



Sun Cluster 3.x With Sun StorEdge D2 Array Manual

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

Part No: 817-0172-10
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 D2 Array 11

Installing Storage Arrays 11

▼ How to Install a Storage Array With a Split-Bus or Single-Bus Configuration 13

Maintaining Storage Arrays 16

▼ How to Add a Storage Array 16

▼ How to Remove a Storage Array 22

▼ How to Replace the Chassis 23

▼ How to Add a Disk Drive 24

▼ How to Remove a Disk Drive 25

▼ How to Replace a Disk Drive 26

▼ How to Replace an Environmental Service Module (ESM) 29

▼ How to Upgrade Disk Drive Firmware 30

▼ How to Upgrade Host Adapter Firmware 30

▼ How to Replace a Host Adapter 31

Index 33

Preface

The *Sun Cluster 3.x With Sun StorEdge D2 Array Manual* provides procedures specific to Sun StorEdge™ D2 arrays 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 two major sections.

Section 1 discusses how to install Sun StorEdge D2 storage arrays.

Section 2 describes how to maintain Sun StorEdge D2 storage arrays in a running cluster.

Revision History

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

TABLE P-1 Sun Cluster 3.x With Sun StorEdge D2 Array Manual

Revision Date	New Information
July 2003	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 Storage D2 Array Installation, Operation, and Service Manual</i>	816-2578
<i>Sun StorEdge PCI Dual Ultra 3 SCSI Host Adapter Release Notes</i>	816-2157

TABLE P-3 Sun Cluster Documentation

Application	Title	Part Number
Concepts	<i>Sun Cluster 3.1 Concepts Guide</i>	816-3383

TABLE P-3 Sun Cluster Documentation (Continued)

Application	Title	Part Number
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% su</code> 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 D2 Array

This chapter provides the procedures to install and maintain a Sun StorEdge™ D2 array in a Sun™ Cluster environment.

This chapter contains the following procedures.

- “How to Install a Storage Array With a Split-Bus or Single-Bus Configuration” on page 13
- “How to Add a Storage Array” on page 16
- “How to Replace the Chassis” on page 23
- “How to Remove a Storage Array” on page 22
- “How to Add a Disk Drive” on page 24
- “How to Replace a Disk Drive” on page 26
- “How to Remove a Disk Drive” on page 25
- “How to Replace an Environmental Service Module (ESM)” on page 29
- “How to Upgrade Disk Drive Firmware” on page 30
- “How to Upgrade Host Adapter Firmware” on page 30
- “How to Replace a Host Adapter” on page 31

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

Installing Storage Arrays

This section contains the procedure for an initial installation of the storage array in a new cluster configuration.

You can install a single storage array by using a split-bus configuration. You can install two storage arrays by using a single-bus configuration.




FIGURE 1-1 Installing a Split-Bus Configuration: An Example

To use a single storage array in a Sun Cluster environment you must use the split-bus configuration.

The split-bus configuration requires that the storage array contain two Environmental Service Modules (ESM) that are each connected to two nodes. Each ESM controls half the disks in the storage array and the data is mirrored using a volume manager (see Figure 1-1).




FIGURE 1-2 Installing a Single-Bus Configuration: An Example

The single-bus configuration requires two storage arrays that each contain a single Environmental Service Module (ESM). The ESM controls all disks on the storage array that it is installed in and the volume manager implements mirroring between the two storage arrays. Each ESM must be connected to two nodes (see Figure 1-2).

▼ How to Install a Storage Array With a Split-Bus or Single-Bus Configuration

Perform this procedure in conjunction with the procedures in the Sun Cluster software installation documentation and your server hardware manual.

Multihost storage in clusters uses the multi-initiator capability of the SCSI specification. For conceptual information on multi-initiator capability, see your Sun Cluster concepts documentation document.

- 1. Verify that the storage array is configured for the split-bus or single-bus configuration.**

See the *Sun Storage D2 Array Installation, Operation, and Service Manual* for procedures about how to configure the ESMs.

