



Sun Cluster 3.0-3.1 With Sun StorEdge 6320 System Manual for Solaris OS

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

Part No: 817-1702-13
November 2005, Revision A

Copyright 2005 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. 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 2005 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.

Certaines 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. 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. Sun StorEdge,

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.



051024@13215



Contents

Preface	5
1 Installing and Maintaining a Sun StorEdge 6320 System	11
Installing Storage Systems	11
▼ How to Install Storage Systems in a New Cluster	12
How to Add Storage Systems to an Existing Cluster	14
▼ How to Perform Initial Configuration Tasks on the Storage Array	14
▼ How to Connect the Node to the FC Switches	15
Configuring Storage Systems	19
▼ How to Create a Logical Volume	19
▼ How to Remove a Logical Volume	21
Maintaining Storage Systems	23
StorEdge 6320 System FRUs	24
▼ How to Upgrade Storage Array Firmware	25
▼ How to Remove a Storage System	25
Replacing a Node-to-Switch Component	26
▼ How to Replace a Node-to-Switch Component in a Cluster That Uses Multipathing	27
▼ How to Replace a Node-to-Switch Component in a Cluster Without Multipathing	27
▼ How to Replace a Host Adapter	28
Index	31

Preface

The *Sun Cluster 3.0-3.1 With Sun StorEdge 6320 System Manual for Solaris OS* provides procedures that are specific to a Sun™ StorEdge™ 6320 system that is placed in a Sun™ Cluster environment.

Use this manual with any version of Sun Cluster 3.0 or 3.1 software in SPARC® based clusters and any supported version of Sun Cluster 3.1 software in x86 based clusters. Unless otherwise noted, procedures are the same for all supported Sun Cluster 3.x versions. See the [“Revision History” on page 6](#) for a list of changes to this manual.

Note – In this document, the term “x86” refers to the Intel 32-bit family of microprocessor chips and compatible microprocessor chips made by AMD.

Sun Cluster software runs on two platforms, SPARC and x86. 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 presales guide. You should have already determined your system requirements and purchased the appropriate equipment and software before reading this document.

How This Book Is Organized

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

Section 1 discusses how to install Sun StorEdge 6320 storage arrays.

Section 2 discusses how to configure logical units on Sun StorEdge 6320 storage arrays.

Section 3 describes how to maintain Sun StorEdge 6320 storage arrays in a running cluster.

Revision History

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

TABLE P-1 Sun Cluster 3.0-3.1 With Sun StorEdge 6320 System Manual for Solaris OS

Revision Date	New Information
November 2005	Expanded support to x86 based clusters.

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, 2004, and 2005 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 Traffic Manager Software Release Notes: For the Solaris Operating System</i>	817-0385
<i>Sun StorEdge Traffic Manager Installation and Configuration Guide</i>	816-1420
<i>Sun StorEdge 6320 System Reference and Service Manual</i>	816-7879
Available on http://www.sun.com/products-n-solutions/hardware/docs	
<i>Sun StorEdge 6320 System Installation Guide</i>	816-7878
Available on http://www.sun.com/products-n-solutions/hardware/docs	
<i>Sun StorEdge 6320 Series Site Preparation Guide</i>	816-7877
Available on http://www.sun.com/products-n-solutions/hardware/docs	
<i>Sun StorEdge 6020 and 6120 Array System Manual</i>	817-0200
Available on http://www.sun.com/products-n-solutions/hardware/docs	
<i>Sun StorEdge 6120 Array Installation Guide</i>	817-0199
Available on http://www.sun.com/products-n-solutions/hardware/docs	
<i>Sun StorEdge 6000 Family Host Installation Software Guide</i>	817-1739
Available on http://www.sun.com/products-n-solutions/hardware/docs	

