



Sun Cluster 3.x With Sun StorEdge A3500/A3500FC System Manual

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

Part No: 817-0174-10
April 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.



030313@5533



Contents

Preface	5
1 Installing and Maintaining a Sun StorEdge A3500 System or a Sun StorEdge A3500FC System	11
Installing Storage Systems	12
▼ How to Install a Storage System	12
Configuring Storage Systems	18
▼ How to Create a LUN	19
▼ How to Delete a LUN	21
▼ How to Reset the LUN Configuration	23
▼ How to Correct Mismatched DID Numbers	25
Maintaining Storage Systems	26
▼ How to Add a Storage System	30
▼ How to Remove a Storage System	37
▼ How to Replace a Failed Controller or Restore an Offline Controller	39
▼ How to Upgrade Controller Module Firmware in a Running Cluster	41
▼ How to Add a Disk Drive in a Running Cluster	42
▼ How to Replace a Failed Disk Drive in a Running Cluster	43
▼ How to Remove a Disk Drive From a Running Cluster	44
How to Upgrade Disk Drive Firmware in a Running Cluster	44
▼ How to Replace a Host Adapter in a Node (StorEdge A3500 Storage System Only)	44
▼ How to Replace a Host Adapter in a Node (StorEdge A3500FC Storage System Only)	46
SAN Considerations	47
SAN Clustering Considerations	48

Preface

The *Sun Cluster 3.x With Sun StorEdge A3500/A3500FC System Manual* provides procedures specific to Sun StorEdge™ A3500/A3500FC systems that are placed in a Sun™ Cluster environment.

Who Should Use This Book

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

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

How This Book Is Organized

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

Section 1 discusses how to install Sun StorEdge A3500/A3500FC storage system.

Section 2 discusses how to configure logical units on Sun StorEdge A3500/A3500FC storage system.

Section 3 describes how to maintain Sun StorEdge A3500/A3500FC storage system 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 A3500/A3500FC System Manual

Revision Date	Information Added
	No revision.

Related Documentation

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

The following Sun Cluster books support 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>OpenBoot 2.x Command Reference Manual</i>	806-2906
<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 A3500/A3500FC Hardware Configuration Guide</i>	805-4981
<i>Sun StorEdge D1000 Storage Guide</i>	805-4013
<i>Sun StorEdge RAID Manager Installation and Support Guide</i>	805-7756
<i>Sun StorEdge RAID Manager Release Notes</i>	805-7758

TABLE P-2 Hardware Documentation (Continued)

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

Using UNIX Commands

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

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

- Online documentation for the Solaris software environment

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

Getting Help

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

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

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

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

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

Accessing Sun Documentation Online

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

Typographic Conventions

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

TABLE P-4 Typographic Conventions

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

Shell Prompts in Command Examples

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

TABLE P-5 Shell Prompts

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

Installing and Maintaining a Sun StorEdge A3500 System or a Sun StorEdge A3500FC System

This chapter describes the procedures about how to install, configure, and maintain a SunStorEdge™ A3500 (SCSI-based) system and a Sun StorEdge A3500FC (Fibre Channel-based) system in a Sun™ Cluster environment.

This chapter contains the following procedures.

- “Installing Storage Systems” on page 12
- “How to Create a LUN” on page 19
- “How to Delete a LUN” on page 21
- “How to Reset the LUN Configuration” on page 23
- “How to Correct Mismatched DID Numbers” on page 25
- “How to Add a Storage System” on page 30
- “How to Remove a Storage System” on page 37
- “How to Replace a Failed Controller or Restore an Offline Controller” on page 39
- “How to Upgrade Controller Module Firmware in a Running Cluster” on page 41
- “How to Add a Disk Drive in a Running Cluster” on page 42
- “How to Replace a Failed Disk Drive in a Running Cluster” on page 43
- “How to Remove a Disk Drive From a Running Cluster” on page 44
- “How to Upgrade Disk Drive Firmware in a Running Cluster” on page 44
- “How to Replace a Host Adapter in a Node (StorEdge A3500 Storage System Only)” on page 44
- “How to Replace a Host Adapter in a Node (StorEdge A3500FC Storage System Only)” on page 46

For information about how to use storage arrays in a storage area network (SAN), see “SAN Considerations” on page 47.