- 2. Install the host adapters in the node that you are connecting to the storage array.**

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

- 3. Connect the cables to the storage arrays, as shown in Figure 1–1 and Figure 1–2.**

Make sure that the *entire* bus length that is connected to the storage array is less than 25m. The bus length includes the cables from the ESM to the host adapter. See the documentation that shipped with your hardware for internal bus length.

- 4. Connect the AC power cords for each storage array to a different power source.**

- 5. Ensure that each device in the SCSI chain has a unique SCSI address.**

To avoid SCSI-chain conflicts, this procedure instructs you to reserve SCSI address 7 for one host adapter in the SCSI chain and change the other host adapter's global `scsi-initiator-id` to an available SCSI address. Then the procedure instructs you to change the `scsi-initiator-id` for local devices back to 7.

Note – A slot in the storage array might not be in use. However, do not set the `scsi-initiator-id` to a SCSI address for that disk slot. This precaution minimizes future complications if you install additional disk drives.

- a. If necessary, power on the storage devices.**

For the procedure on powering on a storage device, see the service manual that shipped with your storage device.

- b. If necessary, power on a node, but do not allow it to boot. If necessary, halt the system to continue with OBP Monitor tasks.**

- c. Set the `scsi-initiator-id` for one node to 6.**

```
{1} ok setenv scsi-initiator-id 6
scsi-initiator-id = 6
```

d. Find the paths to the host adapters that connect to the local disk drives.

```
{0} ok show-disks
```

Use this information to change the SCSI addresses in the `nvrामrc` script. Do not include the `/sd` directories in the device paths.

e. Edit the `nvrामrc` script to set the `scsi-initiator-id` for the local devices on the first node to 7.

For a full list of commands, see the *OpenBoot 2.x Command Reference Manual*.



Caution – Insert exactly one space after the first double quote and before `scsi-initiator-id`.

```
{0} ok nvedit
0: probe-all
1: cd /pci@1f,4000/scsi@2
2: 7 " scsi-initiator-id" integer-property
3: device-end
4: cd /pci@1f,4000/scsi@3
5: 7 " scsi-initiator-id" integer-property
6: device-end
7: install-console
8: banner[Control C]
{0} ok
```

f. Store the changes.

The changes you make through the `nvedit` command are recorded on a temporary copy of the `nvrामrc` script. You can continue to edit this copy without risk. After you complete your edits, save the changes. If you are not sure about the changes, discard them.

- To store the changes, type:

```
{0} ok nvstore
{1} ok
```

- To discard the changes, type:

```
{0} ok nvquit
{1} ok
```

g. Verify the contents of the `nvrामrc` script that you created, as shown in the following example.

If the contents of the `nvrामrc` script are incorrect, use the `nvedit` command to make corrections.

```
{0} ok printenv nvrामrc
nvrामrc = probe-all
```

```

cd /pci@1f,4000/scsi@2
7 " scsi-initiator-id" integer-property
device-end
cd /pci@1f,4000/scsi@3
7 " scsi-initiator-id" integer-property
device-end
install-console
banner
{1} ok

```

- h. Instruct the OpenBoot PROM (OBP) Monitor to use the `nvrsrc` script, as shown in the following example.

```

{0} ok setenv use-nvrsrc? true
use-nvrsrc? = true
{1} ok

```

6. Verify that the `scsi-initiator-id` is set correctly on the second node.

- a. If necessary, power on the second node, but do not allow it to boot. If necessary, halt the system to continue with OBP Monitor tasks.
- b. Verify that the `scsi-initiator-id` for each host adapter on the second node is set to 7

Use the `show-disks` command to find the paths to the host adapters that are connected to these enclosures. Select each host adapter's device tree node, and display the node's properties to confirm that the `scsi-initiator-id` for each host adapter is set to 7.

```

{0} ok cd /pci@6,4000/pci@3/scsi@5
{0} ok .properties
scsi-initiator-id      00000007
...

