

Sun Cluster 3.1 - 3.2 With Fibre Channel JBOD Storage Device Manual

SPARC Platform Edition



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

Part No: 819-2994-10
December 2006, Revision A

Copyright 2006 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, Solstice, Solaris DiskSuite, Sun Cluster, SunSolve, Sun StorEdge, 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 2006 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 États-Unis et dans d'autres pays.

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

Certains composants de ce produit peuvent être dérivés du logiciel Berkeley BSD, licenciés par l'Université de Californie. UNIX est une marque déposée aux États-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, Solstice, Solaris DiskSuite, Sun Cluster, SunSolve, Sun StorEdge, Java et Solaris sont des marques de fabrique ou des marques déposées de Sun Microsystems, Inc. aux États-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 États-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 États-Unis, ou vers des entités figurant sur les listes d'exclusion d'exportation américaines, y compris, mais de manière non exclusive, la liste de personnes qui font objet d'un ordre de ne pas participer, d'une façon directe ou indirecte, aux exportations des produits ou des services qui sont régis par la législation américaine en matière de contrôle des exportations et la liste de ressortissants spécifiquement désignés, sont rigoureusement interdites.

LA DOCUMENTATION EST FOURNIE "EN L'ETAT" ET TOUTES AUTRES CONDITIONS, DECLARATIONS ET GARANTIES EXPRESSES OU TACITES SONT FORMELLEMENT EXCLUES, DANS LA MESURE AUTORISEE PAR LA LOI APPLICABLE, Y COMPRIS NOTAMMENT TOUTE GARANTIE IMPLICITE RELATIVE A LA QUALITE MARCHANDE, A L'APTITUDE A UNE UTILISATION PARTICULIERE OU A L'ABSENCE DE CONTREFACON.

Contents

Preface	5
1 Installing and Maintaining a Fibre Channel JBOD Storage Device	11
Installing Storage Arrays	11
▼ How to Install a Storage Array in a New Cluster	12
▼ How to Add the <i>First</i> Storage Array to an Existing Cluster	12
▼ How to Add a <i>Subsequent</i> Storage Array to an Existing Cluster	14
Maintaining Storage Arrays	16
FRUs That <i>Do Not</i> Require Sun Cluster Maintenance Procedures	16
▼ How to Replace a Storage Array	17
▼ How to Remove a Storage Array	19
▼ How to Add a Disk Drive	22
▼ How to Remove a Disk Drive	23
▼ How to Replace a Disk Drive	26
A Cabling Diagrams	31
Sun StorEdge A5x00 Cabling Diagrams	31
Installing a Storage Array	31
Adding a Storage System	32
Index	33

Preface

The *Sun Cluster 3.1 - 3.2 With Fibre Channel JBOD Storage Device Manual* provides procedures specific to FC JBOD storage devices that are placed in a Sun™ Cluster environment.

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 and one appendix.

- [Chapter 1](#) contains:
 - “[Installing Storage Arrays](#)” on [page 11](#) discusses how to install FC JBOD storage devices.
 - “[Maintaining Storage Arrays](#)” on [page 16](#) describes how to maintain FC JBOD storage devices in a running cluster.
- [Appendix A](#) illustrates installation and maintenance configurations.

Revision History

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

TABLE P-1 Sun Cluster 3.1 - 3.2 With Fibre Channel JBOD Storage Device Manual

Revision Date	New Information
	No revisions.

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.

The related books are organized by the following products and subjects:

- [Sun StorEdge A5x00](#)
- [Sun Cluster software and data services](#)

TABLE P-2 Sun StorEdge A5x00 Storage Array

Title	Part Number
<i>Sun StorEdge A5000 Configuration Guide</i>	805-0264
<i>Sun StorEdge A5000 Installation and Service Manual</i>	802-7573
Available at http://www.sun.com/products-n-solutions/hardware/docs/pdf/805-0070-11.pdf	
<i>Sun StorEdge A5000 Product Note</i>	805-1018

