



Sun Cluster 3.1 - 3.2 With StorageTek 2540 RAID Arrays Manual for Solaris OS



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

Part No: 820-1692-11
January 2009, Revision B

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, StorageTek, SunSolve, 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, StorageTek, SunSolve, 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 Requirements and Restrictions	11
Requirements for This Array	11
Restrictions for This Array	11
2 Installing and Configuring a StorageTek Array	13
Installing Storage Arrays	13
Storage Array Cabling Configurations	13
How to Install Storage Arrays in a New Cluster	14
▼ How to Add Storage Arrays to an Existing Cluster	16
Configuring Storage Arrays	17
▼ How to Create a Logical Volume	18
▼ How to Remove a Logical Volume	20
3 Maintaining a StorageTek Array	25
FRUs That <i>Do Not</i> Require Sun Cluster Maintenance Procedures	25
Maintaining Storage Arrays	26
▼ How to Upgrade Storage Array Firmware	26
▼ How to Remove a Storage Array	27
Replacing a Node-to-Switch Component	28
▼ How to Replace a Host Adapter	31
▼ How to Replace a Disk Drive	34
Index	37

Preface

The *Sun Cluster 3.1 - 3.2 With StorageTek 2540 RAID Arrays Manual for Solaris OS* provides procedures that are specific to a Sun StorageTek™ 2540 array that is 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 on any supported version of Sun Cluster 3.1 or 3.2 software on x86 based clusters. Use this manual with any version of Sun Cluster 3.1 or 3.2 software. Unless otherwise noted, procedures are the same for all Sun Cluster 3.1 or 3.2 versions.

This book assumes that you are performing one or more of the following tasks:

- You want to replace an array component to *prevent* a failure.
- You want to replace an array component because you have an *existing* failure.
- You want to add (to an *established* cluster) or install (to a *new* cluster) a storage array.

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 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

This book contains the following chapters:

- [Chapter 1, “Requirements and Restrictions,”](#) lists limitations of your use of Sun StorageTek 2540 storage arrays in a Sun Cluster environment.
- [Chapter 2, “Installing and Configuring a StorageTek Array,”](#) discusses how to install Sun StorageTek 2540 storage arrays and how to configure logical units on them.
- [Chapter 3, “Maintaining a StorageTek Array,”](#) describes how to maintain Sun StorageTek 2540 storage arrays in a running cluster.

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-1 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-2 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%</code> you have mail.
AaBbCc123	What you type, contrasted with onscreen computer output	<code>machine_name%</code> 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> . <i>A 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-3 Shell Prompts

Shell	Prompt
C shell	<code>machine_name%</code>
C shell for superuser	<code>machine_name#</code>

TABLE P-3 Shell Prompts (Continued)

Shell	Prompt
Bourne shell and Korn shell	\$
Bourne shell and Korn shell for superuser	#

◆ ◆ ◆ 1
CHAPTER 1

Requirements and Restrictions

This chapter includes only requirements and restrictions that have a direct impact on the procedures in this book. For general support information, contact your Sun service provider.

Requirements for This Array

Requirements do not exist for the StorageTek™ FC RAID array.

Restrictions for This Array

When using storage-based replication, do not configure a replicated volume as a quorum device. Locate any quorum devices on an unreplicated volume. See [“Using Storage-Based Data Replication”](#) in *Sun Cluster System Administration Guide for Solaris OS* for more information on storage-based replication.

Installing and Configuring a StorageTek Array

This chapter contains the procedures about how to install and configure StorageTek FC RAID arrays. These procedures are specific to a Sun™ Cluster environment.

This chapter contains the following main topics:

- “Installing Storage Arrays” on page 13
- “Configuring Storage Arrays” on page 17

For detailed information about storage array architecture, features, configuration utilities, and installation, see the StorageTek documentation listed in “[Related Documentation](#)” on page 6.

Installing Storage Arrays

This section contains the procedures that are listed in [Table 2-1](#).

TABLE 2-1 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 Storage Arrays in a New Cluster ” on page 14
Add a storage array to an existing cluster.	“ How to Add Storage Arrays to an Existing Cluster ” on page 16

Storage Array Cabling Configurations

You can install your storage array in several different configurations: [Figure 2-1](#) and [Figure 2-2](#) are two examples.

