



Sun Cluster 3.x With Sun StorEdge 9900 Series Storage Device Manual

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

Part No: 817-0177-10
July 2003, Revision A

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

This product or document is protected by copyright and distributed under licenses restricting its use, copying, distribution, and decompilation. No part of this product or document may be reproduced in any form by any means without prior written authorization of Sun and its licensors, if any. Third-party software, including font technology, is copyrighted and licensed from Sun suppliers.

Parts of the product may be derived from Berkeley BSD systems, licensed from the University of California. UNIX is a registered trademark in the U.S. and other countries, exclusively licensed through X/Open Company, Ltd.

Sun, Sun Microsystems, the Sun logo, docs.sun.com, AnswerBook, AnswerBook2, and Solaris are trademarks, registered trademarks, or service marks of Sun Microsystems, Inc. in the U.S. and other countries. All SPARC trademarks are used under license and are trademarks or registered trademarks of SPARC International, Inc. in the U.S. and other countries. Products bearing SPARC trademarks are based upon an architecture developed by Sun Microsystems, Inc.

The OPEN LOOK and Sun™ Graphical User Interface was developed by Sun Microsystems, Inc. for its users and licensees. Sun acknowledges the pioneering efforts of Xerox in researching and developing the concept of visual or graphical user interfaces for the computer industry. Sun holds a non-exclusive license from Xerox to the Xerox Graphical User Interface, which license also covers Sun's licensees who implement OPEN LOOK GUIs and otherwise comply with Sun's written license agreements.

Federal Acquisitions: Commercial Software—Government Users Subject to Standard License Terms and Conditions.

DOCUMENTATION IS PROVIDED "AS IS" AND ALL EXPRESS OR IMPLIED CONDITIONS, REPRESENTATIONS AND WARRANTIES, INCLUDING ANY IMPLIED WARRANTY OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE OR NON-INFRINGEMENT, ARE DISCLAIMED, EXCEPT TO THE EXTENT THAT SUCH DISCLAIMERS ARE HELD TO BE LEGALLY INVALID.

Copyright 2003 Sun Microsystems, Inc. 4150 Network Circle, Santa Clara, CA 95054 U.S.A. Tous droits réservés.

Ce produit ou document est protégé par un copyright et distribué avec des licences qui en restreignent l'utilisation, la copie, la distribution, et la décompilation. Aucune partie de ce produit ou document ne peut être reproduite sous aucune forme, par quelque moyen que ce soit, sans l'autorisation préalable et écrite de Sun et de ses bailleurs de licence, s'il y en a. Le logiciel détenu par des tiers, et qui comprend la technologie relative aux polices de caractères, est protégé par un copyright et licencié par des fournisseurs de Sun.

Des parties de ce produit pourront être dérivées du système Berkeley BSD licenciés par l'Université de Californie. UNIX est une marque déposée aux Etats-Unis et dans d'autres pays et licenciée exclusivement par X/Open Company, Ltd.

Sun, Sun Microsystems, le logo Sun, docs.sun.com, AnswerBook, AnswerBook2, et Solaris sont des marques de fabrique ou des marques déposées, ou marques de service, de Sun Microsystems, Inc. aux Etats-Unis et dans d'autres pays. Toutes les marques SPARC sont utilisées sous licence et sont des marques de fabrique ou des marques déposées de SPARC International, Inc. aux Etats-Unis et dans d'autres pays. Les produits portant les marques SPARC sont basés sur une architecture développée par Sun Microsystems, Inc.

L'interface d'utilisation graphique OPEN LOOK et Sun™ a été développée par Sun Microsystems, Inc. pour ses utilisateurs et licenciés. Sun reconnaît les efforts de pionniers de Xerox pour la recherche et le développement du concept des interfaces d'utilisation visuelle ou graphique pour l'industrie de l'informatique. Sun détient une licence non exclusive de Xerox sur l'interface d'utilisation graphique Xerox, cette licence couvrant également les licenciés de Sun qui mettent en place l'interface d'utilisation graphique OPEN LOOK et qui en outre se conforment aux licences écrites de Sun.

