



Sun Cluster 3.x With Sun StorEdge A1000 or Netra st A1000 Array Manual

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

Part No: 817-0171-10
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 Sun StorEdge A1000 or Netra st A1000 Arrays	11
Installing Storage Arrays	11
▼ How to Install a Pair of Storage Arrays	12
Configuring Storage Arrays	16
▼ How to Create a LUN	17
▼ How to Delete a LUN	19
▼ How to Reset a LUN Configuration	20
▼ How to Correct Mismatched DID Numbers	22
Maintaining Storage Arrays	23
▼ How to Add a Pair of Storage Arrays	25
▼ How to Remove a Storage Array	31
▼ How to Replace a Failed Controller or Restore an Offline Controller	33
▼ How to Add a Disk Drive	34
▼ How to Replace a Disk Drive	35
▼ How to Remove a Disk Drive	36
How to Upgrade Disk Drive Firmware	36
▼ How to Replace a Node's Host Adapter	36
Index	39

Preface

The *Sun Cluster 3.x With Sun StorEdge A1000 or Netra st A1000 Array Manual* provides procedures specific to Sun StorEdge™ A1000 and Netra™ st A1000 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 three major sections.

Section 1 discusses how to install Sun StorEdge A1000 and Netra st A1000 storage arrays

Section 2 discusses how to configure logical units on Sun StorEdge A1000 and Netra st A1000 storage arrays

Section 3 describes how to maintain Sun StorEdge A1000 and Netra st A1000 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 A1000 or Netra st A1000 Array Manual

Revision Date	New Information
July 2003	Updated illustrations.
	Updated documentation with editorial changes.
	Provided a list of procedures that require no cluster-specific procedures.

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>Netscaler A1000/D1000 Installation and Maintenance Manual</i>	805-7147
<i>OpenBoot 2.x Command Reference Manual</i>	806-2906
<i>Sun StorEdge A1000 and D1000 Installation, Operations, and Service Manual</i>	805-2624
<i>Sun StorEdge A3500/A3500FC Controller Module Guide</i>	805-4980
Available on http://www.sun.com/products-n-solutions/hardware/docs	
<i>Sun StorEdge D1000 Storage Guide</i>	805-4013
<i>Sun StorEdge MultiPack User's Guide</i>	Unavailable online
Shipped with your storage	
<i>Sun StorEdge RAID Manager Release Notes</i>	805-7758
<i>Sun StorEdge RAID Manager User's Guide</i>	806-0478
<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

TABLE P-3 Sun Cluster Documentation (Continued)

Application	Title	Part Number
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

Command	Function
<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 <i>rm</i> <i>filename</i> .

TABLE P-4 Typographic Conventions (Continued)

Typeface or Symbol	Meaning	Example
<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	machine_name%
C shell superuser prompt	machine_name#
Bourne shell and Korn shell prompt	\$
Bourne shell and Korn shell superuser prompt	#

Installing and Maintaining Sun StorEdge A1000 or Netra st A1000 Arrays

This chapter describes the procedures about how to install, configure, and maintain a Sun StorEdge™ A1000 array and a Netra™ st A1000 array in a Sun™ Cluster environment. The procedures in this chapter apply to both the Sun StorEdge A1000 and the Netra st A1000 arrays.

This chapter contains the following procedures.

- “How to Install a Pair of Storage Arrays” on page 12
- “How to Create a LUN” on page 17
- “How to Delete a LUN” on page 19
- “How to Reset a LUN Configuration” on page 20
- “How to Correct Mismatched DID Numbers” on page 22
- “How to Add a Pair of Storage Arrays” on page 25
- “How to Remove a Storage Array” on page 31
- “How to Add a Disk Drive” on page 34
- “How to Replace a Disk Drive” on page 35
- “How to Remove a Disk Drive” on page 36
- “How to Upgrade Disk Drive Firmware” on page 36
- “How to Replace a Failed Controller or Restore an Offline Controller” on page 33
- “How to Replace a Node’s Host Adapter” on page 36

Installing Storage Arrays

This section provides the procedure for an initial installation of a pair of storage arrays to a nonconfigured cluster. To add storage arrays to an operating cluster, use the procedure in “How to Add a Pair of Storage Arrays” on page 25.

▼ How to Install a Pair of Storage Arrays

Use this procedure to install and configure a pair of storage arrays, *before* you install the Solaris operating environment and Sun Cluster software on your nodes.