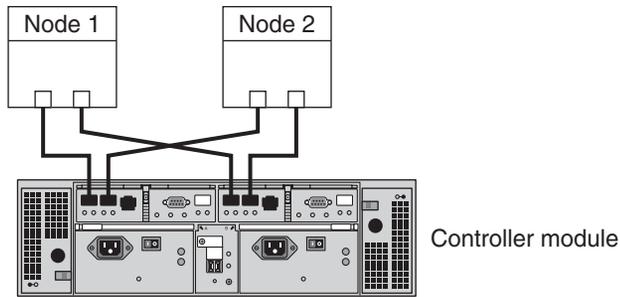


FIGURE 2-1 StorageTek 2540 Array Direct-Connect Configuration

Note – The StorageTek 2540 array houses two controllers, each having two host ports. The cabling approach is the same as shown in [Figure 2-1](#).

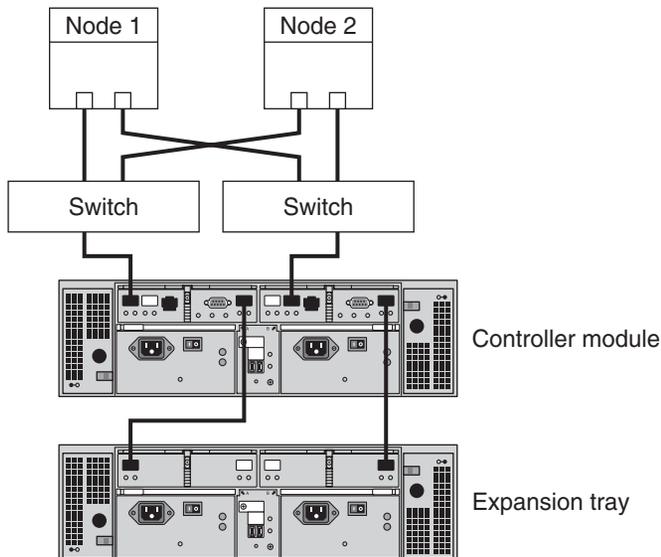


FIGURE 2-2 StorageTek 2540 Array Switched Configuration

How to Install Storage Arrays in a New Cluster

Use this procedure to install a storage array in a new cluster. To add a storage array to an existing cluster, use the procedure in [“How to Add Storage Arrays to an Existing Cluster”](#) on page 16.

This procedure relies on the following assumptions:

- You have *not* installed the Solaris Operating System.
- You have *not* installed the Sun Cluster software.
- Your cluster has *not* been established.
- You have enough host adapters to connect the nodes and the storage array to the service processor panel.

If you need to install host adapters, see “[How to Replace a Host Adapter](#)” on page 31. When this procedure asks you to replace the failed host adapter, simply install the new host adapter instead.

▼ Install and Cable the Hardware

1 **Unpack, place, and level your storage array.**

For the procedure about how to unpack, place, and level your storage array, see [Sun StorageTek 2500 Series Array Hardware Installation Guide, v1.0](#).

2 **(Optional) If you do not have a Fibre Channel (FC) switch installed, install it for the storage array.**

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

3 **Connect the nodes.**

- *(Direct-Attached Configuration)* **To see how to connect each node to the host ports on the array, see [Figure 2-1](#).**
- *(SAN Configuration)* **To see how to connect the nodes to the ports of the FC switches, see [Figure 2-2](#).**

For instructions, see [Sun StorageTek 2500 Series Array Hardware Installation Guide, v1.0](#).

4 **Install the storage array.**

For instructions, see [Sun StorageTek 2500 Series Array Hardware Installation Guide, v1.0](#).

5 **Power on the storage array nodes, and if necessary, the FC switches.**

For instructions, see [Sun StorageTek 2500 Series Array Hardware Installation Guide, v1.0](#).

▼ Install the Solaris Operating System and Configure Multipathing

- 1 **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*.

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

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*.

- See Also**
- To create a logical volume, see [“How to Create a Logical Volume”](#) on page 18.
 - To continue with Sun Cluster software installation tasks, see your Sun Cluster software installation documentation.

▼ How to Add Storage Arrays to an Existing Cluster

Use this procedure to add a new storage array to a *running* cluster. To install a new storage array in a Sun Cluster configuration that is not running, use the procedure in [“How to Install Storage Arrays in a New Cluster”](#) on page 14.

Before You Begin This procedure relies on the following assumptions:

- *(Veritas Volume Manager Only)* You have a version of Veritas Volume Manager that includes Array Support Library (ASL). For more information on this requirement, see [“Requirements for This Array”](#) on page 11.
- You have enough host adapters to connect the nodes and the storage array to the service processor panel.
If you need to install host adapters, see [“How to Replace a Host Adapter”](#) on page 31. When this procedure asks you to replace the failed host adapter, simply install the new host adapter instead.
- You do *not* need to install a new node in your cluster.
If you need to add a node to your cluster, see your Sun Cluster system administration documentation. Ensure that you install the required Solaris patches for storage array support.