TABLE P-3 Sun Cluster Documentation

Application	Title
Concepts	<i>Sun Cluster Concepts Guide for Solaris OS</i>
Overview	<i>Sun Cluster Overview for Solaris OS</i>
Hardware administration	<i>Sun Cluster 3.0-3.1 Hardware Administration Manual for Solaris OS</i> Individual hardware administration guides
Software installation	<i>Sun Cluster Software Installation Guide for Solaris OS</i>
Data service administration	<i>Sun Cluster Data Services Planning and Administration Guide for Solaris OS</i> Individual data service guides

TABLE P-3 Sun Cluster Documentation (Continued)

Application	Title
Data service development	<i>Sun Cluster Data Services Developer's Guide for Solaris OS</i>
System administration	<i>Sun Cluster System Administration Guide for Solaris OS</i>
Error messages	<i>Sun Cluster Error Messages Guide for Solaris OS</i>
Command and function reference	<i>Sun Cluster Reference Manual for Solaris OS</i>
Release Notes	<i>Sun Cluster 3.1 8/05 Release Notes for Solaris OS</i> <i>Sun Cluster 3.0-3.1 Release Notes Supplement</i>

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 9)
- The release number of Sun Cluster (for example, Sun Cluster 3.1 8/05)

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/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/>)
- Support (<http://www.sun.com/support/>)
- Training (<http://www.sun.com/training/>)

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>
AaBbCc123	What you type, contrasted with onscreen computer output	<code>machine_name% su</code> Password:

TABLE P-4 Typographic Conventions (Continued)

Typeface	Meaning	Example
<i>aabbcc123</i>	Placeholder: replace with a real name or value	The command to remove a file is <code>rm filename</code> .
<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-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	#

Product Training

Sun Microsystems offers training in many Sun technologies through a variety of instructor-led courses and self-paced courses. For information about the training courses that Sun offers and to enroll in a class, visit Sun Training at <http://training.sun.com/>.

Installing and Maintaining a Sun StorEdge 6320 System

This chapter contains the procedures about how to install, configure, and maintain a Sun™ StorEdge™ 6320 system. These procedures are specific to a Sun™ Cluster environment.

This chapter contains the following main topics:

- “Installing Storage Systems” on page 11
- “Configuring Storage Systems” on page 19
- “Maintaining Storage Systems” on page 23

For detailed information about storage system architecture, features, and configuration utilities, see the *Sun StorEdge 6320 System Reference and Service Manual* and the *Sun StorEdge 6320 System Installation Guide*.

Installing Storage Systems

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

TABLE 1-1 Task Map: Installing Storage Systems

Task	Information
Install storage systems in a new cluster, before the OS and Sun Cluster software is installed.	“How to Install Storage Systems in a New Cluster” on page 12
Add storage systems to an existing cluster.	“How to Add Storage Systems to an Existing Cluster” on page 14

You can install your storage system in several different configurations. Evaluate your needs. Determine which configuration is best for your situation. See the *Sun StorEdge 6320 System Installation Guide*, and “Installing Storage Arrays” in *Sun Cluster 3.0-3.1 With Sun StorEdge 6120 Array Manual for Solaris OS*.

▼ How to Install Storage Systems in a New Cluster

Use this procedure to install a storage system *before* you install the Solaris operating environment and Sun Cluster software on your nodes. To add a storage system to an existing cluster, use the procedure in “How to Add Storage Systems to an Existing Cluster” on page 14.

- Steps**
- 1. If necessary, install host adapters in the nodes to be connected to the storage system.**
For the procedure about how to install host adapters, see the documentation that shipped with your host adapters and nodes.
 - 2. (StorEdge 6320SL storage system ONLY) Install the Fibre Channel (FC) switch for the storage system if you do not have a switch installed.**

Note – In a StorEdge 6320SL storage system, the customer provides the switch.

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

- 3. Unpack, place, and level the storage system.**
For instructions, see the *Sun StorEdge 6320 System Installation Guide*.
- 4. Install the system power cord and the system grounding strap.**
For instructions, see the *Sun StorEdge 6320 System Installation Guide*.
- 5. (StorEdge 6320SL storage system ONLY) Connect the storage arrays to the FC switches by using fiber-optic cables.**



Caution – Do not connect the switch’s Ethernet port to the storage system’s private LAN.