```

7. Install the Solaris operating environment.

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

8. Install the `qus` driver and appropriate driver patches.

For driver installation procedures, see the *Sun StorEdge PCI Dual Ultra 3 SCSI Host Adapter Release Notes*.

9. Install the Sun Cluster software and volume management software.

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

Maintaining Storage Arrays

This section provides cluster-specific procedures about how to maintain a storage array. See the documentation that shipped with your storage array for general maintenance procedures.

TABLE 1-1 Task Map: Maintaining a Storage Array

Task	For Instructions, Go To
Add an storage array.	"How to Add a Storage Array" on page 16
Remove an storage array.	"How to Remove a Storage Array" on page 22
Replace an storage array.	"How to Remove a Storage Array" on page 22
To replace a failed storage array, remove the failed storage array and add a new storage array to the configuration.	"How to Add a Storage Array" on page 16
Replace a chassis.	"How to Replace the Chassis" on page 23
Add a disk drive.	"How to Add a Disk Drive" on page 24
Remove a disk drive.	"How to Remove a Disk Drive" on page 25
Replace a disk drive.	"How to Replace a Disk Drive" on page 26
Replace an Environmental Service Module (ESM).	"How to Replace an Environmental Service Module (ESM)" on page 29
Upgrade disk drive firmware.	"How to Upgrade Disk Drive Firmware" on page 30
Update host adapter firmware.	"How to Upgrade Host Adapter Firmware" on page 30
Replace a host adapter card.	"How to Replace a Host Adapter" on page 31

▼ How to Add a Storage Array

Use this procedure to add a storage array to a running cluster. Perform the steps in this procedure in conjunction with the procedures in the Sun Cluster system administration documentation and your server hardware manual.

This procedure uses an updated method for setting the `scsi-initiator-id`. This new method must be used with the StorEdge PCI Dual Ultra3 SCSI host adapter. If you are adding this storage array to a cluster configuration that uses the previous

method of setting the `scsi-initiator-id`, you will need to update the `nvrsrc` script and set the `scsi-initiator-id` by using this procedure. The previous method relied on setting each node's `scsi-initiator-id` for shared storage in the `nvrsrc` script.

Note – To determine if the cluster configuration uses the previous method of setting the `scsi-initiator-id`, run the following command on the two nodes that you plan to connect to this storage array.

```
# eeprom scsi-initiator-id
```

If `scsi-initiator-id` equals 7 on both nodes, you must update the `nvrsrc` script and set the `scsi-initiator-id` for each node that connects to this storage array.

For conceptual information on multi-initiator SCSI and device IDs, see the Sun Cluster concepts documentation document.

This procedure defines Node 1 as the node with which you begin working. Node 2 is the remaining node.

1. Verify that the storage array is configured for the split-bus or single-bus configuration.

See the *Sun Storage D2 Array Installation, Operation, and Service Manual* for procedures about how to configure the ESMs.

2. Shut down Node 1.

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

3. If necessary, install the host adapters in Node 1.

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

4. Connect the SCSI cable between Node 1 and the storage array, as shown in Figure 1–3 for split-bus configurations. See Figure 1–4 for the single-bus configuration cabling diagram.

Make sure that the *entire* bus length that is connected to the storage array is less than 12m. The bus length includes the cables from the ESM to the host adapter. See the documentation that shipped with your hardware for internal bus length.




FIGURE 1-3 Adding a Split-Bus Configuration: Part I




FIGURE 1-4 Adding a Single-Bus Configuration: Part I

5. Ensure that each device in the SCSI chain has a unique SCSI address.

To avoid SCSI-chain conflicts, this procedure instructs you to reserve SCSI address 7 for one host adapter in the SCSI chain and change the other host adapter's global `scsi-initiator-id` to an available SCSI address. Then the procedure instructs you to change the `scsi-initiator-id` for local devices back to 7.