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



030807 @ 5943



Contents

Preface	5
1 Installing and Maintaining a Sun StorEdge 9900 Series Storage Array	11
Installing Storage Arrays	11
How to Install a Storage Array in a Campus Cluster Configuration (StorEdge 9910 and StorEdge 9960 Only)	12
Configuring Storage Arrays	12
▼ How to Add a Logical Volume	13
▼ How to Remove a Logical Volume	14
Maintaining Storage Arrays	15
▼ How to Add a Storage Array	17
▼ How to Remove a Storage Array	22
▼ How to Replace a Node's Host Adapter	24
▼ How to Replace an FC Switch or Storage Array-to-Switch Component	25
▼ How to Replace a Node-to-Switch/Storage Component	26
2 Installing Multipathing Software in a Sun StorEdge 9900 Series Storage Array	29
Installing Multipathing Software	29
▼ How to Install Sun StorEdge 9900 Dynamic Link Manager Software	30
▼ How to Install Sun StorEdge Traffic Manager Software	31
Adding Multipathing Software	33
▼ How to Add Sun StorEdge 9900 Dynamic Link Manager Software	33
▼ How to Add Sun StorEdge Traffic Manager Software	35

Preface

The *Sun Cluster 3.x With Sun StorEdge 9900 Series Manual* provides procedures specific to Sun StorEdge™ 9900 Series storage devices that are placed in a Sun™ Cluster environment.

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

Who Should Use This Book

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

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

How This Book Is Organized

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

Section 1 discusses how to install Sun StorEdge 9900 storage systems

Section 2 discusses how to configure logical units on Sun StorEdge 9900 storage systems

Section 3 describes how to maintain Sun StorEdge 9900 storage systems in a running cluster.

Revision History

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

TABLE P-1 Sun Cluster 3.x With Sun StorEdge 9900 Series Manual

Revision Date	New Information
July 2003	Updated illustrations.
	Updated documentation with editorial changes.

Related Documentation

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

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

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

TABLE P-2 Hardware Documentation

Title	Part Number
<i>Sun StorEdge Traffic Manager Installation and Configuration Guide</i>	816-1420
Available on http://www.sun.com/products-n-solutions/hardware/docs	

TABLE P-3 Sun Cluster Documentation

Application	Title	Part Number
Concepts	<i>Sun Cluster 3.1 Concepts Guide</i>	816-3383
Hardware	<i>Sun Cluster 3.x Hardware Administration Manual</i>	817-0168
	Sun Cluster 3.x Hardware Administration Collection at http://docs.sun.com/db/coll/1024.1/	
Software Installation	<i>Sun Cluster 3.1 Software Installation Guide</i>	816-3388
Data Services	<i>Sun Cluster 3.1 Data Service Planning and Administration Guide</i>	817-1526
	Sun Cluster 3.1 Data Service Collection at http://docs.sun.com/db/coll/573.10/	
API Development	<i>Sun Cluster 3.1 Data Services Developer's Guide</i>	816-3385
Administration	<i>Sun Cluster 3.1 System Administration Guide</i>	816-3384
Error Messages	<i>Sun Cluster 3.1 Error Messages Guide</i>	816-3382
Man Pages	<i>Sun Cluster 3.1 Man Page Reference Manual</i>	816-5251
Release Notes	<i>Sun Cluster 3.1 Release Notes</i>	816-5317
	<i>Sun Cluster 3.1 Release Notes Supplement</i>	816-3381

Using UNIX Commands

This document contains information on commands used to install, configure, or upgrade a Sun Cluster configuration. This document might not contain complete information on basic UNIX[®] commands and procedures such as shutting down the system, booting the system, and configuring devices.

See one or more of the following sources for this information.

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

Getting Help

If you have problems installing or using Sun Cluster, contact your service provider and provide the following information.

- Your name and email address (if available)
- Your company name, address, and phone number
- The model number and serial number of your systems
- The release number of the operating environment (for example, Solaris 8)
- The release number of Sun Cluster (for example, Sun Cluster 3.0)

Use the following commands to gather information on your system for your service provider.

Command	Function
<code>prtconf -v</code>	Displays the size of the system memory and reports information about peripheral devices
<code>psrinfo -v</code>	Displays information about processors
<code>showrev -p</code>	Reports which patches are installed
<code>prtdiag -v</code>	Displays system diagnostic information
<code>/usr/cluster/bin/scinstall -pv</code>	Displays Sun Cluster release and package version information

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

Accessing Sun Documentation Online

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

Typographic Conventions

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

TABLE P-4 Typographic Conventions

Typeface or Symbol	Meaning	Example
AaBbCc123	The names of commands, files, and directories; on-screen computer output	Edit your <code>.login</code> file. Use <code>ls -a</code> to list all files. <code>machine_name%</code> you have mail.
AaBbCc123	What you type, contrasted with on-screen computer output	<code>machine_name%</code> su Password:
<i>AaBbCc123</i>	Command-line placeholder: replace with a real name or value	To delete a file, type rm <i>filename</i> .
<i>AaBbCc123</i>	Book titles, new words, or terms, or words to be emphasized.	Read Chapter 6 in <i>User's Guide</i> . These are called <i>class</i> options. You must be <i>root</i> to do this.