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 3.1 - 3.2 Hardware Administration Manual for Solaris OS</i> Individual hardware administration guides
Software installation	<i>Sun Cluster Software Installation Guide for Solaris OS</i>

TABLE P-3 Sun Cluster Documentation (Continued)

Application	Title
Data service administration	<i>Sun Cluster Data Services Planning and Administration Guide for Solaris OS</i> Individual data service guides
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 3.2 Release Notes for Solaris OS</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 8)
- 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/>)

Typographic Conventions

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

TABLE P-4 Typographic Conventions

Typeface	Meaning	Example
<code>AaBbCc123</code>	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:
<i>aabbcc123</i>	Placeholder: replace with a real name or value	The command to remove a file is <code>rm filename.</code>

TABLE P-4 Typographic Conventions (Continued)

Typeface	Meaning	Example
<i>AaBbCc123</i>	Book titles, new terms, and terms to be emphasized	Read Chapter 6 in the <i>User's Guide</i> . A <i>cache</i> is a copy that is stored locally. Do <i>not</i> save the file. Note: Some emphasized items appear bold online.

Shell Prompts in Command Examples

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

TABLE P-5 Shell Prompts

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

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 Fibre Channel JBOD Storage Device

This chapter describes the procedures about how to install, configure, and maintain fibre channel (FC) JBOD storage devices in a Sun™ Cluster environment.

The procedures in this chapter apply to the Sun StorEdge A5x00.

This chapter contains the following main sections:

- “Installing Storage Arrays” on page 11
- “Maintaining Storage Arrays” on page 16

For information about how to use storage arrays in a storage area network (SAN), see “SAN Solutions in a Sun Cluster Environment” in *Sun Cluster 3.1 - 3.2 Hardware Administration Manual for Solaris OS*.

Installing Storage Arrays

This section contains instructions on installing arrays both to new clusters and operational clusters.

TABLE 1-1 Task Map: Installing Storage Arrays

Task	Information
Install a storage array in a new cluster, before the OS and Sun Cluster software are installed.	“How to Install a Storage Array in a New Cluster” on page 12
Add a storage array to an operational cluster.	“How to Add the <i>First</i> Storage Array to an Existing Cluster” on page 12 “How to Add a <i>Subsequent</i> Storage Array to an Existing Cluster” on page 14

▼ How to Install a Storage Array in a New Cluster

This procedure assumes you are installing one or more storage arrays at initial installation of a cluster.

1 Install host adapters in the nodes that are to be connected to the storage array.

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

Note – To ensure maximum redundancy, put each host adapter on a separate I/O board, if possible.

2 Cable the storage arrays to the nodes.

For cabling diagrams, see [Appendix A](#).

3 Check the revision number for the storage array's controller firmware. If necessary, install the most recent firmware.

For more information, see your storage documentation. For a list of storage documentation, see “[Related Documentation](#)” on page 6.

▼ How to Add the *First* Storage Array to an Existing Cluster

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

- Your cluster is operational.
- You do *not* have an existing storage array that is installed and configured.

If you are installing a storage array in a running cluster that already has storage arrays installed and configured, use the procedure in “[How to Add a *Subsequent* Storage Array to an Existing Cluster](#)” on page 14.

1 Determine if the storage array packages need to be installed on the nodes. These nodes are the nodes to which you are connecting the storage array. This product requires the following packages.

```
# pkginfo | egrep wLux
system SUNWld Sun Enterprise Network Array sf Device Driver
system SUNWluxdx Sun Enterprise Network Array sf Device Driver
(64-bit)
system SUNWluxl Sun Enterprise Network Array socl Device Driver
system SUNWluxlx Sun Enterprise Network Array socl Device Driver
(64-bit)
system SUNWluxop Sun Enterprise Network Array firmware and utilities
system SUNWluxox Sun Enterprise Network array libraries (64 bit)
```

2 On each node, install any necessary packages for the Solaris Operating System.

The storage array packages are located in the Product directory of the CD-ROM. Use the pkgadd command to add any necessary packages.