This procedure uses an updated method for setting the `scsi-initiator-id`. For this storage device, the method published in earlier documentation is still applicable. However, if your cluster configuration uses a Sun StorEdge PCI Dual Ultra3 SCSI host adapter to connect to any other shared storage, you will need to update your `nvrnrc` script and set the `scsi-initiator-id` by using this procedure.

1. Install the host adapters in the nodes that connect to the storage arrays.

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

2. Cable the storage arrays.

The storage arrays must be configured in pairs for the Sun Cluster environment.

To cable the storage arrays, connect the differential SCSI cable between the node and the storage array. Make sure that the *entire* SCSI bus length in each SCSI chain is less than 25 m. This measurement includes the cables to both nodes, as well as the bus length that is internal to each enclosure, node, host adapter. Figure 1–1 and Figure 1–2 illustrate a storage array that is cabled in a Sun Cluster environment.




FIGURE 1–1 Storage Array Cabling

Figure 1–2 illustrates how to cable the storage array in a daisy-chain configuration. If you place two storage arrays on a SCSI chain, ensure that each storage array's ID dial is set to a unique number.




FIGURE 1-2 Storage Array With DaisyChaining

3. 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 `nvrnrc` 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 `nvrnrc` 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 `nvrnrc` script that you created, as shown in the following example.

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

```
{0} ok printenv nvrnrc
nvrnrc =
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
```

4. 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
...
```

5. Install the Solaris operating environment, then apply any required Solaris patches.

Note – For the current list of patches that are required for the Solaris operating environment, refer to SunSolve. SunSolve is available online to Sun service providers and to customers with SunSolve service contracts at the SunSolve site. <http://sunsolve.sun.com>.

6. Install the RAID Manager software.

For the procedure about how to install the RAID Manager software, see the *Sun StorEdge RAID Manager User's Guide*.

Note – RAID Manager 6.22.1 is required for clustering the storage array with Sun Cluster 3.0.

7. Install any storage array or RAID Manager patches.

Note – For the most current list of software, firmware, and patches that are required for the storage array, refer to *Info Docs 20029, A1000/A3x00/A1000FC Software/Firmware Configuration Matrix*. This document is available online to Sun service providers and to customers with SunSolve service contracts at the SunSolve site, <http://sunsolve.sun.com>, under advanced search.

8. Check the storage array NVSRAM file revision. If necessary, install the most recent revision.

For the NVSRAM file revision number, boot level, and procedure about how to upgrade the NVSRAM file, see the *Sun StorEdge RAID Manager Release Notes*.

9. Set the Rdac parameters in the `/etc/osa/rmparams` file on both nodes.

```
Rdac_RetryCount=1
Rdac_NoAltOffline=TRUE
```

10. Set up the storage arrays with logical units (LUNs) and hot spares.

For the procedure about how to set up the storage array with LUNs and hot spares, see the *Sun StorEdge RAID Manager User's Guide*.

Note – Use the `format` command to verify Solaris logical device names.

11. Ensure that the new logical name for the LUN that you created in Step 10 appears in the `/dev/rdsd` directory on both nodes.

```
# /etc/raid/bin/hot_add
```

Where to Go From Here

To continue with Sun Cluster software and data services installation tasks, see your Sun Cluster software installation documentation and the Sun Cluster data services developer's documentation.

Configuring Storage Arrays

This section contains the procedures about how to configure a storage array *after* you install Sun Cluster software. Table 1–1 lists these procedures.

To configure a storage array *before* you install Sun Cluster software, use the same procedures that you use in a noncluster environment. For the procedures about how to configure storage arrays before you install Sun Cluster, see the *Sun StorEdge RAID Manager User's Guide*.

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 Storage Array Disk Drives

Task	Information
Create a logical unit number (LUN).	"How to Create a LUN" on page 17
Remove a LUN.	"How to Delete a LUN" on page 19
Reset the storage array configuration.	"How to Reset a LUN Configuration" on page 20

The following is a list of administrative tasks that require no cluster-specific procedures. See the *Sun StorEdge RAID Manager User's Guide* and the *Sun StorEdge RAID Manager Release Notes* for the following procedures.

- Create a hot spare.
- Delete a hot spare.
- Increase the size of a drive group.

▼ How to Create a LUN

Use this procedure to create a logical unit (LUN) from unassigned disk drives or remaining capacity. See the *Sun StorEdge RAID Manager Release Notes* for the latest information about LUN administration.

This product supports the use of hardware RAID and host-based software RAID. For host-based software RAID, this product supports RAID levels 0+1 and 1+0.

Note – When you use host-based software RAID with hardware RAID, the hardware RAID levels you use affect hardware maintenance. If you use hardware RAID level 1, 3, or 5, you can perform most maintenance procedures in “Maintaining Storage Arrays” on page 23 without volume management disruptions. If you use hardware RAID level 0, some maintenance procedures in “Maintaining Storage Arrays” on page 23 require additional volume management administration because the availability of the LUNs is impacted.

1. **With all nodes booted and attached to the storage array, create the LUN on one node.**

After the LUN formatting completes, a logical name for the new LUN appears in `/dev/rdisk` on all nodes. These nodes are attached to the storage array.

For the procedure about how to create a LUN, see the *Sun StorEdge RAID Manager User's Guide*.

If the following warning message displays, ignore the message. Continue with the next step.