Note – A slot in the storage array might not be in use. However, do not set the `scsi-initiator-id` to a SCSI address for that disk slot. This precaution minimizes future complications if you install additional disk drives.

a. If necessary, power on the storage devices.

For the procedure on powering on a storage device, see the service manual that shipped with your storage device.

- b. If necessary, power on a node, but do not allow it to boot. If necessary, halt the system to continue with OBP Monitor tasks.
- c. Set the `scsi-initiator-id` for one node to 6.

```
{1} ok setenv scsi-initiator-id 6
scsi-initiator-id = 6
```

- d. Find the paths to the host adapters that connect to the local disk drives.

```
{0} ok show-disks
```

Use this information to change the SCSI addresses in the `nvrामrc` script. Do not include the `/sd` directories in the device paths.

- e. Edit the `nvrामrc` script to set the `scsi-initiator-id` for the local devices on the first node to 7.

For a full list of commands, see the *OpenBoot 2.x Command Reference Manual*.



Caution – Insert exactly one space after the first double quote and before `scsi-initiator-id`.

```
{0} ok nvedit
0: probe-all
1: cd /pci@1f,4000/scsi@2
2: 7 "scsi-initiator-id" integer-property
3: device-end
4: cd /pci@1f,4000/scsi@3
5: 7 "scsi-initiator-id" integer-property
6: device-end
7: install-console
8: banner[Control C]
{0} ok
```

- f. Store the changes.

The changes you make through the `nvedit` command are recorded on a temporary copy of the `nvrामrc` script. You can continue to edit this copy without risk. After you complete your edits, save the changes. If you are not sure about the changes, discard them.

- To store the changes, type:

```
{0} ok nvstore
{1} ok
```

- To discard the changes, type:

```
{0} ok nvquit
{1} ok
```

- g. Verify the contents of the `nvrarc` script that you created, as shown in the following example.

If the contents of the `nvrarc` script are incorrect, use the `nvedit` command to make corrections.

```
{0} ok printenv nvrarc
nvrarc =
        probe-all
        cd /pci@1f,4000/scsi@2
        7 " scsi-initiator-id" integer-property
        device-end
        cd /pci@1f,4000/scsi@3
        7 " scsi-initiator-id" integer-property
        device-end
        install-console
        banner

{1} ok
```

- h. Instruct the OpenBoot PROM (OBP) Monitor to use the `nvrarc` script, as shown in the following example.

```
{0} ok setenv use-nvrarc? true
use-nvrarc? = true
{1} ok
```

6. Perform a reconfiguration boot on Node 1 to create the new Solaris device files and links.

7. If necessary, format and label the disks.

8. On Node 1, verify that the DIDs have been assigned to the disk drives in the storage array.

```
# sddidadm -l
```

9. Shut down Node 2.

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

10. If necessary, install the host adapters in Node 2.

For the procedure about how to install a host adapter, see the documentation that shipped with your host adapter and node.

11. Connect the storage array to the host adapters by using SCSI cables.

See Figure 1–5 for a split-bus configuration diagram and Figure 1–6 for a single-bus configuration diagram.




FIGURE 1-5 Adding a Split-Bus Configuration: Part II




FIGURE 1-6 Adding a Single-Bus Configuration: Part II

12. Verify that the `scsi-initiator-id` is set correctly on the second node.

- a. If necessary, power on the second node, but do not allow it to boot. If necessary, halt the system to continue with OBP Monitor tasks.
- b. Verify that the `scsi-initiator-id` for each host adapter on the second node is set to 7

Use the `show-disks` command to find the paths to the host adapters that are connected to these enclosures. Select each host adapter's device tree node, and display the node's properties to confirm that the `scsi-initiator-id` for each host adapter is set to 7.

```
{0} ok cd /pci@6,4000/pci@3/scsi@5
{0} ok .properties
scsi-initiator-id      00000007
...
```

13. Perform a reconfiguration boot on Node 2 to create the new Solaris device files and links.

14. On Node 2, verify that the DIDs have been assigned to the disk drives in the storage array.

```
# sddidadm -L
```

15. 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 Remove a Storage Array

Use this procedure to remove a storage array from a cluster. This procedure assumes that you are removing the references to the disk drives in the enclosure.