Note – The -G option applies only if you are using the Solaris 10 OS. Omit this option if you are using Solaris 8 or 9 OS.

```
# pkgadd -G -d path_to_Solaris/Product Pkg1 Pkg2 Pkg3 ... PkgN
```

-G Add package(s) in the current zone only. When used in the global zone, the package is added to the global zone only and is not propagated to any existing or yet-to-be-created non-global zone. When used in non-global zone, the package(s) are added to the non-global zone only.

path_to_Solaris Path to the Solaris Operating System

Pkg1 Pkg2 The packages to be added

3 Shut down and power off any node that is connected to the storage array.

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

4 Install host adapters in the node that is to be connected to the storage array.

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

5 Cable, configure, and power on the storage array.

For cabling diagrams, see [Appendix A](#).

6 To create the new Solaris device files and links, perform a reconfiguration boot by adding - r to your boot instruction.

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

7 Determine if any patches need to be installed on nodes that are to be connected to the storage array.

For a list of patches specific to Sun Cluster, see your Sun Cluster release notes documentation.

8 Obtain and install any necessary patches on the nodes that are to be connected to the storage array.

For procedures about how to apply patches, see your Sun Cluster system administration documentation.

Note – Read any README files that accompany the patches before you begin this installation. Some patches must be installed in a specific order.

9 If required by the patch README instructions, shut down and reboot the node.

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*.

10 Perform Step 3 through Step 9 for each node that is attached to the storage array.**11 Perform volume management administration to add the disk drives in the storage array to the volume management configuration.**

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

▼ How to Add a *Subsequent* Storage Array to an Existing Cluster

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

- Your cluster is operational.
- You have an existing storage array that is installed and configured.

If you are installing a storage array in a running cluster that does not yet have a storage array that is installed, use the procedure in “[How to Add the *First* Storage Array to an Existing Cluster](#)” on page 12.

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

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

1 Configure the new storage array.

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

2 On both nodes, insert the new storage array into the cluster. Add paths to the disk drives.

```
# luxadm insert_device
```

Please hit <RETURN> when you have finished adding

```
Fibre Channel Enclosure(s)/Device(s):
```

Note – Do not press the Return key until you complete [Step 3](#).

3 Cable the new storage array to a spare port in the existing hub, switch, or host adapter in your cluster.

For cabling diagrams, see [Appendix A](#).

Note – You must use FC switches when installing storage arrays in a partner-group configuration. If you want to create a storage area network (SAN) by using two FC switches and Sun SAN software, see “SAN Solutions in a Sun Cluster Environment” in *Sun Cluster 3.1 - 3.2 Hardware Administration Manual for Solaris OS* for more information.

4 After you cable the new storage array, press the Return key to complete the `luxadm insert_device` operation.

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

5 On both nodes, verify that the new storage array is visible to both nodes.

```
# luxadm probe
```

6 On one node, update the DID database.

- 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
```

Maintaining Storage Arrays

The maintenance procedures in “FRUs That *Do Not* Require Sun Cluster Maintenance Procedures” on page 16 are performed the same as in a noncluster environment. Table 1–2 lists the procedures that require cluster-specific steps.

TABLE 1–2 Task Map: Maintaining a Storage Array

Task	Information
Remove a storage array	“How to Remove a Storage Array” on page 19
Replace a storage array	“How to Replace a Storage Array” on page 17
Add a disk drive	“How to Add a Disk Drive” on page 22
Remove a disk drive	“How to Remove a Disk Drive” on page 23
Replace a disk drive	“How to Replace a Disk Drive” on page 26

FRUs That *Do Not* Require Sun Cluster Maintenance Procedures

Each storage device has a different set of FRUs that do not require cluster-specific procedures.

Sun StorEdge A5x00 FRUs

The following is a list of administrative tasks that require no cluster-specific procedures. See the *Sun StorEdge A5000 Installation and Service Manual* for the following procedures.

- Replace an AC power filter
- Replace a backplane
- Replace a chassis
- Replace a door panel assembly
- Replace a fan tray
- Remove a fiber-optic cable
- Replace an FPM
- Remove a GBIC
- Replace an interconnect assembly
- Replace an interface board
- Replace a power supply

▼ How to Replace a Storage Array

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

- Your cluster is operational.
- You want to retain the existing disk drives in the storage array.

If you want to replace your disk drives, see [“How to Replace a Disk Drive” on page 26](#).

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

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

[Example 1–1](#) shows you how to apply this procedure.

1 If possible, back up the metadevices or volumes that reside in the storage array.

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

2 Perform volume management administration to remove the storage array from the configuration.

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

3 On all nodes that are connected to the storage array, run the `luxadm remove_device` command.