Installing Storage Systems

This section describes the procedure for an *initial installation* of a storage system.

▼ How to Install a Storage System

Use this procedure for an initial installation and an initial configuration, *before* you install the Solaris operating environment and Sun Cluster software.

This procedure uses an updated method for setting the `scsi-initiator-id` on the SCSI version of the storage system. For this storage system, 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 `nvrामrc` script and set the `scsi-initiator-id` by using this procedure.

1. Install the host adapters in the nodes. These nodes are to be connected to the storage system.

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

2. Cable the storage system.

- See Figure 1-1 for an example about how to cable a StorEdge A3500 system.
- See Figure 1-2 for an example about how to cable a StorEdge A3500FC system.

For more sample configurations, see the *Sun StorEdge A3500/A3500FC Hardware Configuration Guide*.

For the procedure about how to install the cables, see the *Sun StorEdge A3500/A3500FC Controller Module Guide*.

Note – If you want to create a SAN by using two FC switches and Sun SAN software, see “SAN Considerations” on page 47 for more information. Sun SAN software does not support storage array.

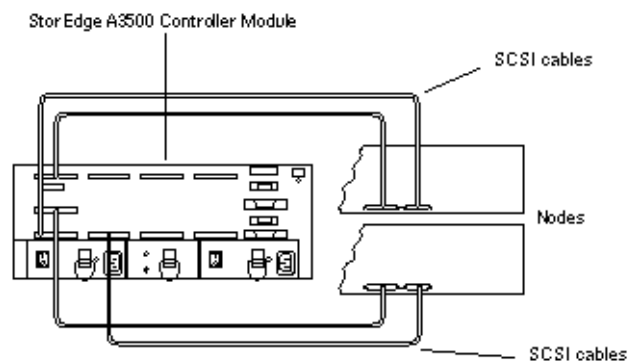


FIGURE 1-1 Installing a StorEdge A3500 Storage System: An Example




FIGURE 1-2 Installing a StorEdge A3500FC Storage System: An Example

3. Depending on which type of controller module you are installing.

- If you are installing a StorEdge A3500 controller module, proceed to Step 4.
- If you are installing a StorEdge A3500FC controller module, set the loop ID of the controller module. Set the loop ID by installing jumpers to the appropriate pins on the rear of the controller module.

For diagrams and information about how to set FC-AL ID settings, see the *Sun StorEdge A3500/A3500FC Controller Module Guide*.

4. Power on the storage system and nodes.

Note – For StorEdge A3500 controller modules only: When you power on the nodes, do not enable the nodes to boot. If necessary, halt the nodes so that you can perform OpenBoot™ PROM (OBP) Monitor tasks at the `ok` prompt.

For the procedure about how to power on the storage system, see the *Sun StorEdge A3500/A3500FC Controller Module Guide*.

5. Which type of controller module you are installing?

- For a StorEdge A3500FC controller module, skip to Step 10.
- For a StorEdge A3500 controller module, proceed to Step 6.

6. 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 `nvrarc` 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-nvramrc? true
use-nvramrc? = true
{1} ok
```

7. 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
...
```

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

For the procedure about how to install the Solaris operating environment, see your Sun Cluster software installation documentation. For the location of patches and installation instructions, see your Sun Cluster release notes documentation.

9. Read the following two conditions carefully to determine whether you must reboot the nodes.

- If you are using a version of RAID Manager later than 6.22, proceed to Step 10. If you are using a version of the Solaris operating environment earlier than Solaris 8 Update 4, proceed to Step 10.
- If you are using RAID Manager 6.22 *and* the Solaris 8 Update 4 or later operating environment, reboot both nodes.

```
# reboot
```

10. Install the RAID Manager software.

For the procedure about how to install the 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 software.

Note – For the most current list of software, firmware, and patches that are required for the storage system, refer to *EarlyNotifier 20029, A1000/A3x00/A3500FC 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>.

11. Install the storage system controller module and RAID Manager patches.

For more information, see the *Sun StorEdge RAID Manager Release Notes*.