1. Determine if the storage array you want to remove contains any quorum devices.

```
# scstat -q
```

- If the storage array does not contain any quorum devices, go to Step 3.
- If the storage array you want to remove contains a quorum device and you need a quorum device to have quorum, choose another disk drive to configure as the 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.

2. If necessary, back up the metadvice or volume.

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

3. 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.

4. Disconnect the SCSI cables from the storage array.

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

For the procedure about how to power off a storage array, see the *Sun StorEdge D2 Installation, Operation, and Service Manual*.

6. Identify the disk drives you need to remove.

```
# cfgadm -al
```

7. Remove the storage array.

For the procedure about how to remove a storage array, see the *Sun StorEdge D2 Installation, Operation, and Service Manual*.

8. **On all nodes, remove references to the disk drives in the storage array that you removed.**

```
# cftadm -c unconfigure cN::dsk/cNtXdY
# devfsadm -C
# scdidadm -C
```

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

For the procedure about how to remove a host adapter, see “How to Replace a Host Adapter” on page 31 and the documentation that shipped with your host adapter and node.

10. **From any node, verify that the configuration is correct.**

```
# scdidadm -L
```

▼ How to Replace the Chassis

Use this procedure to replace the chassis of a storage array in a running cluster. This procedure assumes that you want to retain the disk drives in the storage array you are replacing. This procedure also assumes that you want to retain the references to these disk drives.

1. **Determine if the chassis you want to replace contains a quorum device.**

```
# scstat -q
```

- If the chassis doesn’t contain a quorum device, go to Step 2.
- If the chassis you want to replace contains a quorum device and you need the quorum device to maintain quorum, choose another disk drive on another storage array to configure as the 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.

2. **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.

3. **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.

4. **Disconnect the SCSI cables from the storage array.**

5. **Power off and disconnect the storage array from the AC power source.**
For the procedure about how to power off a storage array, see the *Sun StorEdge D2 Array Installation, Operation, and Service Manual*.
6. **Connect the new storage array to an AC power source.**
7. **Connect the SCSI cables to the new storage array.**
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.**
Move your other components to the new chassis as well. For the procedure about how to replace your storage array's components, see the *Sun StorEdge D2 Array Installation, Operation, and Service Manual*.
9. **Power on the storage array.**
To power on the storage arrays, press the power switches on the power supplies to the On position.
10. **On all nodes that are attached to the storage array, run the `devfsadm(1M)` command.**

```
# devfsadm
```
11. **From one node, attach the new storage array to the global device namespace.**

```
# scgdevs
```
12. **One at a time, shut down and reboot the nodes that are connected to the storage array.**
For the procedure about how to shut down and power off a node, see your Sun Cluster system administration documentation.
13. **Perform volume management administration to add the storage array back into the configuration.**
For more information, see your Solstice DiskSuite/Solaris Volume Manager or VERITAS Volume Manager documentation.

▼ How to Add a Disk Drive

Use this procedure to add a disk drive to a running cluster. Perform the steps in this procedure in conjunction with the procedures in the Sun Cluster system administration documentation and your server hardware manual.

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

1. **Record target number of the empty disk slot to which you are adding the disk drive.**

Identify the disk slot in the storage array for the disk drive that you are adding and note the target number. Refer to the documentation that shipped with your storage array.

2. Install the disk drive.

For the procedure about how to install a disk drive, see the *Sun StorEdge D2 Array Installation, Operation, and Service Manual*.