Shell Prompts in Command Examples

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

TABLE P-5 Shell Prompts

Shell	Prompt
C shell prompt	<code>machine_name%</code>
C shell superuser prompt	<code>machine_name#</code>
Bourne shell and Korn shell prompt	<code>\$</code>
Bourne shell and Korn shell superuser prompt	<code>#</code>

Installing and Maintaining a Sun StorEdge 9900 Series Storage Array

This chapter contains a limited set of procedures about how to install, configure, and maintain Sun StorEdge 9900 Series storage arrays in a cluster environment. The StorEdge 9900 Series storage arrays include the StorEdge 9910, StorEdge 9960, StorEdge 9970 and StorEdge 9980. Procedures in this chapter can be performed on all StorEdge 9900 Series storage arrays unless noted otherwise. Contact your service provider to perform tasks not documented in this chapter.

This chapter contains the following procedures.

- “How to Add a Logical Volume” on page 13
- “How to Remove a Logical Volume” on page 14
- “How to Add a Storage Array” on page 17
- “How to Remove a Storage Array” on page 22
- “How to Replace a Node’s Host Adapter” on page 24
- “How to Replace an FC Switch or Storage Array-to-Switch Component” on page 25
- “How to Replace a Node-to-Switch/Storage Component” on page 26

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

Installing Storage Arrays

This section references information about how to install a cluster in a campus cluster configuration. The installation of the storage array should only be performed by your certified service provider.

How to Install a Storage Array in a Campus Cluster Configuration (StorEdge 9910 and StorEdge 9960 Only)

You can configure storage arrays in a two-room or three-room campus cluster.

Note – This functionality is not supported for the StorEdge 9970 and StorEdge 9980.

For supported campus cluster hardware configurations, contact your Sun sales representative.

Configuring Storage Arrays

This section contains the procedures about how to configure a storage array in a Sun Cluster environment. The following table lists these procedures. For configuration tasks that are not cluster-specific, see the documentation that shipped with your storage array.

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

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

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

TABLE 1-1 Task Map: Configuring a Storage Array

Task	Information
Add a logical volume.	See “How to Add a Logical Volume” on page 13.

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

Task	Information
Remove a logical volume.	See “How to Remove a Logical Volume” on page 14.

▼ How to Add a Logical Volume

Use this procedure to add a logical volume to a cluster. This procedure assumes that your service provider created your logical volume. This procedure also assumes that all nodes are booted and are attached to the storage array.

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

```
# devfsadm
```

2. **On each node connected to the storage array, verify that the same set of LUNs is visible to the expected controllers.**

```
# format
```

See the `format` command man page for more information about how to use the command.

3. **Determine if you are running VERITAS Volume Manager.**

- If not, proceed to Step 4
- If you are running VERITAS Volume Manager, update the list of devices on all nodes that are attached to the logical volume that you created in Step 2.

See your VERITAS Volume Manager documentation for information about how to use the `vxctl enable` command. Use this command to update new devices (volumes) in your VERITAS Volume Manager list of devices.

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

```
# scgdevs
```

If a volume management daemon such as `vold` 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 expected behavior.

Where to Go From Here

To create a new resource or reconfigure a running resource to use the new logical volume, see your Sun Cluster data services collection.

▼ How to Remove a Logical Volume

Use this procedure to remove a logical volume. This procedure assumes all nodes are booted and are connected to the storage array. This storage array hosts the logical volume that you are removing.

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

If you need to remove a storage array from more than two nodes, repeat Step 9 through Step 11 for each additional node. Each node connects to the logical volume.



Caution – This procedure destroys all data on the logical volume that you are removing.

1. **If necessary, back up all data. Migrate all resource groups and disk device groups to another node.**
2. **Is the logical volume that you plan to remove configured as a quorum device?**
 - # `scstat -q`
 - If no, proceed to Step 3.
 - If yes, 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 your Sun Cluster system administration documentation.
3. **Run the appropriate Solstice DiskSuite/Solaris Volume Manager commands or VERITAS Volume Manager commands to remove the reference to the logical volume from any diskset or disk group.**

For more information, see your Solstice DiskSuite/Solaris Volume Manager or VERITAS Volume Manager documentation.
4. **If the cluster is running VERITAS Volume Manager, update the list of devices on all nodes. These nodes are attached to the logical volume that you are removing.**

See your VERITAS Volume Manager documentation for information about how to use the `vxdisk rm` command to remove devices (volumes) in your VERITAS Volume Manager device list.
5. **Remove the logical volume.**