For the procedure about how to cable the storage system, see the *Sun StorEdge 6320 System Installation Guide*.

- 6. Power on the storage system and the nodes.**
For instructions about how to power on the storage system, see the *Sun StorEdge 6320 System Installation Guide*. For instructions about how to power on a node, see the documentation that shipped with your node hardware.
- 7. Configure the service processor.**
For more information, see the *Sun StorEdge 6320 System Installation Guide*.

8. Create a volume.

For the procedure about how to create a volume, see the *Sun StorEdge 6320 System Reference and Service Manual*.

9. (Optional) Specify initiator groups for the volume.

For the procedure about how to specify initiator groups, see the *Sun StorEdge 6320 System Reference and Service Manual*.

10. If necessary, reconfigure the storage system's FC switches to ensure that all nodes can access each storage array.

The following configurations might prevent some nodes from accessing each storage array in the cluster.

- Zone configuration
- Multiple clusters that use same switch
- Unconfigured ports or misconfigured ports

11. 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*.

12. Install any required patches or software for Sun StorEdge Traffic Manager software support to nodes and enable multipathing.

For the procedure about how to install the Sun StorEdge Traffic Manager software, see "How to Install Sun Multipathing Software" in *Sun Cluster Software Installation Guide for Solaris OS*.

13. Configure the STMS paths.

```
cfgadm -c configure controllerinstance
```

For the procedure about how to configure STMS paths, see the *Sun StorEdge Traffic Manager Installation and Configuration Guide*.

14. Update the Solaris device files and links.

```
# devfsadm
```

Note – You can wait for the `devfsadm` daemon to automatically update the Solaris device files and links, or you can run the `devfsadm` command to immediately update the Solaris device files and links.

15. Confirm that all storage arrays that you installed are visible to all nodes.

```
# luxadm probe
```

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

How to Add Storage Systems to an Existing Cluster

Use this procedure to add a new storage system to a running cluster. To install systems to a new Sun Cluster configuration that is not running, use the procedure in [“How to Install Storage Systems in a New Cluster”](#) on page 12.

This procedure defines `Node N` as the node to be connected to the storage system you are adding and the node with which you begin working.

▼ How to Perform Initial Configuration Tasks on the Storage Array

- Steps**
1. *(StorEdge 6320SL storage system ONLY)* **Install the Fibre Channel (FC) switch for the storage system if you do not have a switch installed.**

Note – In a StorEdge 6320SL storage system, the customer provides the switch.

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

2. **Configure the service processor.**
For more information, see the *Sun StorEdge 6320 System Installation Guide*.
3. **Create a volume.**
For the procedure about how to create a volume, see the *Sun StorEdge 6320 System Reference and Service Manual*.
4. **(Optional) Specify initiator groups for the volume.**
For the procedure about how to specify initiator groups, see the *Sun StorEdge 6320 System Reference and Service Manual*.
5. **Unpack, place, and level the storage system.**
For instructions, see the *Sun StorEdge 6320 System Installation Guide*.
6. **Install the system power cord and the system grounding strap.**
For instructions, see the *Sun StorEdge 6320 System Installation Guide*.
7. *(StorEdge 6320SL storage system ONLY)* **Connect the storage arrays to the FC switches by using fiber-optic cables.**



Caution – Do not connect the switch’s Ethernet port to the storage system’s private LAN.

For the procedure about how to cable the storage system, see the *Sun StorEdge 6320 System Installation Guide*.

8. Power on the new storage system.

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

For the procedure about how to power on the storage system, see the *Sun StorEdge 6320 System Installation Guide*.

9. If necessary, reconfigure the storage system’s FC switches to ensure that all nodes can access each storage array.

The following configurations might prevent some nodes from accessing each storage array in the cluster:

- Zone configuration
- Multiple clusters that use the same switch
- Unconfigured ports or misconfigured ports

▼ How to Connect the Node to the FC Switches

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

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

```
# scstat
```

2. Move all resource groups and device groups off Node N.

```
# scswitch -s -h from-node
```

3. Do you need to install one or more host adapters in Node N?

- If no, skip to [Step 10](#).
- If yes, proceed to [Step 4](#).