```
scsi. WARNING. /sbus@e,0/QLGC,isp@1,10000/sd@2,1
(sd153):corrupt label - wrong magic number
```

Note – Use the `format(1M)` command to verify Solaris logical device names. If necessary, label the LUN.

2. **Ensure that the new logical name for the LUN that you created in Step 1 appears in the `/dev/rdisk` directory on both nodes.**

```
# /etc/raid/bin/hot_add
```

3. **On one node, update the global device namespace.**

```
# scgdevs
```

4. **Use the `sddidadm` command to verify that the DID numbers for the LUNs are the same on both nodes. In the sample output that follows, the DID numbers are different.**

```
# sddidadm -L
...
33      e07a:/dev/rdisk/c1t4d2      /dev/did/rdsk/d33
33      e07c:/dev/rdisk/c0t4d2      /dev/did/rdsk/d33
```

5. **Are the DID numbers you received from running the `sddidadm` command in Step 4 the same for both your nodes?**

- If yes, proceed to Step 6.

- If no, perform the procedure in “How to Correct Mismatched DID Numbers” on page 22 before you continue with Step 6 of this procedure.
6. **If you want a volume manager to manage the new LUN, incorporate the LUN into a diskset or disk group.**
For more information, see your Solstice DiskSuite/Solaris Volume Manager or VERITAS Volume Manager documentation.
 7. **If you want the new LUN to be a quorum device, add the quorum device.**
For the procedure about how to add a quorum device, see the *Sun Cluster 3.0 U2 System Administration Guide*

▼ How to Delete a LUN

Use this procedure to delete one or more LUNs. See the *Sun StorEdge RAID Manager Release Notes* for the latest information about LUN administration.



Caution – This procedure removes all data on the LUN that you delete.



Caution – Do not delete LUN 0.

1. **From one node that is connected to the storage array, determine the paths to the LUN. This LUN is the LUN that you are deleting.**

```
f28c# format
Searching for disks...done
AVAILABLE DISK SELECTIONS:
  0. c0t10d0 <SUN18G cyl 7506 alt 2 hd 19 sec 248>
     /sbus@3,0/SUNW,fas@3,8800000/sd@a,0
  1. c1t5d0 <Symbios-StorEDGEA1000-0301 cyl 12160 alt 2 hd 64 sec 64>
     /pseudo/rdnexus@1/rdriver@5,0
  2. c2t2d0 <Symbios-StorEDGEA1000-0301 cyl 12160 alt 2 hd 64 sec 64>
     /pseudo/rdnexus@2/rdriver@2,0
```

2. **Is the LUN a quorum device? This LUN is the LUN that you are removing.**

```
# scstat -q
```

- If no, proceed to Step 3.
- If yes, relocate that quorum device to another suitable storage array.

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

3. Remove the LUN from disksets or disk groups.

Remove the LUN from any diskset or disk group. For more information, see your Solstice DiskSuite/Solaris Volume Manager or VERITAS Volume Manager documentation.

LUNs that were managed by VERITAS Volume Manager must be removed from VERITAS Volume Manager control before you can delete the LUNs. To remove the LUNs, after you delete the LUNs from any disk group, use the following commands.

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

4. From one node, delete the LUN.

For the procedure about how to delete a LUN, see the *Sun StorEdge RAID Manager User's Guide*.

5. From the same node, remove the paths to the LUNs you are deleting.

```
# rm /dev/rdisk/cNtXdY*
# rm /dev/dsk/cNtXdY*
# rm /dev/osa/dev/dsk/cNtXdY*
# rm /dev/osa/dev/rdisk/cNtXdY*
```

6. From the same node, remove all obsolete device IDs (DID).