Contact your service provider to remove the logical volume.
6. **Determine the resource groups and device groups that are running on Node A and Node B.**

Record this information because you use this information in Step 11 of this procedure to return resource groups and device groups to these nodes.

```
# scstat
```

7. **Shut down and reboot Node A by using the `shutdown` command with the `-i6` option.**

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

8. **On Node A, update the `/devices` and `/dev` entries.**

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

9. **Shut down and reboot Node B by using the `shutdown` command with the `-i6` option.**

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

10. **On Node B, update the `/devices` and `/dev` entries.**

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

11. **Return the resource groups and device groups that you identified in Step 6 to Node B.**

```
# 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.

12. **Repeat Step 9 through Step 11 for each additional node that connects to the logical volume.**

Where to Go From Here

To create a logical volume, see “How to Add a Logical Volume” on page 13.

Maintaining Storage Arrays

This section contains a limited set of procedures about how to maintain a storage array. Contact your service provider to add, remove, or replace any storage array components.

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

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

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

TABLE 1-2 Task Map: Maintaining a Storage Array

Task	Information
Add a storage array.	"How to Add a Storage Array" on page 17
Remove a storage array.	"How to Remove a Storage Array" on page 22
Add a node to the storage array.	Sun Cluster system administration documentation
Remove a node from the storage array.	Sun Cluster system administration documentation
Replace a node's host adapter.	"How to Replace a Node's Host Adapter" on page 24
Replace an FC switch or storage array-to-switch component.	"How to Replace an FC Switch or Storage Array-to-Switch Component" on page 25
Replace a node-to-switch/storage component.	"How to Replace a Node-to-Switch/Storage Component" on page 26

▼ How to Add a Storage Array

Use this procedure to add a new storage array to a running cluster.

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

If you need to add a storage array to more than two nodes, repeat Step 24 through Step 40 for each additional node that connects to the storage array.

1. Power on the storage array.

Note – The storage array requires approximately 10 minutes to boot.

Contact your service provider to power on the storage array.

2. If you plan to use multipathing software, verify that the storage array is configured for multipathing.

Contact your service provider to verify that the storage array is configured for multipathing.

3. Configure the new storage array.

Contact your service provider to create the desired logical volumes.

4. Do you need to install a host adapter in Node A?

Note – If you use multipathing software, each node requires two paths to the same set of LUNs.

- If no, skip to Step 12.
- If yes, proceed to Step 5.

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

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

- If no, proceed to Step 6.
- If yes, contact your service provider to install the support packages and configure the drivers before you proceed to Step 6.

6. Is your node enabled with the Solaris dynamic reconfiguration (DR) feature?

- If yes, install the host adapter.

For the procedure about how to install a host adapter, see the documentation that shipped with your host adapter or updated information on the manufacturer's web site.

- If no, shut down this node to install the host adapter(s). Proceed to Step 7.

7. Determine the resource groups and device groups that are running on Node A and Node B.

Record this information because you use this information in Step 40 of this procedure to return resource groups and device groups to these nodes.

```
# scstat
```

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

8. Shut down and power off Node A.

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

9. Install the host adapter in Node A.

For the procedure about how to install a host adapter, see the documentation that shipped with your host adapter or updated information on the manufacturer's web site.

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

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

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

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

12. Attach the storage array to Node A.

Contact your service provider to install a fiber-optic cable between the storage array and your node.

13. Configure the host adapter and the storage array.

Contact your service provider to configure the adapter and storage array.

14. If you plan to install multipathing software, which multipathing solution do you plan to install?

- If Sun StorEdge Traffic Manager software, proceed to Step 15.
- If Sun StorEdge 9900 Dynamic Link Manager (Sun SDLM) software, skip to Step 17.

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

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

required patches for Sun StorEdge Traffic Manager software, see the *Sun StorEdge Traffic Manager Software Release Notes* on <http://www.sun.com/storage/san>.

16. Edit the `/kernel/drv/scsi_vhci.conf` file to activate the Sun StorEdge Traffic Manager software functionality. Set the `mpxio-disable` parameter to `no`. Skip to Step 18.

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

17. Install Sun StorEdge 9900 Dynamic Link Manager (Sun SDLM) software and any required patches for Sun SDLM software support on Node A.

For the procedure about how to install the Sun SDLM software, see the documentation that shipped with your storage array.

18. Shut down Node A.

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

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

20. On Node A, configure all controllers that are affected by the new physical path.

```
# cfgadm -c configure cN
```

21. On each node connected to the storage array, verify that the same set of LUNs is visible to the expected controllers.

```
# format
```

See the `format` command man page for more information about how to use the command.

22. On Node A, update the paths to the DID instances.