```
# luxadm remove_device -F boxname
```

See [Example 1–1](#) for an example of this command and its use.

4 Disconnect the fiber-optic cables from the storage array.

5 Power off and disconnect the storage array from the AC power source.

For more information, see your storage documentation. For a list of storage documentation, see [“Related Documentation” on page 6](#).

6 Connect the fiber optic cables to the new storage array.

7 Connect the new storage array to an AC power source.

8 One disk drive at a time, remove the disk drives from the old storage array. Insert the disk drives into the same slots in the new storage array.

9 Power on the storage array.

10 Use the `luxadm insert_device` command to find the new storage array.

Repeat this step for each node that is connected to the storage array.

```
# luxadm insert_device
```

See [Example 1–1](#) for an example of this command and its use.

11 On all nodes that are connected to the new storage array, upload the new information to the DID driver.

If a volume management daemon such as `voId` is running on your node, and you have a CD-ROM drive that is connected to the node, a device busy error might be returned even if no disk is in the drive. This error is an expected behavior.

- **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
```

12 Perform volume management administration to add the new storage array to the configuration.

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

Example 1–1 Replacing a Sun StorEdge A5x00 Storage Array When Using Sun Cluster 3.2 Software

The following example shows how to replace a Sun StorEdge A5x00 storage array when using Sun Cluster 3.2 software. The storage array to be replaced is `venus1`.

```
# luxadm remove_device -F venus1
```

WARNING!!! Please ensure that no filesystems are mounted on these device(s).
All data on these devices should have been backed up.

The list of devices that will be removed is:

```
1: Box name:      venus1
   Node WWN:      123456789abcdef
   Device Type:   SENA (SES device)
   SES Paths:
           /devices/nodes@1/sbus@1f,0/SUNW,socal@1,0/sf@0,0/ \
           ses@w123456789abcdf03,0:0
           /devices/nodes@1/sbus@1f,0/SUNW,socal@1,0/sf@1,0/ \
           ses@w123456789abcdf00,0:0
```

Please verify the above list of devices and then enter 'c' or <CR> to Continue or 'q' to Quit. [Default: c]:

```

<Return>
Hit <Return> after removing the device(s). <Return>

# luxadm insert_device
Please hit <RETURN> when you have finished adding Fibre Channel
Enclosure(s)/Device(s): <Return>
# cldevice populate      If using Sun Cluster 3.1, use scgdevs.

```

Example 1-2 Replacing a Sun StorEdge A5x00 Storage Array When Using Sun Cluster 3.1 Software

The following example shows how to replace a Sun StorEdge A5x00 storage array when using Sun Cluster 3.1 software. The storage array to be replaced is venus1.

```
# luxadm remove_device -F venus1
```

WARNING!!! Please ensure that no filesystems are mounted on these device(s).
All data on these devices should have been backed up.

The list of devices that will be removed is:

```

1: Box name:      venus1
   Node WWN:      123456789abcdeff
   Device Type:   SENA (SES device)
   SES Paths:
       /devices/nodes@1/sbus@1f,0/SUNW,socal@1,0/sf@0,0/ \
       ses@w123456789abcdf03,0:0
       /devices/nodes@1/sbus@1f,0/SUNW,socal@1,0/sf@1,0/ \
       ses@w123456789abcdf00,0:0

```

Please verify the above list of devices and then enter 'c' or <CR> to Continue or 'q' to Quit. [Default: c]:

```

<Return>
Hit <Return> after removing the device(s). <Return>

```

```

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

```

▼ How to Remove a Storage Array

Use this procedure to remove a storage array from a cluster. [Example 1-3](#) shows you how to apply this procedure. Use the procedures in your server hardware manual to identify the storage array.