- 1 **Unpack, place, and level your storage array.**

For the procedure about how to unpack, place, and level your storage array, see *Sun StorageTek 2500 Series Array Hardware Installation Guide, v1.0*.

2 (Optional) If you do not have a Fibre Channel (FC) switch installed, install it for the storage array.

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

3 Connect the nodes.

- *(Direct-Attached Configuration)* To see how to connect each node to the host ports on the array, see [Figure 2–1](#).
- *(SAN Configuration)* To see how to connect the nodes to the ports of the FC switches, see [Figure 2–2](#).

For instructions, see *Sun StorageTek 2500 Series Array Hardware Installation Guide, v1.0*.

4 Install the storage array.

For instructions, see *Sun StorageTek 2500 Series Array Hardware Installation Guide, v1.0*.

5 Power on the storage array nodes, and if necessary, the FC switches.

For instructions, see *Sun StorageTek 2500 Series Array Hardware Installation Guide, v1.0*.

- See Also**
- To create a logical volume, see [“How to Create a Logical Volume”](#) on page 18.
 - If you added a new node, perform the steps in [“Install the Solaris Operating System and Configure Multipathing”](#) on page 16 as appropriate. Apply the correct patches for storage array support enable multipathing.

Configuring Storage Arrays

This section contains the procedures about how to configure a storage array in a running cluster. [Table 2–2](#) lists these procedures.

TABLE 2–2 Task Map: Configuring a Storage Array

Task	Information
Create a logical volume.	“How to Create a Logical Volume” on page 18
Remove a logical volume.	“How to Remove a Logical Volume” on page 20

The following list outlines administrative tasks that require no cluster-specific procedures. See the storage array's online help for the following procedures.

- Creating a storage pool
- Removing a storage pool
- Creating a host group
- Removing a host group
- Adding a host to a host group
- Removing a host from a host group
- Adding initiators
- Removing initiators

▼ How to Create a Logical Volume

Use this procedure to create a logical volume from unassigned storage capacity.

Note – Sun storage documentation uses the following terms:

- Logical volume
- Logical device
- Logical unit number (LUN)

This manual uses *logical volume* to refer to all such logical constructs.

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, the storage device is configured as described in the installation procedure.
- If you are using Solaris I/O multipathing (MPxIO) for the Solaris 10 OS, previously called Sun StorEdge Traffic Manager in the Solaris 9 OS, verify that it is installed and configured and the path to the storage device is functioning. To configure the Traffic Manager for the Solaris 9 OS, see the *Sun StorEdge Traffic Manager Installation and Configuration Guide*. To configure multipathing for the Solaris 10 OS, see the *Solaris Fibre Channel Storage Configuration and Multipathing Support Guide*.

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.modify` role-based access control (RBAC) authorization.**

- 2 **Follow the instructions in your storage device's documentation to create and map the logical volume. For a URL to this storage documentation, see “[Related Documentation](#)” on page 6.**
 - Completely set up the logical volume. When you are finished, the volume must be created, mapped, mounted, and initialized.
 - If necessary, partition the volume.
 - To allow multiple clusters and nonclustered nodes to access the storage device, create initiator groups by using LUN masking.

3 **If you are not using multipathing, skip to [Step 5](#).**

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

To determine whether any devices that are associated with the volume you created are at an unconfigured state, use the following command.

```
# cfmadm -al | grep disk
```

Note – To configure the Solaris I/O multipathing paths on each node that is connected to the storage device, use the following command.

```
# cfmadm -o force_update -c configure controllerinstance
```

To configure the Traffic Manager for the Solaris 9 OS, see the *Sun StorEdge Traffic Manager Installation and Configuration Guide*. To configure multipathing for the Solaris 10 OS, see the *Solaris Fibre Channel Storage Configuration and Multipathing Support Guide*.

- 5 **On one node that is connected to the storage device, use the `format` command to label the new logical volume.**
- 6 **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 DVD drive connected to the node. Under these conditions, a device busy error might be returned even if no disk is inserted in the drive. This error is expected behavior. You can safely ignore this error message.

7 To manage this volume with volume management software, use 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 logical volume 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 logical volume, see [Chapter 2, “Administering Data Service Resources,” in *Sun Cluster Data Services Planning and Administration Guide for Solaris OS*](#).