```
# sddidadm -C  
# sddidadm -r
```

23. (Optional) On Node A, verify that the DIDs are assigned to the new storage array.

```
# sddidadm -l
```

24. Do you need to install a host adapter in Node B?

Note – If you use multipathing software, each node requires two paths to the same set of LUNs.

- If no, skip to Step 31.
- If yes, proceed to Step 25.

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

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

- If no, skip to Step 26.
- If yes, contact your service provider to install the support packages and configure the drivers before you proceed to Step 26.

26. Is your node enabled with the Solaris dynamic reconfiguration (DR) feature?

- If yes, install the host adapter.
For the procedure about how to install a host adapter, see the documentation that shipped with your host adapter or updated information on the manufacturer's web site.
- If no, proceed to Step 27.

27. Shut down and power off Node B.

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

28. Install the host adapter in Node B.

For the procedure about how to install a host adapter, see the documentation that shipped with your host adapter or updated information on the manufacturer's web site.

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

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

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

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

31. Attach the storage array to Node B.

Contact your service provider to install a fiber-optic cable between the storage array and your node.

32. If you plan to install multipathing software, which multipathing solution do you plan to install?

- If Sun StorEdge Traffic Manager software, proceed to Step 33.
- If Sun StorEdge 9900 Dynamic Link Manager (Sun SDLM) software, skip to Step 35.

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

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

- 34. Edit the `/kernel/drv/scsi_vhci.conf` file to activate the Sun StorEdge Traffic Manager software functionality. Set the `mpxio-disable` parameter to `no`. Skip to Step 36.**

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

- 35. Install Sun StorEdge 9900 Dynamic Link Manager (Sun SDLM) software and any required patches for Sun SDLM software support on Node B.**

For the procedure about how to install the Sun SDLM software, see the documentation that shipped with your storage array.

- 36. Shut down Node B.**

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

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

- 38. On Node B, update the paths to the DID instances.**

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

- 39. (Optional) On Node B, verify that the device IDs (DIDs) are assigned to the new storage array.**

```
# scdidadm -l
```

- 40. Return the resource groups and device groups that you identified in Step 7 to Node B.**

```
# 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.

- 41. Repeat Step 24 through Step 40 for each additional node that connects to the storage array.**

- 42. Perform volume management administration to incorporate the new logical volumes into the cluster.**

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

▼ How to Remove a Storage Array

Use this procedure to permanently remove a storage array. This procedure provides the flexibility to remove the host adapters from the nodes that are attached to the storage array that you are removing.

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

If you need to remove a storage array from more than two nodes, repeat Step 15 through Step 22 for each additional node that connects to the storage array.



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

1. **If necessary, back up all data and migrate all resource groups and disk device groups to another node.**
2. **Determine if the storage array that you plan to remove is configured as a quorum device.**

`scstat -q`

- If the storage array is not a quorum device, proceed to Step 3.
- If the storage array is configured as a quorum device, choose and configure another device to be the new quorum device. Then remove the old quorum device.

To add and remove a quorum device in your configuration, see your Sun Cluster system administration documentation.

3. **If necessary, detach the submirrors from the storage array that you are removing in order to stop all I/O activity to the storage array.**

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

4. **Run the appropriate Solstice DiskSuite/Solaris Volume Manager or VERITAS Volume Manager commands to remove the references to the logical volume(s) from any diskset or disk group.**

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

5. **Are your nodes enabled with the Solaris dynamic reconfiguration (DR) feature?**

- If yes, disconnect the fiber-optic cables and, if desired, remove the host adapters from both nodes. Then perform Step 22 on each node that was connected to the storage array.

For the procedure about how to remove a host adapter, see the documentation that shipped with your host adapter or updated information on the manufacturer's web site.

- If no, proceed to Step 6.

6. Determine the resource groups and device groups that are running on Node A and Node B.

Record this information because that you use this information in Step 21 of this procedure to return resource groups and device groups to these nodes.

```
# scstat
```

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

7. Do you want to remove any multipathing software?

- If no, proceed to Step 8.
- If yes, remove the multipathing software packages.

8. Stop the Sun Cluster software on Node A, and then shut down Node A.

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

9. Disconnect the fiber-optic cable between the storage array and Node A.

10. Do you want to remove the host adapter from Node A?

- If no, skip to Step 13.
- If yes, power off Node A.

11. Remove the host adapter from Node A.

For the procedure about how to remove host adapters, see the documentation that shipped with your host adapter or updated information on the manufacturer's web site.

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

For more information, see the documentation that shipped with your server.