4. Is the host adapter that you that are installing the first host adapter on Node N?

- If no, skip to [Step 6](#).

- If yes, determine whether the required drivers for the host adapter are already installed on this node. For the required packages, see the documentation that shipped with your host adapters.

5. Are the required support packages already installed?

- If yes, proceed to [Step 6](#).
- If no, install the packages.

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

6. Shut down and power off Node N.

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

7. Install one or more host adapters in Node N.

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

8. Power on and boot Node N into noncluster mode.

For more information about how to boot nodes, see your Sun Cluster system administration documentation.

9. If necessary, upgrade the host adapter firmware on Node N.

PatchPro is a patch-management tool that eases the selection and download of patches required for installation or maintenance of Sun Cluster software. PatchPro provides an Interactive Mode tool especially for Sun Cluster. The Interactive Tool makes the installation of patches easier. PatchPro's Expert Mode tool helps you to maintain your configuration with the latest set of patches. Expert Mode is especially useful for obtaining all of the latest patches, not just the high availability and security patches.

To access the PatchPro tool for Sun Cluster software, go to <http://www.sun.com/PatchPro/>, click Sun Cluster, then choose either Interactive Mode or Expert Mode. Follow the instructions in the PatchPro tool to describe your cluster configuration and download the patches.

For third-party firmware patches, see the SunSolveSM Online site at <http://sunsolve.ebay.sun.com>.

10. If necessary, install a GBIC or an SFP in the FC switch or the storage array.

For the procedure about how to install a GBIC or an SFP, see the documentation that shipped with your FC switch hardware or the *Sun StorEdge 6320 System Installation Guide*.

11. Connect a fiber-optic cable between the FC switch and Node N.

For the procedure about how to install a fiber-optic cable, see the *Sun StorEdge 6320 System Installation Guide*.

12. Install the required Solaris patches for storage array support on Node N.

PatchPro is a patch-management tool that eases the selection and download of patches required for installation or maintenance of Sun Cluster software. PatchPro provides an Interactive Mode tool especially for Sun Cluster. The Interactive Tool makes the installation of patches easier. PatchPro's Expert Mode tool helps you to maintain your configuration with the latest set of patches. Expert Mode is especially useful for obtaining all of the latest patches, not just the high availability and security patches.

To access the PatchPro tool for Sun Cluster software, go to <http://www.sun.com/PatchPro/>, click Sun Cluster, then choose either Interactive Mode or Expert Mode. Follow the instructions in the PatchPro tool to describe your cluster configuration and download the patches.

For third-party firmware patches, see the SunSolveSM Online site at <http://sunsolve.ebay.sun.com>.

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

```
# boot -r
```

14. Are any devices that are associated with the volume you created at an unconfigured state?

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

- If no, proceed to [Step 15](#).
- If yes, configure the STMS paths.

```
cfgadm -c configure controllerinstance
```

For the procedure about how to configure STMS paths, see the *Sun StorEdge Traffic Manager Installation and Configuration Guide*.

Note – You need to reboot if the `cfgadm` command does not configure the unconfigured devices that are associated with the volume you are creating. See the *Sun StorEdge Traffic Manager Installation and Configuration Guide* for more information.

15. Update the Solaris device files and links.

```
# devfsadm
```

Note – You can wait for the `devfsadm` daemon to automatically update the Solaris device files and links, or you can run the `devfsadm` command to immediately update the Solaris device files and links.

16. On Node N, update the paths to the DID instances.

```
# scgdevs
```

17. If necessary, label the new storage array logical volume.

For the procedure about how to label a logical volume, see the *Sun StorEdge 6320 System Reference and Service Manual*.

18. (Optional) On Node N, verify that the device IDs (DIDs) are assigned to the new storage array.

```
# scdidadm -l
```

19. Repeat Step 2 through Step 18 for each remaining node that you plan to connect to the storage array.

20. (Optional) Return the resource groups and device groups that you identified in Step 1 to the original nodes.