▼ How to Remove a Logical Volume

Use this procedure to remove a logical volume. This procedure defines Node A as the node with which you begin working.

Note – Sun storage documentation uses the following terms:

- Logical volume
- Logical device
- Logical unit number (LUN)

This manual uses *logical volume* to refer to all such logical constructs.

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 logical volume and the path between the nodes and the storage device are both operational.

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 Identify the logical volume that you are removing.**

Refer to your Solaris Volume Manager or Veritas Volume Manager documentation for more information.
- 3 (Optional) Migrate all data off the logical volume that you are removing. Alternatively, back up that data.**
- 4 If the LUN that you are removing 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*.
- 5 If you are using volume management software, use that software to update the list of devices on all nodes that are attached to the logical volume that you are removing.**

For instructions about how to update the list of devices, see your Solaris Volume Manager or Veritas Volume Manager documentation.
- 6 If you are using volume management software, run the appropriate Solaris Volume Manager or Veritas Volume Manager commands to remove the logical volume from any diskset or disk group.**

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

Note – Volumes 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 volume from any disk group, use the following commands on both nodes to remove the volume from Veritas Volume Manager control.

```
# vxdisk offline Accessname
```

```
# vxdisk rm Accessname
```

Accessname Disk access name

7 If you are using multipathing, unconfigure the volume in Solaris I/O multipathing.

```
# cfgadm -o force_update -c unconfigure Logical_Volume
```

8 Access the storage device and remove the logical volume.

To remove the volume, see your storage documentation. For a list of storage documentation, see [“Related Documentation” on page 6](#).

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

Record this information because you use it in [Step 14](#) and [Step 15](#) 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
```

10 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
```

11 Shut down and reboot Node A.

To shut down and boot a node, see [Chapter 3, “Shutting Down and Booting a Cluster,” in *Sun Cluster System Administration Guide for Solaris OS*](#).

- 12** On Node A, remove the paths to the logical volume that you removed. Remove obsolete device IDs.
- If you are using Sun Cluster 3.2, use the following command:


```
# devfsadm -C
# cldevice clear
```
 - If you are using Sun Cluster 3.1, use the following command:


```
# devfsadm -C
# scdidadm -C
```
- 13** For each additional node that is connected to the shared storage that hosted the logical volume, repeat [Step 9](#) to [Step 12](#).
- 14** (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
```
- 15** (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
```


Maintaining a StorageTek Array

This chapter contains the procedures about how to maintain a Sun StorageTek™ 2540 Fibre Channel RAID array. These procedures are specific to a Sun™ Cluster environment.

This chapter contains the following procedures:

- “How to Upgrade Storage Array Firmware” on page 26
- “How to Remove a Storage Array” on page 27
- “Replacing a Node-to-Switch Component” on page 28
- “How to Replace a Host Adapter” on page 31
- “How to Replace a Disk Drive” on page 34

For detailed information about storage array architecture, features, and configuration utilities, see the StorageTek documentation listed in “[Related Documentation](#)” on page 6.

FRUs That *Do Not* Require Sun Cluster Maintenance Procedures

In general, the following is a list of administrative tasks that require no cluster-specific procedures. See the base-product documentation for these procedures.

Note – Depending on your configuration type and the state of your cluster, a few of the following FRUs might require cluster-specific steps. Such FRUs include the DSP and service processor panel.

- Adding a disk drive.
- Replacing a storage array's chassis.
- Replacing an Ethernet cable.
- Replacing an Ethernet hub in a cabinet.
- Replacing an expansion cabinet.
- Replacing a fan assembly.

- Replacing the power cable on the storage array.
- Replacing a power and cooling unit (PCU).
- Replacing a controller.
- Replacing a battery.

Before removing a disk drive, stop all I/O to the storage array. There are no other cluster-specific steps.

Maintaining Storage Arrays

This section contains the procedures about how to maintain a storage system in a running cluster. [Table 3-1](#) lists these procedures.

TABLE 3-1 Task Map: Maintaining a Storage System

Task	Information
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 26
Replace a node-to-switch component.	“How to Replace a Node-to-Switch Component in a Cluster Without Multipathing” on page 29
Replace a node's host adapter.	“How to Replace a Host Adapter” on page 31
Replace a disk drive.	“How to Replace a Disk Drive” on page 34
Add a node to the storage array.	Sun Cluster system administration documentation
Remove a node from the storage array.	Sun Cluster system administration documentation

▼ How to Upgrade Storage Array Firmware

Use this procedure to upgrade storage array firmware in a running cluster. Storage array firmware includes controller firmware, unit interconnect card (UIC) firmware, EPROM firmware, and disk drive firmware.