```
# scdidadm -C
```

7. Move all resource groups and device groups off the node.

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

8. Shut down the node.

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

9. Perform a reconfiguration boot to create the new Solaris device files and links.

10. Repeat Step 7 through Step 9 on the other node. This node is attached to the storage array.

▼ How to Reset a LUN Configuration

Use this procedure to reset a storage array LUN configuration.



Caution – If you reset LUN configuration, a new DID number is assigned to LUN 0. This change occurs because the software assigns a new worldwide number (WWN) to the new LUN.

1. From one node that is connected to the storage array, determine the paths to the LUN. This LUN is the LUN that you are resetting.

```
f28c# format
Searching for disks...done
AVAILABLE DISK SELECTIONS:
  0. c0t10d0 <SUN18G cyl 7506 alt 2 hd 19 sec 248>
     /sbus@3,0/SUNW,fas@3,8800000/sd@a,0
  1. c1t5d0 <Symbios-StorEDGEA1000-0301 cyl 12160 alt 2 hd 64 sec 64>
     /pseudo/rdnexus@1/rdriver@5,0
  2. c2t2d0 <Symbios-StorEDGEA1000-0301 cyl 12160 alt 2 hd 64 sec 64>
     /pseudo/rdnexus@2/rdriver@2,0
```

2. Is the LUN a quorum device? This LUN is the LUN that you are resetting.

```
# scstat -q
```

- If no, proceed to Step 3.
- If yes, relocate that quorum device to another suitable storage array.
For procedures about how to add and remove quorum devices, see your Sun Cluster system administration documentation.

3. Remove the LUN from disksets or disk groups.

Remove the LUN from any diskset or disk group. For more information, see your Solstice DiskSuite/Solaris Volume Manager or VERITAS Volume Manager documentation.

You must completely remove LUNs that were managed by VERITAS Volume from VERITAS Volume Manager control before you can delete the LUNs.

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

4. On one node, reset the LUN configuration.

For the procedure about how to reset storage array LUN configuration, see the *Sun StorEdge RAID Manager User's Guide*.

Note – Use the `format` command to verify Solaris logical device names.

5. By using the `format` command, label the new LUN 0.

6. Remove the paths to the old LUN that you reset.

```
# rm /dev/rdisk/cNtXdY*
# rm /dev/dsk/cNtXdY*

# rm /dev/osa/dev/dsk/cNtXdY*
# rm /dev/osa/dev/rdisk/cNtXdY*
```

7. Update device namespaces on both nodes.

```
devfsadm -C
```

8. Remove all obsolete DIDs on both nodes.

```
# sddidadm -C
```

9. Move all resource groups and device groups off the node.

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

10. Shut down the node.

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

11. Perform a reconfiguration boot to create the new Solaris device files and links.

If the following error message appears, ignore this message. Continue with the next step.

```
device id for '/dev/rdisk/c0t5d0' does not match physical disk's id.
```

12. After the node reboots and joins the cluster, repeat Step 6 through Step 11 on the other node. This node is attached to the storage array.

The DID number for the original LUN 0 is removed. A new DID is assigned to LUN 0.

▼ How to Correct Mismatched DID Numbers

Use this section to correct mismatched device ID (DID) numbers that might appear during the creation of storage array LUNs. You correct the mismatch by deleting Solaris and Sun Cluster paths to the LUNs that have DID numbers that are different. After rebooting, the paths are corrected.

Note – Use this procedure only if you are directed to do so from “How to Create a LUN” on page 17.

1. From one node that is connected to the storage array, determine the paths to the LUNs. These LUNs have different DID numbers.

```
# format
```

2. Remove the paths to the LUNs that have different DID numbers.

```
# rm /dev/rdisk/cNtXdY*
```

```
# rm /dev/dsk/cNtXdY*
```

```
# rm /dev/osa/dev/dsk/cNtXdY*
```

```
# rm /dev/osa/dev/rdisk/cNtXdY*
```

3. Use the `lad` command to determine the *alternate* paths to the LUNs that have different DID numbers.

The RAID Manager software creates two paths to the LUN in the `/dev/osa/dev/rdisk` directory. Substitute the `cNtXdY` number from the other storage array to determine the alternate path.

For example, with this configuration.

```
# lad
c0t5d0 1T93600714 LUNS: 0 1
c1t4d0 1T93500595 LUNS: 2
```

The alternate paths would be as follows.

```
/dev/osa/dev/dsk/c1t4d1*
/dev/osa/dev/rdisk/c1t4d1*
```

4. Remove the *alternate* paths to the LUNs that have different DID numbers.