3. On all nodes attached to the storage array, configure the disk drive into the cluster.

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

4. If necessary, use the `format(1M)` command or the `fmthard(1M)` command to partition and label the disk drive.

5. From any node, 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 an expected behavior.

6. On one node, verify that a device ID (DID) has been assigned to the disk drive.

```
# scdidadm -L
```

7. Perform volume management administration to add the new disk drive to the configuration.

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

8. 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 your Sun Cluster system administration documentation.

▼ How to Remove a Disk Drive

Use this procedure to remove a disk drive from a storage array in a running cluster. Perform the steps in this procedure in conjunction with the procedures in the Sun Cluster system administration documentation and the *Sun StorEdge D2 Array Installation, Operation, and Service Manual*.

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

1. Identify the disk drive you are removing. Identify the slot from which you are removing the disk drive.

If the disk error message reports the drive problem by DID, use the `scdidadm -l` command to determine the Solaris device name. To list all configurable hardware information, use the `cfgadm -al` command.

```
# scdidadm -l deviceID
# cfgadm -al
```

2. Is the disk drive you are removing a quorum device?

```
# scstat -q
```

- If no, proceed to Step 3.
- If yes, choose another disk drive to configure as the quorum device. Remove the old quorum device.

For procedures about how to add and remove quorum devices, see your 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. Remove the disk drive.

For the procedure about how to remove a disk drive, see the *Sun StorEdge D2 Array Installation, Operation, and Service Manual*.

6. On all nodes, unconfigure the disk drive and remove references to the disk drive from the operating system and clustering environment.

```
# cfgadm -c unconfigure cN::disk/cNtXdY
# devfsadm -C
# scdidadm -C
```

▼ How to Replace a Disk Drive

Use this procedure to replace a disk drive in a storage array. Perform the steps in this procedure in conjunction with the procedures in the Sun Cluster system administration documentation and the *Sun StorEdge D2 Array Installation, Operation, and Service Manual*.

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

1. Identify the failed disk drive.

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

2. Determine if the disk drive you want to replace is a quorum device.

```
# scstat -q
```

- If the disk is not a quorum device, skip this step.
- If the disk drive you want to replace is a quorum device, add a new quorum device on a different storage device and remove the old quorum device.

For the procedure on how to add and remove a quorum device, see your 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. If you are using Solstice DiskSuite/Solaris Volume Manager as your volume manager, save the disk partitioning for use when you partition the new disk drive.

```
# prtvtoc /dev/rdsk/cNtXdYsZ > filename
```

Note – Do not save this file under /tmp because you will lose this file when you reboot. Instead, save this file under /usr/tmp.

5. Replace the failed disk.

a. Determine which node owns the device group.

```
# scstat -D
```

b. If you are using VERITAS Volume Manager, remove the LUNs from the VERITAS Volume Manager control on a node that does not have ownership of the device group.

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

c. On a node that does not have ownership of the device group, suspend activity on the SCSI bus.

```
# cfgadm -x replace_device cX::disk/cXtYdZ
```

When prompted, enter **y** to suspend activity on the SCSI bus.

d. After SCSI bus activity stops, replace the disk and enter y at the prompt.

After replacing the disk, warning messages might display. These messages can be ignored.

e. Switch the primary for the device group.

```
# scswitch -z -D disk-device-group -h node
```

- f. If you are using VERITAS Volume Manager, remove the LUNs from the VERITAS Volume Manager control on a node that does not have ownership of the device group.

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

- g. Suspend activity on the device group's original primary.

```
# cfgadm -x replace_device cX::disk/cXtYdZ
```

When prompted, enter **y** to suspend activity on the SCSI bus.

- h. When prompted to proceed with the disk replacement, enter **y** to proceed.

Warning messages might display. These messages can be ignored.

6. On all nodes that are attached to the storage device, run the **devfsadm(1M)** command to probe all devices and to write the new disk drive to the **/dev/rdisk** directory.

Depending on the number of devices connected to the node, the **devfsadm** command can require at least five minutes to complete.

```
# devfsadm
```

7. Label the new disk drive by using the **format** command.

8. If you are using Solstice DiskSuite/Solaris Volume Manager as your volume manager, from any node that is connected to the storage device, partition the new disk drive by using the partitioning you saved.