This procedure provides the long forms of the Sun Cluster commands. Most commands also have short forms. Except for the forms of the command names, the commands are identical. For a list of

the commands and their short forms, see Appendix A, “Sun Cluster Object-Oriented Commands,” in *Sun Cluster 3.1 - 3.2 Hardware Administration Manual for Solaris OS*.

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

1 Perform volume management administration to remove the storage array from the configuration.

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

2 On all nodes that are connected to the storage array, run the `luxadm remove_device` command.

```
# luxadm remove_device -F boxname
```

3 Remove the storage array and the fiber-optic cables that are connected to the storage array.

For more information, see your storage documentation. For a list of storage documentation, see “Related Documentation” on page 6.

Note – If you are using your storage arrays in a SAN-configured cluster, you must keep two FC switches configured in parallel. This configuration maintains cluster availability. See “SAN Solutions in a Sun Cluster Environment” in *Sun Cluster 3.1 - 3.2 Hardware Administration Manual for Solaris OS* for more information.

4 On all nodes, remove references to the storage array.

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

```
# devfsadm -C
# cldevice populate
```

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

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

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

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

Example 1–3 Removing a Sun StorEdge A5x00 Storage Array When Using Sun Cluster 3.2 Software

The following example shows how to remove a Sun StorEdge A5x00 storage array from a cluster running Sun Cluster version 3.2 software. The storage array to be removed is `venus1`.

```
# luxadm remove_device -F venus1
WARNING!!! Please ensure that no file systems are mounted on these device(s).
```

All data on these devices should have been backed up.

The list of devices that will be removed is:

```
1: Storage Array:   venus1
   Node WWN:       123456789abcdeff
   Device Type:    SENA (SES device)
   SES Paths:
     /devices/nodes@1/sbus@1f,0/SUNW,socal@1,0/sf@0,0/ \
       ses@w123456789abcdef03,0:0
     /devices/nodes@1/sbus@1f,0/SUNW,socal@1,0/sf@1,0/ \
       ses@w123456789abcdef00,0:0
```

Please verify the above list of devices and then enter 'c' or <CR> to Continue or 'q' to Quit. [Default: c]: <Return>

Hit <Return> after removing the device(s). <Return>

```
# devfsadm -C
# cldevice populate
```

Example 1-4 Removing a Sun StorEdge A5x00 Storage Array When Using Sun Cluster 3.1 Software

The following example shows how to remove a Sun StorEdge A5x00 storage array from a cluster running Sun Cluster version 3.1 software. The storage array to be removed is venus1.

```
# luxadm remove_device -F venus1
WARNING!!! Please ensure that no file systems are mounted on these device(s).
All data on these devices should have been backed up.
```

The list of devices that will be removed is:

```
1: Storage Array:   venus1
   Node WWN:       123456789abcdeff
   Device Type:    SENA (SES device)
   SES Paths:
     /devices/nodes@1/sbus@1f,0/SUNW,socal@1,0/sf@0,0/ \
       ses@w123456789abcdef03,0:0
     /devices/nodes@1/sbus@1f,0/SUNW,socal@1,0/sf@1,0/ \
       ses@w123456789abcdef00,0:0
```

Please verify the above list of devices and then enter 'c' or <CR> to Continue or 'q' to Quit. [Default: c]: <Return>

Hit <Return> after removing the device(s). <Return>

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

▼ How to Add a Disk Drive

For conceptual information about quorums, quorum devices, global devices, and device IDs, see your Sun Cluster concepts documentation. For a list of Sun Cluster documentation, see [“Related Documentation”](#) on page 6.

Before You Begin This procedure assumes that your cluster is 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*.

