

- 9 If necessary, label the LUNs on the new storage array.**

```
# format
```

- 10 (Optional) On the node, verify that the device IDs are assigned to the new LUNs.**

- **If you are using Sun Cluster 3.2, use the following command:**

```
# cldevice list -v
```

- If you are using Sun Cluster 3.1, use the following commands:

```
# scdidadm -C
# scdidadm -l
```

- 11 Repeat [Step 2](#) to [Step 10](#) for each remaining node that you plan to connect to the storage array.

- 12 (Optional) Restore the device groups to the original node.

Perform the following step for each device group you want to return to the original node.

- If you are using Sun Cluster 3.2, use the following command:

```
# cldevicegroup switch -n nodename devicegroup1[ devicegroup2 ...]
```

`-n nodename` The node to which you are restoring device groups.

`devicegroup1[devicegroup2 ...]` The device group or groups that you are restoring to the node.

- If you are using Sun Cluster 3.1, use the following command:

```
# scswitch -z -D devicegroup -h nodename
```

- 13 (Optional) Restore the resource groups to the original node.

Perform the following step for each resource group you want to return to the original node.

- If you are using Sun Cluster 3.2, use the following command:

```
# clresourcegroup switch -n nodename resourcegroup1[ resourcegroup2 ...]
```

`nodename` For failover resource groups, the node to which the groups are returned. For scalable resource groups, the node list to which the groups are returned.

`resourcegroup1[resourcegroup2 ...]` The resource group or groups that you are returning to the node or nodes.

- If you are using Sun Cluster 3.1, use the following command:

```
# scswitch -z -g resourcegroup -h nodename
```

- 14 Perform volume management administration to incorporate the new logical drives into the cluster.

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

Next Steps The best way to enable multipathing for a cluster is to install the multipathing software and enable multipathing before installing the Sun Cluster software and establishing the cluster. For this procedure, see “[How to Install Sun Multipathing Software](#)” in *Sun Cluster Software Installation Guide for Solaris OS*. If you need to add multipathing software to an established

cluster, see “How to Install Sun Multipathing Software” in *Sun Cluster Software Installation Guide for Solaris OS* and follow the troubleshooting steps to clean up the device IDs.

Configuring Storage Arrays in a Running Cluster

This section contains the procedures for configuring a storage array in a running cluster. [Table 1–2](#) lists these procedures.

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 might change unexpectedly. When you perform a check of the device ID configuration by running the `cldevice check` or `scdidadm -c` command, the following error message appears on your console if the device ID changed unexpectedly.

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

To fix device IDs that report this error, run the `cldevice repair` or `scdidadm -R` command for each affected device.

Note – Logical volumes are not supported in a Sun Cluster environment. Use logical drives as an alternative.

TABLE 1–2 Task Map: Configuring a Fibre-Channel Storage Array

Task	Information
Create a LUN	“How to Create and Map a LUN” on page 23
Remove a LUN	“How to Unmap and Remove a LUN” on page 24

▼ How to Create and Map a LUN

Use this procedure to create a LUN from unassigned storage capacity.

Before You Begin This procedure relies on the following prerequisites and assumptions.

- All nodes are booted in cluster mode and attached to the storage device.
- The storage device is installed and configured. If you are using multipathing, it is configured as described in the installation procedure.
- If you are using Sun StorEdge Traffic Manager software for the Solaris 9 OS or Solaris I/O multipathing software for the Solaris 10 OS, the software is installed and configured and the path between the software and the storage device is functioning.

This procedure provides the long forms of the Sun Cluster commands. Most commands also have short forms. Except for the forms of the command names, the commands are identical. For a list of the commands and their short forms, see [Appendix A, “Sun Cluster Object-Oriented Commands,”](#) in *Sun Cluster 3.1 - 3.2 Hardware Administration Manual for Solaris OS*.

To perform this procedure, become superuser or assume a role that provides `solaris.cluster.modify` RBAC authorization.