13. Boot Node A into cluster mode.

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

14. On Node A, update the device namespace.

```
# devfsadm -C
```

15. Stop the Sun Cluster software on Node B, and then shut down Node B.

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

16. Disconnect the fiber-optic cable between the storage array and Node B.

17. Do you want to remove the host adapter from Node B?

- If no, skip to Step 20.
- If yes, power off Node B.

18. Remove the host adapter from Node B.

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

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

For more information, see the documentation that shipped with your server.

20. Boot Node B into cluster mode.

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

21. Return the resource groups and device groups that you identified in Step 6 to Node B.

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

22. On Node B, update the device namespace.

```
# devfsadm -C
```

23. Repeat Step 15 through Step 22 for each additional node that connects to the storage array.

24. From one node, remove DID references to the storage array that was removed.

```
# scdidadm -C
```

▼ How to Replace a Node's Host Adapter

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

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

Record this information because you use this information in Step 11 of this procedure to return resource groups and device groups to these nodes.

```
# scstat
```

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

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

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

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

3. Shut down Node N.

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

4. Power off Node N.

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

5. Replace the failed host adapter.

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

6. Power on Node N.

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

7. Boot Node N into single-user mode.

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

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

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

9. Boot Node N into cluster mode.

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

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

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

11. Return the resource groups and device groups that you identified in Step 1 to Node N.

```
# 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.

▼ How to Replace an FC Switch or Storage Array-to-Switch Component

Use this procedure to replace an FC switch, or the following storage array-to-switch components in a running cluster.

- Fiber-optic cable that connects an FC switch to a storage array

- GBIC on an FC switch, connecting to a storage array
- FC switch
- **Replace the component by using the following references.**
 - For the procedure about how to replace a fiber-optic cable between a storage array and an FC switch, see the documentation that shipped with your switch hardware.
 - For the procedure about how to replace a GBIC on an FC switch, see the documentation that shipped with your switch hardware.
 - For the procedure about how to replace an FC switch, see the documentation that shipped with your switch hardware.

Note – If you are replacing an FC switch and you intend to save the switch configuration for restoration to the replacement switch, do not connect the cables to the replacement switch until *after* you recall the Fabric configuration to the replacement switch. For more information about how to save and recall switch configurations see the documentation that shipped with your switch hardware.

Note – Before you replace an FC switch, be sure that the `probe_timeout` parameter of your data service software is set to more than 90 seconds. Increasing the value of the `probe_timeout` parameter to more than 90 seconds avoids unnecessary resource group restarts when one of the FC switches is powered off.

▼ How to Replace a Node-to-Switch/Storage Component

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

- Node-to-switch fiber-optic cable
- Node-to-storage fiber-optic cable, assuming you have direct attached storage
- GBIC (gigabit interface converter) on an FC (Fibre Channel) switch, connecting to a node

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

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

```
# scstat
```

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

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

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

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

3. Replace the node-to-switch/storage component.

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

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

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

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


Installing Multipathing Software in a Sun StorEdge 9900 Series Storage Array

This chapter contains procedures on how to install and add multipathing software. Multipathing software enables you to define and control redundant physical paths to I/O devices such as storage arrays and networks. If the active path to a device becomes unavailable, the multipathing software can automatically switch to an alternate path to maintain availability. This capability is known as automatic failover. To maximize multipathing capabilities, your servers must be configured with redundant hardware. Redundant hardware includes redundant network interfaces or two host bus adapters that are connected to the same dual-ported storage array.

This chapter contains the following procedures.

- “How to Install Sun StorEdge 9900 Dynamic Link Manager Software” on page 30
- “How to Install Sun StorEdge Traffic Manager Software” on page 31
- “How to Add Sun StorEdge 9900 Dynamic Link Manager Software” on page 33
- “How to Add Sun StorEdge Traffic Manager Software” on page 35

Installing Multipathing Software

This section contains procedures on how to install multipathing software in a cluster that *is not* yet installed with Sun Cluster software. Choose one of the following multipathing solutions.

- Sun StorEdge 9900 Dynamic Link Manager software
- Sun StorEdge Traffic Manager software

▼ How to Install Sun StorEdge 9900 Dynamic Link Manager Software

Use this procedure to install Sun StorEdge 9900 Dynamic Link Manager (Sun SDLM) software in a cluster that is *not* yet installed with Sun Cluster software. This procedure assumes that you installed Solaris operating environment. If your cluster is already running Sun Cluster software, see “Adding Multipathing Software” on page 33 to add multipathing software to a running cluster.

Perform this procedure on one node at a time. This procedure defines Node N as the node on which you are installing the multipathing software.

1. Install the new storage array.

For instructions on how to install the storage array, see the documentation that shipped with your storage array.

2. Configure the storage array such that each node has two paths to the same set of LUNs. For an example configuration, see Figure 2–1.

Contact your service provider to create the desired logical volumes. These new logical volumes can be configured by using LUN Expansion (LUSE) or Virtual LVI/LUN.




FIGURE 2–1 Switched Configuration

3. On each node that is connected to the storage array, use the `format (1M)` command. Use this command to verify that the same set of LUNs is visible to the expected controllers.