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

1 On one node that is connected to the storage array, install the new disk.

Install the new disk drive. Press the Return key when prompted. You can insert multiple disk drives at the same time.

```
# luxadm insert_device enclosure,slot
```

2 On all other nodes that are attached to the storage array, probe all devices. Write the new disk drive to the `/dev/rdisk` directory.

The amount of time that the `devfsadm` command requires to complete its processing depends on the number of devices that are connected to the node. Expect at least five minutes.

```
# devfsadm -C
```

3 Ensure that entries for the disk drive have been added to the `/dev/rdisk` directory.

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

4 If necessary, partition the disk drive.

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

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

If a volume management daemon such as `vold` is running on your node, and you have a CD-ROM drive connected to the node, a device busy error might be returned even if no disk is in the drive. This error is an expected behavior.

- 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`

6 Verify that a device ID (DID) has been assigned to the disk drive.

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

- **If you are using Sun Cluster 3.2, use the following command:**
`cldevice list -v`
- **If you are using Sun Cluster 3.1, use the following command:**
`scdidadm -l`

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

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

▼ How to Remove a Disk Drive

For conceptual information about quorum, quorum devices, global devices, and device IDs, see your Sun Cluster concepts documentation. For a list of Sun Cluster documentation, see [“Related Documentation” on page 6](#).

[Example 1–5](#) shows you how to apply this procedure.

Before You Begin This procedure assumes that your cluster is 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*.

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

1 Determine whether the disk drive that you want to remove is configured as a quorum device?

- **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`

- 2 If the disk drive you want to remove is configured as a quorum device, choose and configure another device to be the new quorum device. Then remove the old quorum device.**

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

- 3 If possible, back up the metadvice or volume.**

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

- 4 Perform volume management administration to remove the disk drive from the configuration.**

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

- 5 Identify the disk drive that needs to be removed.**

If the disk error message reports the drive problem by DID, determine the Solaris device name.

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

```
# cldevice list -v
```

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

```
# sctdidadm -l deviceID
```

- 6 On any node that is connected to the storage array, run the `luxadm remove_device` command.**

Remove the disk drive. Press the Return key when prompted.

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

- 7 On all connected nodes, remove references to the disk drive.**

- **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  
# sctdidadm -C
```

Example 1–5 Removing a Disk Drive in a Sun StorEdge A5x00 Storage Array When Using Sun Cluster 3.2 Software

The following example shows how to remove a disk drive from a Sun StorEdge A5x00 storage array in a cluster running Sun Cluster version 3.2 software. The disk drive to be removed is d4 and is a virtual table of contents (VTOC) labelled device.

```

# cldevice list -v

=== DID Device Instances ===

DID Device Name:                /dev/did/rdisk/d4
Full Device Path:               phys0-schost1:/dev/rdisk/clt1d0
Full Device Path:               phys-schost2:/dev/rdisk/clt1d0
Replication:                    none
default_fencing:                global

# luxadm remove_device -F /dev/rdisk/clt32d0s2

WARNING!!! Please ensure that no file systems are mounted on these device(s).
All data on these devices should have been backed up.

The list of devices that will be removed is:
    1: Box Name "venus1" front slot 0

Please enter 'q' to Quit or <Return> to Continue: <Return>

stopping: Drive in "venus1" front slot 0...Done
offlining: Drive in "venus1" front slot 0...Done

Hit <Return> after removing the device(s). <Return>

Drive in Box Name "venus1" front slot 0
Logical Nodes being removed under /dev/dsk/ and /dev/rdisk:
    clt32d0s0
    clt32d0s1
    clt32d0s2
    clt32d0s3
    clt32d0s4
    clt32d0s5
    clt32d0s6
    clt32d0s7

# devfsadm -C
# cldevice clear

```

Example 1–6 Removing a Disk Drive in a Sun StorEdge A5x00 Storage Array When Using Sun Cluster Version 3.1 Software

The following example shows how to remove a disk drive from a Sun StorEdge A5x00 storage array in a cluster running Sun Cluster version 3.1 software. The disk drive to be removed is d4 and is a virtual table of contents (VTOC) labelled device.