1 Follow the instructions in your storage device's documentation to create and map the LUN.

To allow multiple clusters and nonclustered systems to access the storage device, create initiator groups by using LUN filtering or masking.

2 If you are using multipathing, and if any devices that are associated with the LUN you created are at an unconfigured state, configure the STMS paths on each node that is connected to the storage device.

To determine if any devices are at an unconfigured state, use the following command:

```
# cfgadm -al | grep disk
```

To configure the STMS paths on each node, use the following command:

```
cfgadm -o force_update -c configure controllerinstance
```

To configure STMS paths for the Solaris 9 OS, see the *Sun StorEdge Traffic Manager Installation and Configuration Guide*. For the Solaris 10 OS, see the *Solaris Fibre Channel Storage Configuration and Multipathing Support Guide* for instructions on configuring Solaris I/O multipathing.

3 On one node that is connected to the storage device, use the `format` command to label the new LUN.

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

- If you are using Sun Cluster 3.2, use the following command:
`cldevice populate`
- If you are using Sun Cluster 3.1, use the following command:
`scgdevs`

Note – You might have a volume management daemon such as `vold` running on your node, and have a CD-ROM drive connected to the node. Under these conditions, a device busy error might be returned even if no disk is in the drive. This error is expected behavior. You can safely ignore this error message.

5 If you will manage this LUN with volume management software, use the appropriate Solaris Volume Manager or Veritas Volume Manager commands to update the list of devices on all nodes that are attached to the new volume that you created.

For more information, see your Solaris Volume Manager or Veritas Volume Manager documentation.

- See Also**
- To configure a LUN as a quorum device, see [Chapter 6, “Administering Quorum,”](#) in *Sun Cluster System Administration Guide for Solaris OS*.
 - To create a new resource or configure a running resource to use the new LUN, see [Chapter 2, “Administering Data Service Resources,”](#) in *Sun Cluster Data Services Planning and Administration Guide for Solaris OS*.

▼ How to Unmap and Remove a LUN

Use this procedure to remove one or more LUNs. See the *Sun StorEdge 3000 Family RAID Firmware 3.25 and 3.27 User's Guide* for the latest information about LUN administration.

This procedure assumes that all nodes are booted in cluster mode and attached to the storage device.



Caution – When you delete a LUN, you remove all data on the LUN that you delete.

Before You Begin

This procedure provides the long forms of the Sun Cluster commands. Most commands also have short forms. Except for the forms of the command names, the commands are identical. For a list of the commands and their short forms, see [Appendix A, “Sun Cluster Object-Oriented Commands,”](#) in *Sun Cluster 3.1 - 3.2 Hardware Administration Manual for Solaris OS*.

To perform this procedure, become superuser or assume a role that provides `solaris.cluster.read` and `solaris.cluster.modify` RBAC authorization.

1 Identify the LUN or LUNs that you will remove.

Refer to your Solaris Volume Manager or Veritas Volume Manager documentation for the appropriate commands.

For example, use one of the following pairs of commands.

- **If you are using Sun Cluster 3.2, use the following commands:**

```
# luxadm probe
# cldevice show
```

- **If you are using Sun Cluster 3.1, use the following commands:**

```
# luxadm probe
# scdidadm -L pathname
```

2 If the LUN that you will remove is configured as a quorum device, choose and configure another device as the quorum device. Then remove the old quorum device.

To determine whether the LUN is configured as a quorum device, use one of the following commands.

- **If you are using Sun Cluster 3.2, use the following command:**

```
# clquorum show
```

- **If you are using Sun Cluster 3.1, use the following command:**

```
# scstat -q
```

For procedures about how to add and remove quorum devices, see [Chapter 6, “Administering Quorum,”](#) in *Sun Cluster System Administration Guide for Solaris OS*.