```
# rm /dev/osa/dev/dsk/cNtXdY*
# rm /dev/osa/dev/rdisk/cNtXdY*
```

5. On both nodes, remove all obsolete DIDs.

```
# scdidadm -C
```

6. Move all resource groups and device groups off the node.

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

7. Shut down the node.

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

8. Perform a reconfiguration boot to create the new Solaris device files and links.

9. Repeat Step 1 through Step 8 on the other node that is attached to the storage array.

10. Return to “How to Create a LUN” on page 17.

Maintaining Storage Arrays

This section contains the procedures about how to maintain a storage array in a Sun Cluster environment.

Some maintenance procedures in Table 1–2 are performed the same as in a noncluster environment. This section references these procedures, but this section does not contain these procedures. Table 1–2 lists the procedures about how to maintain the 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 1-2 Tasks. Maintaining a Storage Array

Task	Information
storage array procedures.	
Add a storage array to a running cluster.	"How to Add a Pair of Storage Arrays" on page 25
Remove a storage array from a running cluster.	"How to Remove a Storage Array" on page 31
Replace a failed storage array. To replace a failed storage array, remove the failed storage array. Add a new storage array to the configuration.	"How to Remove a Storage Array" on page 31 "How to Add a Pair of Storage Arrays" on page 25
Add a disk drive.	"How to Add a Disk Drive" on page 34
Replace a disk drive.	"How to Replace a Disk Drive" on page 35
Remove a disk drive.	"How to Remove a Disk Drive" on page 36
Upgrade disk drive firmware.	"How to Upgrade Disk Drive Firmware" on page 36
Replace a failed controller or restore an offline controller.	"How to Replace a Failed Controller or Restore an Offline Controller" on page 33
Replace a node's host adapter.	"How to Replace a Node's Host Adapter" on page 36

The following is a list of administrative tasks that require no cluster-specific procedures. See the *Sun StorEdge A1000 and D1000 Installation, Operations, and Service Manual*, the *Netra st A1000/D1000 Installation and Maintenance Manual* and the *Netra st A1000/D1000 Installation and Maintenance Manual* for the following procedures.

- Replacing a power supply.
- Replacing a battery unit.
- Replacing a storage array cooling canister.
- Replacing a power cable on the storage array.

The following is a list of administrative tasks that require no cluster-specific procedures. See the *Sun StorEdge RAID Manager User's Guide* and the *Sun StorEdge RAID Manager Release Notes* for the following procedures.

- Replacing a storage array-to-host SCSI cable.

▼ How to Add a Pair of Storage Arrays

Use this procedure to add a pair of storage arrays to a running cluster.

This procedure uses an updated method for setting the `scsi-initiator-id`. For this storage device, the method published in earlier documentation is still applicable. However, if your cluster configuration uses a Sun StorEdge PCI Dual Ultra3 SCSI host adapter to connect to any other shared storage, you will need to update your `nvrarc` script and set the `scsi-initiator-id` by using this procedure.

1. Install the RAID Manager software on nodes.

For the procedure about how to install RAID Manager software, see the *Sun StorEdge RAID Manager Installation and Support Guide*.

Note – RAID Manager 6.22 or a compatible version is required for clustering with Sun Cluster 3.0.

2. Install any storage array patches on nodes.

Note – For the most current list of software, firmware, and patches that are required for the storage array, refer to *Info Docs 20029, A1000/A3x00/A1000FC Software/Firmware Configuration Matrix*. This document is available online to Sun service providers and to customers with SunSolve service contracts at the SunSolve site. <http://sunsolve.sun.com>.

3. Set the `Rdac` parameters in the `/etc/osa/rmparams` file on both nodes.