```

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

```

WARNING!!! Please ensure that no file systems are mounted on these device(s).
All data on these devices should have been backed up.

The list of devices that will be removed is:

1: Box Name "venus1" front slot 0

Please enter 'q' to Quit or <Return> to Continue: **<Return>**

stopping: Drive in "venus1" front slot 0...Done

offlining: Drive in "venus1" front slot 0...Done

Hit <Return> after removing the device(s). **<Return>**

Drive in Box Name "venus1" front slot 0

Logical Nodes being removed under /dev/dsk/ and /dev/rdisk:

c1t32d0s0

c1t32d0s1

c1t32d0s2

c1t32d0s3

c1t32d0s4

c1t32d0s5

c1t32d0s6

c1t32d0s7

devfsadm -C

scdidadm -C

▼ How to Replace a Disk Drive

For conceptual information about quorum, quorum devices, global devices, and device IDs, see your Sun Cluster concepts documentation.

Before You Begin This procedure assumes that your cluster is 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*.

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

1 Identify the disk drive that needs replacement.

If the disk error message reports the drive problem by device ID (DID), determine the Solaris logical device name. If the disk error message reports the drive problem by the Solaris physical device name,

use your Solaris documentation to map the Solaris physical device name to the Solaris logical device name. Use this Solaris *logical* device name and DID throughout this procedure.

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

```
# cldevice list -v
```

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

```
# scdidadm -l deviceID
```

2 Determine whether the disk drive that you want to replace is configured as a quorum device.

- **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
```

3 If the disk drive that you want to replace is configured as a quorum device, add a new quorum device on a different storage array. Remove the old quorum device.

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

4 If possible, back up the metadvice or volume.

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

5 Identify the failed disk drive's physical DID.

Use this physical DID in [Step 14](#) to verify that the failed disk drive has been replaced with a new disk drive. The DID and the world wide name (WWN) for the disk drive are the same.

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

```
# cldevice list -v
```

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

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

6 If you are using VERITAS Volume Manager, proceed to [Step 8](#).

7 If you are using Solstice DiskSuite/Solaris Volume Manager, save the disk partitioning information to partition the new disk drive.

```
# prtvtoc /dev/rdisk/cNtXdYs2 > filename
```

Note – You can also use the `format` utility to save the disk's partition information.

- 8** On any node that is connected to the storage array, remove the disk drive when prompted.

```
# luxadm remove_device -F /dev/rdisk/cNtXdYs2
```

After running the command, warning messages might display. These messages can be ignored.

- 9** On any node that is connected to the storage array, run the `luxadm insert_device` command. Add the new disk drive when prompted.

```
# luxadm insert_device boxname,fslotnumber
```

or

```
# luxadm insert_device boxname,fslotnumber
```

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

- 10** On all other nodes that are attached to the storage array, probe all devices. Write the new disk drive to the `/dev/rdisk` directory.

The amount of time that the `devfsadm` command requires to complete depends on the number of devices that are connected to the node. Expect at least five minutes.

```
# devfsadm -C
```

- 11** If you are using If VERITAS Volume Manager, proceed to [Step 13](#).

- 12** If you are using Solstice DiskSuite/Solaris Volume Manager, on one node that is connected to the storage array, partition the new disk drive. Use the partitioning information you saved in [Step 6](#).

```
# fmthard -s filename /dev/rdisk/cNtXdYs2
```

Note – You can also use the `format` utility to partition the new disk drive.

- 13** From all nodes that are connected to the storage array, update the DID database and driver.

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

```
# cldevice repair
```

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

```
# scdidadm -R deviceID
```

Note – After running `sccidadm -R` on the first node, each subsequent node that you run the command on might display the warning, device id for the device matches the database. Ignore this warning.

14 On any node, confirm that the failed disk drive has been replaced. Compare the following physical DID to the physical DID in Step 5.