3 Remove the LUN from disksets or disk groups.

Run the appropriate Solaris Volume Manager or Veritas Volume Manager commands to remove the LUN from any diskset or disk group. For more information, see your Solaris Volume Manager or Veritas Volume Manager documentation for more information. See the following note for additional Veritas Volume Manager commands that are required.

Note – LUNs that were managed by Veritas Volume Manager must be completely removed from Veritas Volume Manager control before you can delete them from the Sun Cluster environment. After you delete the LUN from any disk group, use the following commands *on both nodes* to remove the LUN from Veritas Volume Manager control.

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

4 Unmap the LUN from both host channels.

For the procedure on unmapping a LUN, see the *Sun StorEdge 3000 Family RAID Firmware 3.25 and 3.27 User's Guide*.

5 (Optional) Delete the logical drive.

For more information, see *Sun StorEdge 3000 Family RAID Firmware 3.25 and 3.27 User's Guide*.

6 On both nodes, remove the paths to the LUN that you are deleting.

```
# devfsadm -C
```

7 On both nodes, remove all obsolete device IDs.

- If you are using Sun Cluster 3.2, use the following command:

```
# cldevice clear
```

- If you are using Sun Cluster 3.1, use the following command:

```
# scdidadm -C
```

Maintaining Storage Arrays

This section contains the procedures for maintaining a storage array in a Sun Cluster environment. Maintenance tasks are listed in [Table 1–3](#) contain cluster-specific tasks. Tasks that are not cluster-specific are referenced in a list following the table.

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 might change unexpectedly. When you perform a check of the device ID configuration by running the `cldevice check` or `scdidadm -c` command, the following error message appears on your console if the device ID 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 device IDs that report this error, run the `cldevice repair` or `scdidadm -R` command for each affected device.

TABLE 1-3 Task Map: Maintaining a Storage Array

Task	Information
Remove a storage array from a running cluster.	“How to Remove a Storage Array From a Running Cluster” on page 28
Upgrade array firmware.	“How to Upgrade Storage Array Firmware” on page 31
Replace a disk drive in an storage array.	“How to Replace a Disk Drive” on page 32
Replace a host adapter.	“How to Replace a Host Adapter” on page 33
Replace a node-to-switch fiber optic cable.	“Replacing a Node-to-Switch Component” on page 36
Replace a gigabit interface converter (GBIC) or Small Form-Factor Pluggable (SFP) on a node's host adapter.	“Replacing a Node-to-Switch Component” on page 36
Replace a GBIC or an SFP on an FC switch, connecting to a node.	“Replacing a Node-to-Switch Component” on page 36
Replace a storage array-to-switch fiber-optic cable.	“Replacing a Node-to-Switch Component” on page 36
Replace a GBIC or an SFP on an FC switch, connecting to a storage array.	“Replacing a Node-to-Switch Component” on page 36
Replace an FC switch.	“Replacing a Node-to-Switch Component” on page 36
Replace the power cord of an FC switch.	“Replacing a Node-to-Switch Component” on page 36
Replace the controller.	“Replacing a Node-to-Switch Component” on page 36

TABLE 1-3 Task Map: Maintaining a Storage Array (Continued)

Task	Information
Replace the chassis.	“How to Replace a Chassis in a Running Cluster” on page 38
Add a node to the storage array.	Sun Cluster system administration documentation
Remove a node from the storage array.	Sun Cluster system administration documentation

StorEdge 3510 and 3511 FC RAID Array FRUs

The following is a list of administrative tasks that require no cluster-specific procedures. See the *Sun StorEdge 3000 Family Installation, Operation, and Service Manual* for the following procedures.

- Replacing an Ethernet cable.
- Replacing the power cable on the storage array.
- Replacing a Power and Cooling Unit (PCU).
- Replacing the expansion unit-to-RAID storage array interconnect components such as SFPs and fiber-optic cables.

▼ How to Remove a Storage Array From a Running Cluster

Use this procedure to permanently remove storage arrays and their submirrors from a running cluster.