```
Rdac_RetryCount=1
Rdac_NoAltOffline=TRUE
```

4. Power on the storage array.

To power on the storage array, push the power switch to the momentary on position, which is the right side. Release the power switch.

5. Shut down the Node 1.

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

6. If you are installing new host adapters, power off Node 1.

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


7. 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.

8. Cable the storage array to Node 1, as shown in Figure 1-3.

If you are adding a storage array, connect the differential SCSI cable between Node 1 and the storage array. Verify that the *entire* SCSI bus length in each SCSI chain is less than 25m. This measurement includes the cables to both nodes, as well as the bus length that is internal to each enclosure, node, and host adapter.

Since the storage array is not an auto-terminating device, ensure that you place differential SCSI terminators as shown in the figures below.




Ⓣ Differential SCSI terminator

FIGURE 1-3 Storage Array Cabling

The following figure illustrates how to cable daisy-chained storage arrays to one node in the cluster. If you place two storage arrays on a SCSI chain, ensure that

each storage array's ID dial is set to a unique number.



Ⓣ Differential SCSI terminator

FIGURE 1-4 Storage Array Cabling with Daisy Chaining

9. 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 nvramrc script, as shown in the following example.**

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

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

- 11. Check the storage array NVSRAM file and firmware revisions. If necessary, install the most recent revision.**

To verify that you have the current revision, see the *Sun StorEdge RAID Manager Release Notes*. For the procedure about how to upgrade the NVSRAM file and firmware, see the *Sun StorEdge RAID Manager User's Guide*.

- 12. Shut down Node 2.**

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

- 13. If you are installing new host adapters, power off Node 2.**

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

- 14. Install the host adapters in Node 2.**

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

- 15. Cable the storage array to your node, as shown in Figure 1–5.**

Connect the differential SCSI cable between the node and the storage array. Make sure that the *entire* SCSI bus length in each SCSI chain is less than 25m. This measurement includes the cables to both nodes, as well as the bus length that is internal to each enclosure, node, host adapter.




FIGURE 1-5 Storage Array Cabling

The following figure illustrates how to cable daisy-chained storage arrays to the remaining node in the cluster.




FIGURE 1-6 Storage Array With Daisy Chaining

16. 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
...
```

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

18. On one node, verify that the DIDs have been assigned to the storage array LUNs for all nodes. These nodes are attached to the storage array.

```
# scdidadm -L
```

Where to Go From Here

To create a LUN from disk drives that are unassigned, see “How to Create a LUN” on page 17.

To upgrade storage array firmware, see “How to Upgrade Disk Drive Firmware” on page 36.

▼ How to Remove a Storage Array

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



Caution – This procedure removes all data that is on the storage array you remove.

- 1. Migrate any Oracle Parallel Server/Real Application Clusters (OPS) tables, data services, or volumes off the storage array.**
- 2. Is one of the LUNs in the storage array a quorum device? This storage array is the storage array that you are removing.**

```
# scstat -q
```

- If no, proceed to Step 3.
- If yes, relocate that quorum device to another suitable storage array.

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

3. Halt all activity to the storage array.

For instructions, see the *Sun StorEdge RAID Manager User's Guide* and your operating storage array documentation.

4. Remove the LUN from disksets or disk groups.

If a volume manager manages the LUN, remove the LUN from any diskset or disk group. For more information, see your Solstice DiskSuite/Solaris Volume Manager or VERITAS Volume Manager documentation.

You must completely remove LUNs that were managed by VERITAS Volume from VERITAS Volume Manager control before you can delete the LUNs.

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

5. From one node, delete the LUN.

For the procedure about how to delete a LUN, see the *Sun StorEdge RAID Manager User's Guide*.

6. Disconnect all cables from the storage array. Remove the hardware from your cluster.

7. Remove the paths to the LUNs you are deleting.

```
# rm /dev/rdisk/cNtXdY*
# rm /dev/dsk/cNtXdY*

# rm /dev/osa/dev/dsk/cNtXdY*
# rm /dev/osa/dev/rdisk/cNtXdY*
```

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

```
# scdidadm -C
```

9. If you plan to remove a host adapter that has an entry in the `nvrarmrc` script, delete the references to the host adapters in the `nvrarmrc` script.

Note – If there are no other parallel SCSI devices connected to the nodes, you can delete the contents of the `nvrarmrc` script and, at the OpenBoot PROM, set `setenv use-nvrarmrc? false`.

10. Remove any unused host adapters from nodes that were attached to the storage array.

a. Shut down and power off Node 1.

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

b. Remove the host adapter from the Node 1.

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

- c. **Perform a reconfiguration boot to create the new Solaris device files and links.**
- d. **Repeat Step a through Step c for Node 2 that was attached to the storage array.**

11. Return resource groups to their primary nodes.

```
# scswitch -z
```

12. Are you removing the *last* storage array from your cluster?

- If no, you are finished with this procedure.
- If yes, proceed to Step 13.

13. Remove RAID Manager patches, then remove RAID Manager software packages.



Caution – If you improperly remove RAID Manager packages, the next reboot of the node fails. Before you remove RAID Manager software packages, see the *Sun StorEdge RAID Manager Release Notes*.

For the procedure about how to remove software packages, see the documentation that shipped with your storage array.

▼ How to Replace a Failed Controller or Restore an Offline Controller

Use this procedure to replace a storage array controller, or to restore an offline controller.