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

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

21. Perform volume management administration to incorporate the new volumes into the cluster.

For more information, see your Solstice DiskSuite/Solaris Volume Manager 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 Systems

This section contains the procedures about how to configure a storage system in a running cluster. [Table 1–2](#) lists these procedures.

TABLE 1–2 Task Map: Configuring Storage Arrays

Task	Information
Create a volume.	“How to Create a Logical Volume” on page 19
Remove a volume.	“How to Remove a Logical Volume” on page 21

The following is a list of administrative tasks that require no cluster-specific procedures. See the *Sun StorEdge 6320 System Reference and Service Manual* for the following procedures.

- Creating a storage pool.
- Removing a storage pool.
- Creating a volume group.
- Removing a volume group.
- Creating an initiator group.
- Adding an initiator group.
- Removing an initiator group.

▼ 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 devices is configured as described in the installation procedure.
- If you are using Sun StorEdge Traffic Manager, it is installed and configured and the path between Traffic Manager and the storage device is functioning.

Steps

1. 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.

2. Are you using multipathing?

- If yes, proceed to [Step 3](#).
- If no, skip to [Step 4](#).

3. Are any devices that are associated with the volume you created at an unconfigured state?

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

- If no, proceed to [Step 4](#).
- If yes, configure the Traffic Manager paths on each node that is connected to the storage device.

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

For the procedure about how to configure Traffic Manager paths, see the *Sun StorEdge Traffic Manager Installation and Configuration Guide*.

4. On one node that is connected to the storage device, use the `format` command to label the new logical volume.

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

```
# 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.

6. **To manage this volume with volume management software, use the appropriate Solstice DiskSuite/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 Solstice DiskSuite/Solaris Volume Manager or VERITAS Volume Manager documentation.

- See Also**
- To configure a logical volume as a quorum device, see Chapter 5, “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.

- Steps**
1. **Identify the logical volume that you are removing.**
Refer to your Solstice DiskSuite/Solaris Volume Manager or VERITAS Volume Manager documentation for more information.
 2. **(Optional) Migrate all data off the logical volume that you are removing. Alternatively, back up that data.**

3. Check if the logical volume that you are removing is a quorum device.

```
# scstat -q
```

If yes, choose and configure another device as the quorum device. Then remove the old quorum device.

For procedures about how to add and remove quorum devices, see Chapter 5, “Administering Quorum,” in *Sun Cluster System Administration Guide for Solaris OS*.

4. 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 Solstice DiskSuite/Solaris Volume Manager or VERITAS Volume Manager documentation.

5. If you are using volume management software, run the appropriate Solstice DiskSuite/Solaris Volume Manager or VERITAS Volume Manager commands to remove the logical volume from any diskset or disk group.

For more information, see your Solstice DiskSuite/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

6. If you are using multipathing, unconfigure the volume in Sun StorEdge Traffic Manager.

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

7. Access the storage device and remove the logical volume.

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

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

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

```
# scstat
```

9. Move all resource groups and device groups off Node A.