If you need to remove a storage array from more than two nodes, repeat [Step 6](#) to [Step 13](#) for each additional node that connects to the storage array.



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

Before You Begin

This procedure provides the long forms of the Sun Cluster commands. Most commands also have short forms. Except for the forms of the command names, the commands are identical. For a list of the commands and their short forms, see [Appendix A, “Sun Cluster Object-Oriented Commands,”](#) in *Sun Cluster 3.1 - 3.2 Hardware Administration Manual for Solaris OS*.

To perform this procedure, become superuser or assume a role that provides `solaris.cluster.read` and `solaris.cluster.modify` RBAC authorization.

- 1 If the storage array you are removing contains any quorum devices, choose another disk drive to configure as the quorum device. Then remove the old quorum device.**

To determine whether the LUN is configured as a quorum device, use one of the following commands.

- **If you are using Sun Cluster 3.2, use the following command:**

```
# clquorum show
```

- **If you are using Sun Cluster 3.1, use the following command:**

```
# scstat -q
```

For procedures on adding and removing quorum devices, see [Chapter 6, “Administering Quorum,”](#) in *Sun Cluster System Administration Guide for Solaris OS*.

- 2 If necessary, back up all database tables, data services, and drives associated with each storage array that you are removing.**

- 3 Determine the resource groups and device groups that are running on all nodes.**

Record this information because you will use it in [Step 17](#) and [Step 18](#) of this procedure to return resource groups and device groups to these nodes.

- **If you are using Sun Cluster 3.2, use the following commands:**

```
# clresourcegroup status +  
# cldevicegroup status +
```

- **If you are using Sun Cluster 3.1, use the following command:**

```
# scstat
```

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

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

- 5 Run the appropriate volume manager commands to remove references to each LUN that belongs to the storage array that you are removing.**

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

- 6 Shut down the node.**

For the full procedure on shutting down and powering off a node, see [Chapter 3, “Shutting Down and Booting a Cluster,”](#) in *Sun Cluster System Administration Guide for Solaris OS*.

- 7 If necessary, disconnect the storage arrays from the nodes or the FC switches.**

- 8 If the storage array that you are removing is not the last storage array connected to the node, skip to [Step 10](#).
- 9 If the storage array that you are removing is the last storage array connected to the node, disconnect the fiber-optic cable between the node and the FC switch that was connected to this storage array.
- 10 If you do not want to remove the host adapters from the node, skip to [Step 13](#).
- 11 If you want to remove the host adapters from the node, power off the node.
- 12 **Remove the host adapters from the node.**
For the procedure on removing host adapters, see the documentation that shipped with your host adapter and nodes.
- 13 **Boot the node into cluster mode.**
For more information on booting nodes, see [Chapter 3, “Shutting Down and Booting a Cluster,”](#) in *Sun Cluster System Administration Guide for Solaris OS*.
- 14 **Repeat steps [Step 6](#) to [Step 13](#) on each additional node that you need to disconnect from the storage array.**
- 15 **On all cluster nodes, remove the paths to the devices that you are deleting.**
`# devfsadm -C`
- 16 **On all cluster nodes, remove all obsolete device IDs.**
 - **If you are using Sun Cluster 3.2, use the following command:**
`# cldevice clear`
 - **If you are using Sun Cluster 3.1, use the following command:**
`# scdidadm -C`
- 17 **(Optional) Restore the device groups to the original node.**
Perform the following step for each device group you want to return to the original node.
 - **If you are using Sun Cluster 3.2, use the following command:**

<code># cldevicegroup switch -n nodename devicegroup1[devicegroup2 ...]</code>	
<code>-n nodename</code>	The node to which you are restoring device groups.
<code>devicegroup1[devicegroup2 ...]</code>	The device group or groups that you are restoring to the node.

- If you are using Sun Cluster 3.1, use the following command:

```
# scswitch -z -D devicegroup -h nodename
```

18 (Optional) Restore the resource groups to the original node.