For conceptual information on SCSI reservations and failure fencing, see your Sun Cluster concepts documentation.

1. Is one of the LUNs in the storage array a quorum device?

```
# scstat -q
```

- If no, proceed to Step 2.
- If yes, relocate that quorum device to another suitable storage array.

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

2. Restart the RAID Manager daemon.

```
# /etc/init.d/amdemon stop
# /etc/init.d/amdemon start
```

3. Do you have a failed controller?

- If your storage array is offline, but does not have a failed controller, proceed to Step 4.
- If you have a failed controller, replace the failed controller with a new controller. *Do not bring the controller online.*

For the procedure about how to replace storage array controllers, see the *Sun StorEdge A3500/A3500FC Controller Module Guide* and the *Sun StorEdge RAID Manager Release Notes*

4. On one node, use the RAID Manager GUI's Recovery application to restore the controller online.

Note – You must use the RAID Manager GUI's Recovery application to bring the controller online.

For information on the Recovery application, see the *Sun StorEdge RAID Manager User's Guide*. If you have problems with bringing the controller online, see the *Sun StorEdge RAID Manager Release Notes*.

5. On one node that is connected to the storage array, verify that the controller has the correct SCSI reservation state.

Run the `scdidadm(1M)` repair option (`-R`) on LUN 0 of the controller you want to bring online.

```
# scdidadm -R /dev/dsk/cNtXdY
```

▼ How to Add a Disk Drive

Use this procedure to add a disk drive to a storage array that is in a running cluster.

1. Verify that the new disk drive is formatted. Verify that you are not transferring the new disk from another storage array.

For information about how to move drives between StorEdge/Netra st storage array subsystems, see the *Sun StorEdge RAID Manager Release Notes*.

2. Install the new disk drive to the storage array.

For the procedure about how to install a disk drive, see the *Sun StorEdge A1000 and D1000 Installation, Operations, and Service Manual* or the *Netra st A1000/D1000 Installation and Maintenance Manual*.

3. Enable the disk drive to spin up for approximately 30 seconds.

4. Run Health Check to ensure that the new disk drive is not defective.

For the procedure about how to run Recovery Guru and Health Check, see the *Sun StorEdge RAID Manager User's Guide*

5. Fail the new drive, then revive the drive to update DacStore on the drive.

For the procedure about how to fail and revive drives, see the *Sun StorEdge RAID Manager User's Guide*.

6. Repeat Step 2 through Step 5 for each disk drive you are adding.

Where to Go From Here

To create LUNs for the new drives, see “How to Create a LUN” on page 17 for more information.

▼ How to Replace a Disk Drive

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

1. Does replacing the disk drive affect any LUN's availability?

- If no, proceed to Step 2.
- If yes, remove the LUNs from volume management control. For more information, see your Solstice DiskSuite/Solaris Volume Manager or VERITAS Volume Manager documentation.

2. Replace the disk drive in the storage array.

For the procedure about how to replace a disk drive, see the *Sun StorEdge D1000 Storage Guide*.

3. Run Health Check to ensure that the new disk drive is not defective.

For the procedure about how to run Recovery Guru and Health Check, see the *Sun StorEdge RAID Manager User's Guide*.

4. Does the failed drive belong to a drive group?

- If no, proceed to Step 5.
- If yes, reconstruction starts automatically. If reconstruction does not start automatically for any reason, then select Reconstruct from the Manual Recovery application. Do not select Revive. When reconstruction is complete, skip to Step 6.

5. Fail the new drive, then revive the drive to update DacStore on the drive.

For the procedure about how to fail and revive drives, see the *Sun StorEdge RAID Manager User's Guide*.

6. If you removed LUNs from volume management control in Step 1, return the LUNs to volume management control.

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

▼ How to Remove a Disk Drive

Use this procedure to remove a disk drive from a running cluster.

1. Is the LUN that is associated with the disk drive a quorum device? This disk drive is the disk drive that you are removing.

```
# scstat -q
```

- If no, proceed to Step 2.
- If yes, relocate that quorum device to another suitable storage array.

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

2. Remove the LUN that is associated with the disk drive you are removing.

For the procedure about how to remove a LUN, see “How to Delete a LUN” on page 19.

3. Remove the disk drive from the storage array.

For the procedure about how to remove a disk drive, see the *Sun StorEdge D1000 Storage Guide*.



Caution – After you remove the disk drive, install a dummy drive to maintain proper cooling.