If you are using VERITAS Volume Manager, skip this step.

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

9. From all nodes that are connected to the disk drive, update the DID database and driver.

```
# scdidadm -R deviceID
```

Note – After running **scdidadm -R** on the first node, each subsequent node that you run the command on might display a warning. This warning can be ignored.

10. From any node, confirm that the failed disk drive has been replaced by comparing the new physical DID to the physical DID that was identified at the beginning of this procedure.

If the new physical DID is different, you successfully replaced the failed disk drive with a new disk drive.

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

11. 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.

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

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

▼ How to Replace an Environmental Service Module (ESM)

Use this procedure to replace the Environmental Service Module (ESM) in storage array.

1. Determine if the ESM that you plan to replace provides access to a quorum device.

- If the ESM does not provide access to a quorum device, go to Step 2.
- If the ESM provides access to a quorum device and you need a quorum device to have quorum, choose another disk drive not connected to the ESM which is not connected to the ESM to configure as a quorum device. Then remove the old quorum device.

To remove a quorum device from your configuration, see your Sun Cluster system administration documentation.

2. 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.

3. 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.

4. Replace the ESM

For the procedure about how to replace an ESM, see the *Sun StorEdge D2 Array Installation, Operation, and Service Manual*.

5. Perform volume management administration to add the storage array to the configuration.

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

▼ How to Upgrade Disk Drive Firmware

Use this procedure to upgrade disk drive firmware.

1. **Determine if the disk drive that you want to upgrade the firmware on is a quorum device.**

```
# scstat -q
```

- If the disk is not a quorum device, go to Step 2.
- If the disk drive is a quorum device and you need a quorum device to maintain quorum, choose another disk drive to configure as the 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.

2. **If possible, back up the metadevices or volumes that reside in the disk drive.**

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

3. **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.

4. **Install firmware.**

For more information about installing firmware, see the patch installation instructions.

5. **Perform volume management administration to add the disk drive back into the configuration.**

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

▼ How to Upgrade Host Adapter Firmware

Use this procedure to update host adapter firmware.

1. **Determine the resource groups and device groups that are running on the node that you plan to upgrade the host adapter firmware on.**

```
# scstat
```

2. **Migrate the resource groups and device groups off the node that you plan to upgrade the firmware on.**

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

3. **Perform the firmware upgrade.**

For the procedure about how to upgrade your host adapter firmware, see the patch documentation.

4. Return the resource groups and device groups to their original node.

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

▼ How to Replace a Host Adapter

1. Determine the resource groups and device groups that are running on the node that contains the host adapter that you want to replace.

```
# scstat
```

2. Migrate the resource groups and device groups off the node that contains the host adapter that you want to replace and shut down the node

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

3. Replace the host adapter card.

For the procedure about how to replace the host adapter card, see the documentation that shipped with your hardware.

4. Power on the node with the replaced host adapter card.

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

5. Boot the node into cluster mode.

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

6. Return the resource groups and device groups to their original node.

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


Index

A

adding
 See also installing
 disk drives, 24
 storage arrays, 16
arrays, *See* storage arrays

C

chassis, replacing, 23
configurations, 11

D

disk drive firmware, upgrading, 30
disk drives
 adding, 24
 removing, 25
 replacing, 26

E

environmental service module, replacing, 29

H

help, 8
host adapter, replacing, 31
host adapter firmware, upgrading, 30

I

installing
 See also adding
 storage arrays, 13

R

removing
 disk drives, 25
 storage arrays, 22
replacing
 chassis, 23
 disk drives, 26
 environmental service module, 29
 host adapter, 31

S

storage arrays
 adding, 16
 installing, 13
 removing, 22

T

technical support, 8

U

upgrading

- disk drive firmware, 30

- host adapter firmware, 30