Perform the following step for each resource group you want to return to the original node.

- If you are using Sun Cluster 3.2, use the following command:

```
# clresourcegroup switch -n nodename resourcegroup1 [resourcegroup2 ...]
```

nodename

For failover resource groups, the node to which the groups are returned. For scalable resource groups, the node list to which the groups are returned.

resourcegroup1 [*resourcegroup2* ...]

The resource group or groups that you are returning to the node or nodes.

- If you are using Sun Cluster 3.1, use the following command:

```
# scswitch -z -g resourcegroup -h nodename
```

See Also To prepare the storage array for later use, unmap and delete all LUNs and logical drives. See “[How to Unmap and Remove a LUN](#)” on page 24 for information about LUN removal. For more information about removing logical drives, see the *Sun StorEdge 3000 Family Installation, Operation, and Service Manual*.

▼ How to Upgrade Storage Array Firmware

Use this procedure to upgrade storage array firmware in a running cluster.

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 might change unexpectedly. When you perform a check of the device ID configuration by running the `cldevice check` or `sccdidadm -c` command, the following error message appears on your console if the device ID 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 device IDs that report this error, run the `cldevice repair` or `sccdidadm -R` command for each affected device.

1 Stop all I/O to the storage arrays you are upgrading.

2 Download the firmware to the storage arrays.

Refer to the *Sun StorEdge 3000 Family RAID Firmware 3.25 and 3.27 User's Guide* and to any patch readme files for more information.

3 Confirm that all storage arrays that you upgraded are visible to all nodes.

```
# luxadm probe
```

4 Restart all I/O to the storage arrays.

You stopped I/O to these storage arrays in [Step 1](#).

▼ How to Replace a Disk Drive

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

Before You Begin

This procedure provides the long forms of the Sun Cluster commands. Most commands also have short forms. Except for the forms of the command names, the commands are identical. For a list of the commands and their short forms, see [Appendix A, “Sun Cluster Object-Oriented Commands,”](#) in *Sun Cluster 3.1 - 3.2 Hardware Administration Manual for Solaris OS*.

To perform this procedure, become superuser or assume a role that provides `solaris.cluster.read` RBAC authorization.

1 If the failed disk drive does not affect the storage array LUN's availability, skip to [Step 4](#).**2 If the failed disk drive affects the storage array LUN's availability, use volume manager commands to detach the submirror or plex.**

For more information, see your Solstice DiskSuite/Solaris Volume Manager or Veritas Volume Manager documentation.

3 If the LUN (in [Step 1](#)) is configured as a quorum device, choose and configure another device to be the new quorum device. Remove the old quorum device.

To determine whether the LUN is configured as a quorum device, use one of the following commands.

- **If you are using Sun Cluster 3.2, use the following command:**

```
# clquorum show
```

- **If you are using Sun Cluster 3.1, use the following command:**

```
# scstat -q
```


For procedures about how to add and remove quorum devices, see your Sun Cluster system administration documentation.

4 Replace the failed disk drive.

For instructions, refer to the *Sun StorEdge 3000 Family Installation, Operation, and Service Manual*.

5 (Optional) If you reconfigured a quorum device in Step 3, restore the original quorum configuration.

For the procedure about how to add a quorum device, see your Sun Cluster system administration documentation.

6 If you detached a submirror or plex in Step 1, use volume manager commands to reattach the submirror or plex.

For more information, see your Solstice DiskSuite/Solaris Volume Manager or Veritas Volume Manager documentation.

▼ How to Replace a 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 that you are replacing.

Before You Begin This procedure relies on the following prerequisites and assumptions.

- Except for the failed host adapter, your cluster is operational and all nodes are powered on.
- Your nodes are not configured with dynamic reconfiguration functionality.