```
# format
```

See the `format` command man page for more information about by using the command.

4. Install and configure Sun SDLM software.

For instructions on how to install and configure the Sun SDLM software, see the documentation that shipped with your storage array.

5. Did you perform a reconfiguration reboot when you configured Sun SDLM software?

- If yes, proceed to Step 6.
- If no, shut down Node N, and perform a reconfiguration reboot to create the new Solaris device files and links on Node N.

6. On each node connected to the storage array, use the `format (1M)` command. Use this command to verify that you see half the number of disks you saw in Step 3.

`format`

See the `format` command man page for more information about by using the command.

7. On all remaining nodes and one node at a time, repeat Step 1 through Step 6.

Where to Go From Here

After you install multipathing software, you might want to install Sun Cluster software. For the procedures on how to install Sun Cluster software, see Sun Cluster software installation documentation.

▼ How to Install Sun StorEdge Traffic Manager Software

Use this procedure to install Sun StorEdge Traffic Manager software in a cluster that is *not* yet installed with Sun Cluster software. This procedure assumes that you installed Solaris operating environment. If your cluster is already running Sun Cluster software, see “Adding Multipathing Software” on page 33 to add multipathing software to a running cluster.

Perform this procedure on one node at a time. This procedure defines Node N as the node on which you are installing the multipathing software.

1. Install the new storage array.

For instructions on how to install the storage array, see the documentation that shipped with your storage array.

2. Configure the storage array so that each node has two paths to the same set of LUNs. For an example configuration, see Figure 2–2.

Contact your service provider to create the desired logical volumes. These new logical volumes can be configured by using LUN Expansion (LUSE) or Virtual LVI/LUN.




FIGURE 2-2 Switched Configuration

3. On each node connected to the storage array, use the **format (1M)** command. Use this command to verify that the same set of LUNs is visible to the expected controllers.

```
# format
```

See the `format` command man page for more information about by using the command.

4. Install any required patches or software for Sun StorEdge Traffic Manager software support on Node N.

For instructions on installing the software, see the *Sun StorEdge Traffic Manager Installation and Configuration Guide* at <http://www.sun.com/products-n-solutions/hardware/docs/>.

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

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

6. Boot Node N into cluster mode.

For more information on booting nodes, see Sun Cluster software installation documentation.

7. Did you perform a reconfiguration reboot when you configured Sun StorEdge Traffic Manager software?

- If yes, proceed to Step 8.
- If no, shut down Node N, and perform a reconfiguration reboot to create the new Solaris device files and links on Node N.

8. On each node connected to the storage array, use the `format (1M)` command. Use this command to verify that you see half the number of disks you saw in Step 3.

```
# format
```

See the `format` command man page for more information about by using the command.

9. On all remaining nodes and one node at a time, repeat Step 1 through Step 8.

Where to Go From Here

After you install multipathing software, you might want to install Sun Cluster software. For the procedure on how to install Sun Cluster software, see Sun Cluster software installation documentation .

Adding Multipathing Software

This section contains a procedure on how to add multipathing software in a running cluster. Choose one of the following multipathing solutions.

- Sun StorEdge 9900 Dynamic Link Manager software
- Sun StorEdge Traffic Manager software

▼ How to Add Sun StorEdge 9900 Dynamic Link Manager Software

Use this procedure to add Sun StorEdge 9900 Dynamic Link Manager (Sun SDLM) software in a running cluster. This procedure assumes that you installed and configured your storage array. If you have not installed your storage array, see “How to Install Sun StorEdge 9900 Dynamic Link Manager Software” on page 30 to add multipathing software and the storage array at the same time.

Do not use this procedure to convert from host-based mirroring to a multipathing solution. For the procedure on how to convert from host-based mirroring to a multipathing solution, contact your Sun service provider.

Perform this procedure on one node at a time. This procedure defines Node N as the node on which you are installing the multipathing software.

1. **Determine the resource groups and device groups that are running on Node N.**

Record this information because you use this information in Step 11 of this procedure to return resource groups and device groups to these nodes.