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

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

```
# cldevice list -v
```

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

```
# sccidadm -o diskid -l cNtXdY
```

15 Perform volume management administration to add the disk drive back to its diskset or disk group.

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

16 If you want this new disk drive to be a quorum device, add the quorum device.

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

Cabling Diagrams

This appendix provides cabling diagrams for installation and maintenance of Fibre Channel JBOD storage devices. It contains the following sections:

- [Sun StorEdge A5x00](#)

Sun StorEdge A5x00 Cabling Diagrams

Installing a Storage Array

[Figure A-1](#) shows a sample storage array configuration.

Note – You must use FC switches when you install storage arrays in a partner-group configuration. If you want to create a storage area network (SAN) by using two FC switches and Sun SAN software, see “SAN Solutions in a Sun Cluster Environment” in *Sun Cluster 3.1 - 3.2 Hardware Administration Manual for Solaris OS* for more information.

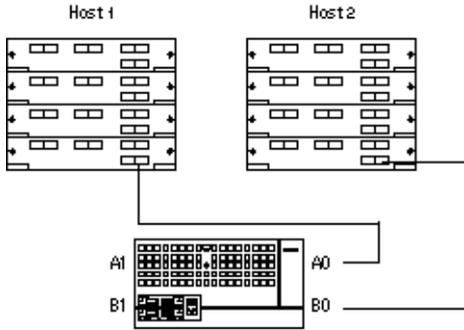


FIGURE A-1 Installing a Sun StorEdge A5x00 Storage Array in a Full-Loop Configuration

For more information about how to cable and configure storage arrays, see the *Sun StorEdge A5000 Installation and Service Manual*.

Adding a Storage System

Figure A-2 shows a sample storage array configuration.

Note – You must use FC switches when installing storage arrays in a partner-group configuration. If you want to create a storage area network (SAN) by using two FC switches and Sun SAN software, see “SAN Solutions in a Sun Cluster Environment” in *Sun Cluster 3.1 - 3.2 Hardware Administration Manual for Solaris OS* for more information.

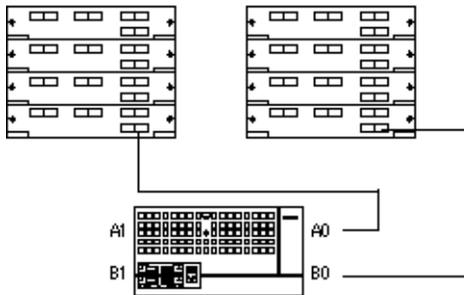


FIGURE A-2 Adding a Sun StorEdge A5x00 Storage Array in a Full-Loop Configuration

For more information about how to cable and configure storage arrays, see the *Sun StorEdge A5000 Configuration Guide* and the *Sun StorEdge A5000 Configuration Guide*.

Index

A

AC power filter, replacing, 16

adding

See also installing

disk drives, 22-23

first storage array, 12-14

subsequent storage arrays, 14-15

B

backplane, replacing, 16

C

cabling, storage arrays, 31-32

chassis, replacing, 16

D

disk drives

adding, 22-23

removing, 23-26

replacing, 26-29

door panel assembly, replacing, 16

F

fan tray, replacing, 16

fiber optic cable, replacing, 16

FPM, replacing, 16

FRUs, 16

G

GBIC, replacing, 16

H

help, 7-8

I

installing

See also adding

storage arrays, 12

interconnect assembly, replacing, 16

interface board, replacing, 16

M

maintaining storage arrays, 16-29

P

power supply, replacing, 16

R

removing

- disk drives, 23-26
- storage arrays, 19-22

replacing

- AC power filter, 16
- backplane, 16
- chassis, 16
- disk drives, 26-29
- door panel assembly, 16
- fan tray, 16
- fiber optic cable, 16
- FPM, 16
- GBIC, 16
- interconnect assembly, 16
- interface board, 16
- power supply, 16
- storage arrays, 17-19

S

storage arrays

- adding subsequent, 14-15
- adding the first, 12-14
- cabling, 31-32
- installing, 12
- maintaining, 16-29
- removing, 19-22
- replacing, 17-19

T

- technical support, 7-8
- training, 9