How to Upgrade Disk Drive Firmware



Caution – You must be a Sun service provider to perform disk drive firmware updates. If you need to upgrade drive firmware, contact your Sun service provider.

▼ How to Replace a Node’s Host Adapter

This section describes the procedure about how to replace a failed host adapter in a running node. This node is attached to a storage array.

In the following procedure, Node 1's host adapter on SCSI bus A needs replacement. Node 2 remains in service.

Note – Several steps in this procedure require that you halt I/O activity. To halt I/O activity, take the storage array offline by using the RAID Manager GUI's manual recovery procedure in the *Sun StorEdge RAID Manager User's Guide*.

1. Without powering off the node, shut down Node 1.

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

2. From Node 2, halt I/O activity to SCSI bus A.

For instructions, see the *Sun StorEdge RAID Manager User's Guide*.

3. From the storage array end of the SCSI cable, disconnect the SCSI bus A cable. This cable connects the storage array to Node 1. Replace this cable with a differential SCSI terminator.

4. Restart I/O activity on SCSI bus A.

For instructions, see the *Sun StorEdge RAID Manager User's Guide*.

5. Does servicing the failed host adapter affect SCSI bus B?

- If no, skip to Step 9.
- If yes, proceed to Step 6.

6. From Node 2, halt I/O activity to the storage array on SCSI bus B.

For instructions, see the *Sun StorEdge RAID Manager User's Guide*.

7. From the storage array end of the SCSI cable, disconnect the SCSI bus B cable. This cable connects the storage array to Node 1. Replace this cable with a differential SCSI terminator.

8. Restart I/O activity on SCSI bus B.

For instructions, see the *Sun StorEdge RAID Manager User's Guide*.

9. Power off Node 1.

10. Replace Node 1's host adapter.

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

11. Power on Node 1. Do not enable the node to boot. If necessary, halt the storage array.

12. From Node 2, halt I/O activity to the storage array on SCSI bus A.

For instructions, see the *Sun StorEdge RAID Manager User's Guide*.

13. Remove the differential SCSI terminator from SCSI bus A. Reinstall the SCSI cable to connect the storage array to Node 1.

14. Restart I/O activity on SCSI bus A.

For instructions, see the *Sun StorEdge RAID Manager User's Guide*.

15. Did you install a differential SCSI terminator to SCSI bus B in Step 7?

- If no, skip to Step 18.
- If yes, halt I/O activity on SCSI bus B, then continue with Step 16.

16. Remove the differential SCSI terminator from SCSI bus B. Reinstall the SCSI cable to connect the storage array to Node 1.

17. Restart I/O activity on SCSI bus B.

For instructions, see the *Sun StorEdge RAID Manager User's Guide*.

18. Bring the storage array back online.

For instructions, see the *Sun StorEdge RAID Manager User's Guide*.

19. Rebalance all logical units (LUNs).

For instructions, see the *Sun StorEdge RAID Manager User's Guide*.

20. Boot Node 1 into cluster mode.

Index

A

adding
 See also installing
 disk drives, 34
 storage arrays, 25
arrays, *See* storage arrays

B

battery units, replacing, 25

C

cables, replacing, 25
controllers
 replacing, 33
 restoring, 33
cooling canisters, replacing, 25
correcting, DID numbers, 22
creating
 hot spares, 17
 LUNs, 17

D

deleting
 hot spares, 17
 LUNs, 19
DID numbers, correcting, 22
disk drive firmware, upgrading, 36

disk drives
 adding, 34
 removing, 36
 replacing, 35
drive groups, expanding, 17

E

expanding, drive groups, 17

H

help, 8
host adapters, replacing, 36
hot spares
 creating, 17
 deleting, 17

I

installing
 See also adding
 storage arrays, 12

L

- logical unit, *See* LUNs
- logical unit number, *See* LUNs
- logical volumes, *See* LUNs
- LUN configurations, resetting, 20
- LUNs
 - creating, 17
 - deleting, 19
 - mismatched DID numbers, 22

P

- power cables, replacing, 25
- power supply, replacing, 25

R

- removing
 - See also* deleting
 - disk drives, 36
 - storage arrays, 31
- replacing
 - battery units, 25
 - cables, 25
 - controllers, 33
 - cooling canisters, 25
 - disk drives, 35
 - host adapters, 36
 - power cables, 25
 - power supply, 25
- resetting, LUN configurations, 20
- restoring, controllers, 33

S

- storage arrays
 - adding, 25
 - installing, 12
 - removing, 31

T

- technical support, 8

U

- upgrading, disk drive firmware, 36