```
# scswitch -s -h from-node
```

10. Shut down and reboot Node A.

For the 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*.

11. On Node A, remove the paths to the logical volume that you removed. Remove obsolete device IDs.

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

12. For each additional node that is connected to the shared storage that hosted the logical volume, repeat Step 8 to Step 11.

13. (Optional) Return the resource groups and device groups that you identified in Step 8 to all cluster nodes.

Maintaining Storage Systems

This section contains the procedures for maintaining a storage system in a running cluster. Table 1-3 lists these procedures.

TABLE 1-3 Task Map: Maintaining Storage Systems

Task	Information
Remove a storage system.	“How to Remove a Storage System” on page 25
Upgrade storage array firmware.	“How to Upgrade Storage Array Firmware” on page 25
Replace a node-to-switch component. <ul style="list-style-type: none">■ Node-to-switch fiber-optic cable■ FC host adapter■ FC switch■ GBIC or SFP	“How to Replace a Node-to-Switch Component in a Cluster Without Multipathing” on page 27
Replace a node’s host adapter.	“How to Replace a Host Adapter” on page 28
Add a node to the storage array.	Sun Cluster system administration documentation

TABLE 1-3 Task Map: Maintaining Storage Systems *(Continued)*

Task	Information
Remove a node from the storage array.	Sun Cluster system administration documentation

StorEdge 6320 System FRUs

The following is a list of administrative tasks that require no cluster-specific procedures. See the *Sun StorEdge 6320 System Reference and Service Manual* for the following procedures.

- Adding a disk drive.
- Replacing a failed disk drive.
- Removing a disk drive.
- Replacing an Ethernet cable.
- Replacing an Ethernet hub in a cabinet.
- Replacing the power cable on the storage array.
- Replacing a power and cooling unit (PCU).
- Replacing a storage service processor.
- Replacing a storage service processor accessory tray.
- Replacing a service processor panel.
- Replacing a USB relay panel.
- Replacing a USB flash disk.
- Extending a dual-controller configuration.
- Removing a dual-controller configuration.
- Replacing a chassis.
- Replacing a controller.
- Replacing a storage array-to-switch component in a dual-controller configuration.

A storage array-to-switch component includes all of the following components:

- Fiber-optic cable that connects an FC switch to a storage array
- FC switch GBIC or an SFP that connects an FC switch to a storage array
- FC switch that connects an FC switch to a storage array
- FC switch power cord
- Interconnect cables that connect two storage arrays
- Interconnect cards

▼ 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 `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 `scdidadm -R` command for each affected device.

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

2. Apply the controller, disk drive, and loop-card firmware patches.

For the list of required patches, see the *Sun StorEdge 6320 System Reference and Service Manual*. For the procedure about how to apply firmware patches, see the firmware patch README file. For the procedure about how to verify the firmware level, see the *Sun StorEdge 6320 System Reference and Service Manual*.

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 System

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

This procedure defines Node N as the node that is connected to the storage system 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 system that you are removing.

- Steps**
1. If necessary, back up all database tables, data services, and volumes that are associated with each partner group that is affected.
 2. Remove references to the volumes that reside on the storage system that you are removing.
For more information, see your Solstice DiskSuite/Solaris Volume Manager or VERITAS Volume Manager documentation.
 3. Disconnect the cables that connected **Node N** to the FC switches in your storage system.
 4. On all nodes, remove the obsolete Solaris links and device IDs (DIDs).

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

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

For the procedure about how to replace a host adapter, see [“How to Replace a Host Adapter”](#) on page 28.

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 27.
- If your cluster does *not* use multipathing, see [“How to Replace a Node-to-Switch Component in a Cluster Without Multipathing”](#) on page 27.

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

- Steps**
1. **Is your configuration active-passive?**
If yes, and 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

- Steps**
1. **Check if the physical data path failed.**
If no, proceed to [Step 2](#).
If yes:
 - a. **Replace the component.**
Refer to your hardware documentation for any component-specific instructions.
 - 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.
 2. **Determine the resource groups and device groups that are running on Node A.**

```
# scstat
```
 3. **Move all resource groups and device groups to another node.**

```
# scswitch -s -h from-node
```
 4. **Replace the node-to-switch component.**
Refer to your hardware documentation for any component-specific instructions.
 5. **(Optional) If necessary, return the resource groups and device groups that you identified in Step 2 to Node A.**

```
# scswitch -z -g resource-group -h nodename  
# swswitch -z -D device-group -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, see your *Sun Cluster Hardware Administration Manual for Solaris OS*, and omit steps that instruct you to shut down and power off the nodes.

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

Record this information because you use this information in [Step 9](#) of this procedure to return resource groups and device groups to Node A.

```
# scstat
```

2. Move all resource groups and device groups off Node A.

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

3. Shut down Node A.

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

4. Power off Node A.

5. Replace the failed host adapter.

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

6. Do you need to upgrade the node's host adapter firmware?

- **If yes, boot Node A into noncluster mode. Proceed to [Step 7](#).**

For more information about how to boot nodes, see your Sun Cluster system administration documentation.