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.

- 1 **Stop all I/O to the storage arrays you are upgrading.**
- 2 **Apply the controller, disk drive, and loop-card firmware patches by using the arrays' GUI tools.**
For specific instructions, see your storage array's documentation.

- 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 Remove a Storage Array

Use this procedure to permanently remove a storage array from a running cluster.

This procedure defines Node N as the node that is connected to the storage array you are removing and the node with which you begin working.



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

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*.

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

- 1 If necessary, back up all database tables, data services, and volumes.
- 2 Remove references to the volumes that reside on the storage array that you are removing.
For more information, see your Sun Cluster, Solaris Volume Manager, or Veritas Volume Manager documentation.
- 3 Disconnect the cables that connected Node N to the FC switches in your storage array.
- 4 On all nodes, remove the obsolete Solaris links and device IDs.
 - If you are using Sun Cluster 3.2, use the following command:

```
# devfsadm -C  
# cldevice clear
```
 - If you are using Sun Cluster 3.1, use the following command:

```
# devfsadm -C  
# scdidadm -C
```
- 5 Repeat [Step 3](#) through [Step 4](#) for each node that is connected to the storage array.

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 31.

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 29.
- If your cluster does *not* use multipathing, see “[How to Replace a Node-to-Switch Component in a Cluster Without Multipathing](#)” on page 29.

▼ 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 ... ]
```

<i>nodename</i>	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.
<i>resourcegroup1</i> [<i>resourcegroup2</i> ...]	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 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>.

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

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

▼ How to Replace a Disk Drive

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

Note – Sun storage documentation uses the following terms:

- Logical volume
- Logical device
- Logical unit number (LUN)

This manual uses *logical volume* to refer to all such logical constructs.

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` RBAC authorization.**
- 2 If the failed disk drive affect the storage array logical volume's availability, if yes, use volume manager commands to detach the submirror or plex.**

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

- 3 If the logical volume (in [Step 1](#)) is configured as a quorum device, choose another volume 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 about how to add and remove quorum devices, see your Sun Cluster system administration documentation.

- 4 Replace the failed disk drive.**
- 5 (Optional) If the new disk drive is part of a logical volume that you want to be a quorum device, add the quorum device.**
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 Solaris Volume Manager or Veritas Volume Manager documentation.

Index

A

adding

See also installing

disk drives, 25

initiators, 18

storage arrays, 16-17

adding host to, host groups, 18

arrays, *See* storage arrays

B

batteries, replacing, 26

C

cables, replacing, 25

chassis, replacing, 25

controllers, replacing, 26

cooling units, replacing, 26

creating

host groups, 18

logical volumes, 18-20

storage pools, 18

D

deleting logical volumes, 20-23

disk drives

adding, 25

removing, 26

disk drives (*Continued*)

replacing, 34-35

E

Ethernet cables, replacing, 25

Ethernet hubs, replacing, 25

expansion cabinets, replacing, 25

F

fan assemblies, replacing, 25

firmware, upgrade storage array firmware, 26-27

FRUs, 25

H

HBAs, replacing, 31-34

help, 7

host adapters, replacing, 31-34

host groups, 18

I

initiators, 18

installing

See also adding

storage arrays, 14-16

L

logical devices, *See* logical volumes
logical unit numbers, *See* logical volumes
logical volumes
 creating, 18-20
 removing, 20-23
LUNs, *See* logical volumes

M

midplane, replacing, 25

N

node-to-switch components, replacing, 28-31
nodes, adding and removing, 26

P

pools, 18
power cables, replacing, 26
power units, replacing, 26

R

removing
 disk drives, 26
 host groups, 18
 initiators, 18
 logical volumes, 20-23
 nodes, 26
 storage arrays, 27-28
 storage pools, 18
removing host from, host groups, 18
replacing
 batteries, 26
 chassis, 25
 controllers, 26
 cooling units, 26
 disk drives, 34-35
 Ethernet cables, 25

replacing (*Continued*)

 Ethernet hubs, 25
 expansion cabinets, 25
 fan assemblies, 25
 host adapters, 31-34
 midplane, 25
 node-to-switch components, 28-31
 power cables, 26
 power units, 26

S

storage array firmware, upgrading, 26-27
storage arrays
 adding, 16-17
 installing, 14-16
 removing, 27-28
storage-based replication, 11
storage pools, 18
systems, *See* storage arrays

T

technical support, 7

U

upgrading, storage array firmware, 26-27