```
# scstat
```

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

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

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

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

3. Do you need to install additional physical paths between Node N and the storage?

Note – If you use multipathing software, each node requires two paths to the same set of LUNs. For an example configuration, see Figure 2–1.

- If no, skip to Step 8.
- If yes, proceed to Step 4.

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

4. Shut down and power off Node N.

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

5. Install the host adapters and the cables between Node N and the storage.

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

6. Power on and boot Node N into non-cluster mode.

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

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

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

8. Install and configure Sun SDLM software on Node N, and apply required patches for Sun SDLM software support on Node N.

For instructions on how to install and configure the Sun SDLM software, see the documentation that shipped with your storage array.

9. Did you perform a reconfiguration reboot when you configured Sun SDLM software?

- If yes, proceed to Step 10.

- If no, perform a reconfiguration reboot to create the new Solaris device files and links on Node N.

10. On Node N, update the paths to the device ID (DID) instances.

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

11. Return the resource groups and device groups that you identified in Step 1 to Node N.

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

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

12. On all remaining nodes and one node at a time, repeat Step 1 through Step 11.

▼ How to Add Sun StorEdge Traffic Manager Software

Use this procedure to add Sun StorEdge Traffic Manager software in a running cluster. This procedure assumes that you installed and configured your storage array. If you have not installed your storage array, see “How to Install Sun StorEdge Traffic Manager Software” on page 31 to add multipathing software and the storage array at the same time.

Do not use this procedure to convert from host-based mirroring to a multipathing solution. For the procedure on how to convert from host-based mirroring to a multipathing solution, contact your Sun service provider.

Perform this procedure on one node at a time. This procedure defines Node N as the node on which you are installing the multipathing software.

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

Record this information because you use this information in Step 14 of this procedure to return resource groups and device groups to these nodes.

```
# scstat
```

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

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

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

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

3. Do you need to install additional physical paths between Node N and the storage?

Note – If you use multipathing software, each node requires two paths to the same set of LUNs. For an example configuration, see Figure 2–2.

- If no, skip to Step 8.
- If yes, proceed to Step 4.

4. Shut down and power off Node N.

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

5. Install the host adapters and the cables between Node N and the storage.

Note – If you use multipathing software, each node requires two paths to the same set of LUNs.

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

6. Power on and boot Node N into non-cluster mode.

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

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

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

8. Install any required patches or software for Sun StorEdge Traffic Manager software support on Node N.

For instructions on installing the software, see the *Sun StorEdge Traffic Manager Installation and Configuration Guide* at <http://www.sun.com/products-n-solutions/hardware/docs/>.

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

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

10. Boot Node N into cluster mode.

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

11. On each node connected to the storage array, use the `format (1M)` command. Use this command to verify that the same set of LUNs is visible to the expected controllers.

```
# format
```

See the `format` command man page for more information about by using the command.

12. On Node N, configure all controllers that are affected by the new physical path.

```
# cfgadm -c configure cN
```

13. On Node N, update the paths to the device ID (DID) instances.

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

14. Return the resource groups and device groups that you identified in Step 1 to Node N.

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

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

15. On each node connected to the storage array, use the `format (1M)` command. Use this command to verify that you see half the number of disks you saw in Step 11.

```
# format
```

See the `format` command man page for more information about by using the command.

16. On all remaining nodes and one node at a time, repeat Step 1 through Step 14.

Index

A

adding
 nodes, 16
 storage arrays, 17
 Sun SDLM, 33, 35
arrays, *See* storage arrays

C

creating, logical volumes, 13

D

deleting
 See also removing

H

help, 8
host adapters
 replacing, 16, 24

I

installing
 storage array
 campus cluster configuration, 12
 noncampus cluster configuration, 11
 Sun SDLM, 30

installing (Continued)

 Sun StorEdge Traffic Manager, 31

L

logical unit numbers, *See* logical volumes
logical units, *See* logical volumes
logical volumes
 creating, 13
 removing, 14
LUNs, *See* logical volumes

M

multipathing software
 adding, 33
 installing, 29

N

node-to-switch/storage components
 replacing, 16, 26
nodes
 adding, 16
 removing, 16

R

removing

- logical volumes, 14
- nodes, 16
- storage arrays, 22

replacing

- host adapters, 16, 24
- node-to-switch/storage components, 16, 26
- storage array-to-switch components, 16, 25

S

storage array

See storage array

installing

- campus cluster configuration, 12
- noncampus cluster configuration, 11

storage array-to-switch components

- replacing, 16, 25

storage arrays

- adding, 17
- removing, 22

Sun SDLM

- adding, 33, 35
- installing, 30

Sun StorEdge Traffic Manager, installing, 31

T

technical support, 8