- **If no, proceed to [Step 8](#).**

7. Upgrade the host adapter firmware on Node A.

PatchPro is a patch-management tool that eases the selection and download of patches required for installation or maintenance of Sun Cluster software. PatchPro provides an Interactive Mode tool especially for Sun Cluster. The Interactive Tool makes the installation of patches easier. PatchPro's Expert Mode tool helps you to maintain your configuration with the latest set of patches. Expert Mode is

especially useful for obtaining all of the latest patches, not just the high availability and security patches.

To access the PatchPro tool for Sun Cluster software, go to <http://www.sun.com/PatchPro/>, click Sun Cluster, then choose either Interactive Mode or Expert Mode. Follow the instructions in the PatchPro tool to describe your cluster configuration and download the patches.

For third-party firmware patches, see the SunSolveSM Online site at <http://sunsolve.ebay.sun.com>.

8. Boot Node A into cluster mode.

For more information about how to boot nodes, see your Sun Cluster system administration documentation.

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

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

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

Index

A

adding
 See also installing
 disk drive, 24
 initiator groups, 19
 nodes, 23
 partner groups, 24
 storage systems, 14-18
arrays, *See* storage systems

C

chassis, replacing, 24
configuring, storage systems, 19-23
controllers, replacing, 24
cooling units, replacing, 24
creating
 initiator groups, 19
 logical volumes, 19-21
 storage pools, 19
 volume groups, 19

D

deleting logical volumes, *See* removing logical volumes
disk drive, adding, 24
disk drives
 removing, 24
 replacing, 24

E

Ethernet cables, replacing, 24
Ethernet hubs, replacing, 24

F

firmware, upgrade storage array firmware, 25
FRUs, 24

H

HBAs, replacing, 28-29
help, 8-9
host adapters, replacing, 28-29

I

initiator groups
 creating, 19
 modifying, 19
 removing, 19
installing
 See also adding
 storage systems, 12-14
interconnect cables, replacing, 24
interconnect cards, replacing, 24

L

logical devices, *See* logical volumes
logical unit numbers
 See logical volumes
 See volumes
logical volumes
 creating, 19-21
 removing, 21-23
loop cables, *See* interconnect cables
loop cards, *See* interconnect cards
LUNs, *See* logical volumes

M

maintaining, storage systems, 23-29
midplane, replacing, 24

N

node-to-switch components, replacing, 26-27
nodes
 adding, 23
 removing, 24

P

partner groups
 adding, 24
 removing, 24
pools, *See* storage pools
power cables, replacing, 24
power units, replacing, 24

R

removing
 disk drives, 24
 initiator groups, 19
 logical volumes, 21-23
 nodes, 24
 partner groups, 24
 storage pools, 19
 storage systems, 25-26
 volume groups, 19

replacing

chassis, 24
controllers, 24
cooling units, 24
disk drives, 24
Ethernet cables, 24
Ethernet hubs, 24
host adapters, 28-29
midplane, 24
node-to-switch components, 26-27
power cables, 24
power units, 24
service processor accessory tray, 24
service processor panel, 24
storage service processor, 24
switch-to-storage array components, 24
USB flash disk, 24
USB relay panel, 24

S

service processor, replacing, 24
service processor accessory tray, replacing, 24
service processor panel, replacing, 24
SPA tray, *See* service processor accessory tray
SSP, *See* service processor
storage array firmware, upgrading, 25
storage pools
 creating, 19
 removing, 19
storage systems
 adding, 14-18
 configuring, 19-23
 installing, 12-14
 maintaining, 23-29
 removing, 25-26
switch-to-storage array components,
 replacing, 24
systems, *See* storage systems

T

technical support, 8-9
training, 10

U

- upgrading, storage array firmware, 25
- USB flash disk, replacing, 24
- USB relay panel, replacing, 24

V

- volume groups
 - creating, 19
 - removing, 19