If your nodes are configured for dynamic reconfiguration **and** you are using two entirely separate hardware paths to your shared data, see the *Sun Cluster Hardware Administration Manual for Solaris OS* and skip steps that instruct you to shut down the cluster.

You cannot replace a single, dual-port HBA that has quorum configured on that storage path by using DR. Follow all steps in the procedure. For the details on the risks and limitations of this configuration, see “[Configuring Cluster Nodes With a Single, Dual-Port HBA](#)” in *Sun Cluster 3.1 - 3.2 Hardware Administration Manual for Solaris OS*.

Exceptions to this restriction include three-node or larger cluster configurations where no storage device has a quorum device configured.

This procedure provides the long forms of the Sun Cluster commands. Most commands also have short forms. Except for the forms of the command names, the commands are identical. For a list of the commands and their short forms, see [Appendix A, “Sun Cluster Object-Oriented Commands,”](#) in *Sun Cluster 3.1 - 3.2 Hardware Administration Manual for Solaris OS*.

- 1 **Become superuser or assume a role that provides `solaris.cluster.read` and `solaris.cluster.modify` RBAC authorization.**
- 2 **Determine the resource groups and device groups that are running on Node A.**
Record this information because you use this information in [Step 10](#) and [Step 11](#) of this procedure to return resource groups and device groups to Node A.
 - **If you are using Sun Cluster 3.2, use the following commands:**

```
# clresourcegroup status -n NodeA
# cldevicegroup status -n NodeA
```

`-n NodeA` The node for which you are determining resource groups and device groups.
 - **If you are using Sun Cluster 3.1, use the following command:**

```
# scstat
```
- 3 **Move all resource groups and device groups off Node A.**
 - **If you are using Sun Cluster 3.2, use the following command:**

```
# clnode evacuate nodename
```
 - **If you are using Sun Cluster 3.1, use the following command:**

```
# scswitch -S -h nodename
```
- 4 **Shut down Node A.**
For the full procedure about how to shut down and power off a node, see [Chapter 3, “Shutting Down and Booting a Cluster,”](#) in *Sun Cluster System Administration Guide for Solaris OS*.
- 5 **Power off Node A.**
- 6 **Replace the failed host adapter.**
To remove and add host adapters, see the documentation that shipped with your nodes.
- 7 **If you need to upgrade the node's host adapter firmware, boot Node A into noncluster mode by adding `-x` to your boot instruction. Proceed to [Step 8](#).**
If you do not need to upgrade firmware, skip to [Step 9](#).
- 8 **Upgrade the host adapter firmware on Node A.**
If you use the Solaris 8, Solaris 9, or Solaris 10 Operating System, [Sun Connection Update Manager](#) keeps you informed of the latest versions of patches and features. Using notifications and intelligent needs-based updating, [Sun Connection](#) helps improve operational efficiency and ensures that you have the latest software patches for your Sun software.

You can download the Sun Connection Update Manager product for free by going to <http://www.sun.com/download/products.xml?id=4457d96d>.

Additional information for using the Sun patch management tools is provided in *Solaris Administration Guide: Basic Administration* at <http://docs.sun.com> (<http://docs.sun.com>). Refer to the version of this manual for the Solaris OS release that you have installed.

If you must apply a patch when a node is in noncluster mode, you can apply it in a rolling fashion, one node at a time, unless instructions for a patch require that you shut down the entire cluster. Follow the procedures in “How to Apply a Rebooting Patch (Node)” in *Sun Cluster System Administration Guide for Solaris OS* to prepare the node and to boot it in noncluster mode. For ease of installation, consider applying all patches at the same time. That is, apply all patches to the node that you place in noncluster mode.

For a list of patches that affect Sun Cluster, see the [Sun Cluster Wiki Patch Klatch](#).

For required firmware, see the *Sun System Handbook*.

9 Boot Node A into cluster mode.

For more information about how to boot nodes, see Chapter 3, “Shutting Down and Booting a Cluster,” in *Sun Cluster System Administration Guide for Solaris OS*.

10 (Optional) Restore the device groups to the original node.

Do the following for each device group that you want to return to the original node.

- **If you are using Sun Cluster 3.2, use the following command:**

```
# cldevicegroup switch -n nodename devicegroup1[ devicegroup2 ...]
```

-n nodename

The node to which you are restoring device groups.

devicegroup1[devicegroup2 ...]

The device group or groups that you are restoring to the node.

- **If you are using Sun Cluster 3.1, use the following command:**

```
# scswitch -z -D devicegroup -h nodename
```

11 (Optional) Restore the resource groups to the original node.

Do the following for each resource group that you want to return to the original node.

- **If you are using Sun Cluster 3.2, use the following command:**

```
# clresourcegroup switch -n nodename resourcegroup1[ resourcegroup2 ...]
```

nodename

For failover resource groups, the node to which the groups are returned. For scalable resource groups, the node list to which the groups are returned.

resourcegroup1 [*resourcegroup2* ...] The resource group or groups that you are returning to the node or nodes.

- **If you are using Sun Cluster 3.1, use the following command:**

```
# scswitch -z -g resourcegroup -h nodename
```

Replacing a Node-to-Switch Component

Use this procedure to replace a node-to-switch component that has failed or that you suspect might be contributing to a problem.

Note – Node-to-switch components that are covered by this procedure include the following components:

- Node-to-switch fiber-optic cables
- Gigabit interface converters (GBICs) or small form-factor pluggables (SFPs) on an FC switch
- FC switches

To replace a host adapter, see [“How to Replace a Host Adapter” on page 33](#).

This procedure defines Node A as the node that is connected to the node-to-switch component that you are replacing. This procedure assumes that, except for the component you are replacing, your cluster is operational.

Ensure that you are following the appropriate instructions:

- If your cluster uses multipathing, see [“How to Replace a Node-to-Switch Component in a Cluster That Uses Multipathing” on page 36](#).
- If your cluster does *not* use multipathing, see [“How to Replace a Node-to-Switch Component in a Cluster Without Multipathing” on page 37](#).

▼ **How to Replace a Node-to-Switch Component in a Cluster That Uses Multipathing**

- 1 **If your configuration is active-passive, and if the active path is the path that needs a component replaced, make that path passive.**
- 2 **Replace the component.**
Refer to your hardware documentation for any component-specific instructions.

- 3 (Optional) If your configuration is active-passive and you changed your configuration in [Step 1](#), switch your original data path back to active.

▼ How to Replace a Node-to-Switch Component in a Cluster Without Multipathing

Before You Begin This procedure provides the long forms of the Sun Cluster commands. Most commands also have short forms. Except for the forms of the command names, the commands are identical. For a list of the commands and their short forms, see [Appendix A, “Sun Cluster Object-Oriented Commands,”](#) in *Sun Cluster 3.1 - 3.2 Hardware Administration Manual for Solaris OS*.

- 1 **Become superuser or assume a role that provides `solaris.cluster.read` and `solaris.cluster.modify` RBAC authorization.**
- 2 **If the physical data path has failed, do the following:**
 - a. Replace the component.
 - b. Fix the volume manager error that was caused by the failed data path.
 - c. (Optional) If necessary, return resource groups and device groups to this node.

You have completed this procedure.

- 3 **If the physical data path has not failed, determine the resource groups and device groups that are running on Node A.**
 - **If you are using Sun Cluster 3.2, use the following commands:**

```
# clresourcegroup status -n NodeA
# cldevicegroup status -n NodeA
```

`-n NodeA` The node for which you are determining resource groups and device groups.
 - **If you are using Sun Cluster 3.1, use the following command:**

```
# scstat
```
- 4 **Move all resource groups and device groups to another node.**
 - **If you are using Sun Cluster 3.2, use the following command:**

```
# clnode evacuate nodename
```
 - **If you are using Sun Cluster 3.1, use the following command:**

```
# scswitch -S -h nodename
```

5 Replace the node-to-switch component.

Refer to your hardware documentation for any component-specific instructions.

6 (Optional) Restore the device groups to the original node.

Do the following for each device group that you want to return to the original node.

- **If you are using Sun Cluster 3.2, use the following command:**

```
# cldevicegroup switch -n nodename devicegroup1[ devicegroup2 ...]
```

-n nodename

The node to which you are restoring device groups.

devicegroup1[devicegroup2 ...]

The device group or groups that you are restoring to the node.

- **If you are using Sun Cluster 3.1, use the following command:**

```
# scswitch -z -D devicegroup -h nodename
```

7 (Optional) Restore the resource groups to the original node.

Do the following for each resource group that you want to return to the original node.

- **If you are using Sun Cluster 3.2, use the following command:**

```
# clresourcegroup switch -n nodename resourcegroup1[ resourcegroup2 ...]
```

nodename

For failover resource groups, the node to which the groups are returned. For scalable resource groups, the node list to which the groups are returned.

resourcegroup1[resourcegroup2 ...]

The resource group or groups that you are returning to the node or nodes.

- **If you are using Sun Cluster 3.1, use the following command:**

```
# scswitch -z -g resourcegroup -h nodename
```

▼ How to Replace a Chassis in a Running Cluster

Use this procedure to replace a storage array chassis in a running cluster. This procedure assumes that you want to retain all FRUs other than the chassis and the backplane.

- 1 To stop all I/O activity to this storage array, detach the submirrors that are connected to the chassis you are replacing.**

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

- 2 If this storage array is not made redundant by host-based mirroring, shut down the cluster.**
For the full procedure on shutting down a cluster, see the Sun Cluster system administration documentation.
- 3 Replace the chassis and backplane.**
For the procedure on replacing a chassis, see the *Sun StorEdge 3000 Family Installation, Operation, and Service Manual*.
- 4 If you shut down the cluster in [Step 2](#), boot it back into cluster mode.**
For the full procedure on booting a cluster, see [Chapter 3, “Shutting Down and Booting a Cluster,”](#) in *Sun Cluster System Administration Guide for Solaris OS*.
- 5 Reattach the submirrors that you detached in [Step 1](#) to resynchronize them.**



Caution – The world wide numbers (WWNs) might change as a result of this procedure. If the WWNs change, and you must reconfigure your volume manager software to recognize the new WWNs.

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

Index

A

adding
 nodes, 28
 storage array s, 17-22
arrays, *See* storage array s

C

configuring, storage arrays, 22-26
cooling units, replacing, 28
creating, LUNs, 23-24

D

disk drives, replacing, 32-33

E

Ethernet cables, replacing, 28

F

firmware, upgrading, 31-32
FRUs, 28

H

HBAs, replacing, 33-36

help, 7-8

host adapters, replacing, 33-36

I

installing
 See also adding
 storage array s, 11-22

L

LUNs
 creating, 23-24
 mapping, 23-24
 removing, 24-26

M

maintaining, storage arrays, 26-39
mapping LUNs, 23-24

N

node-to-switch components, replacing, 36-38
nodes
 adding, 28
 removing, 28

P

power cables, replacing, 28
power units, replacing, 28

R

removing
 LUNs, 24-26
 nodes, 28
 storage arrays, 28-31
replacing
 cooling units, 28
 disk drives, 32-33
 Ethernet cables, 28
 fiber optic cables, 28
 host adapters, 33-36
 node-to-switch components, 36-38
 power cables, 28
 power units, 28
 Small Form-Factor Pluggable (SFP), 28

S

storage arrays
 adding, 17-22
 configuring, 22-26
 installing, 11-22
 maintaining, 26-39
 removing, 28-31
 upgrading firmware, 31-32

T

technical support, 7-8

U

upgrading, storage array firmware, 31-32