12. Check the storage system controller module NVSRAM file revision. If necessary, install the most recent revision.

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

13. Check the storage system C controller module firmware revision. If necessary, install the most recent revision.

For the firmware revision number and boot level, see the *Sun StorEdge RAID Manager Release Notes*. For the procedure about how to upgrade the firmware, see the *Sun StorEdge RAID Manager User's Guide*.

14. Set the Rdac parameters in the `/etc/osa/rmparams` file.

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

15. Ensure that the controller module is set to `active/active` mode.

For more information on controller modes, see the *Sun StorEdge RAID Manager Installation and Support Guide* and the *Sun StorEdge RAID Manager User's Guide*.

16. Set up the storage system controller module with logical unit numbers (LUNs) and hot spares.

For the procedure about how to set up the storage system controller module 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.

17. Copy the `/etc/raid/rdac_address` file from the node on which you created the LUNs to the other node. If you copy this file to the other node, you ensure consistency across both nodes.

18. Ensure that the new logical name for the LUN that you created in Step 16 appears in the `/dev/rdisk` 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 Sun Cluster data services collection.

Configuring Storage Systems

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

To configure a storage system *before* you install Sun Cluster software, use the same procedures you use in a noncluster environment. For the procedures about how to configure a storage system before you install Sun Cluster software, 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's device ID (DID) might change unexpectedly. When you perform a check of the DID configuration by running the `sddidadm -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.
```

Run the `sddidadm -R` command for each affected device.

TABLE 1-1 Task Map: Configuring Disk Drives

Task	Information
Create a logical unit number (LUN).	"How to Create a LUN" on page 19
Remove a LUN.	"How to Delete a LUN" on page 21
Reset the LUN configuration.	"How to Reset the LUN Configuration" on page 23

TABLE 1-1 Task Map: Configuring Disk Drives (Continued)

Task	Information
Rebalance running LUNs. Follow the same procedure that you use in a noncluster environment.	<i>Sun StorEdge RAID Manager User's Guide</i> <i>Sun StorEdge RAID Manager Release Notes</i>
Create a hot spare. Follow the same procedure that you use in a noncluster environment.	<i>Sun StorEdge RAID Manager User's Guide</i> <i>Sun StorEdge RAID Manager Release Notes</i>
Delete a hot spare. Follow the same procedure that you use in a noncluster environment.	<i>Sun StorEdge RAID Manager User's Guide</i> <i>Sun StorEdge RAID Manager Release Notes</i>
Increase the size of a drive group. Follow the same procedure that you use in a noncluster environment.	<i>Sun StorEdge RAID Manager User's Guide</i> <i>Sun StorEdge RAID Manager Release Notes</i>

▼ How to Create a LUN

Use this procedure to create a logical unit number (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 – You must use hardware RAID for Oracle Parallel Server (OPS) data stored on the storage array. Do not place OPS data under volume management control. You must place all non-OPS data that is stored on the storage arrays under volume management control. Use either hardware RAID, host-based software RAID, or both types of RAID to manage your non-OPS data.

Hardware RAID uses the StorEdge A3500 system's or the StorEdge A3500FC system's hardware redundancy to ensure that independent hardware failures do not impact data availability. If you mirror across separate storage arrays, host-based software RAID ensures that independent hardware failures do not impact data availability when an entire storage array is offline. Although you can use hardware RAID and host-based software RAID concurrently, you need only one RAID solution to maintain a high degree of data availability.

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 Systems” on page 26 without volume management disruptions. If you use hardware RAID level 0, some maintenance procedures in “Maintaining Storage Systems” on page 26 require additional volume management administration because the availability of the LUNs is impacted.

1. **With all nodes booted and attached to the storage system, 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 system.

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@40,0/SUNW,socal@0,0/sf@1,0/ssd@w200200a0b80740db,4 (ssd0) :
corrupt label - wrong magic number
```

Note – Use the `format(1M)` command to verify Solaris logical device names.

2. **Copy the `/etc/raid/rdac_address` file from the node on which you created the LUN to the other node. If you copy this file to the other node, you ensure consistency across both nodes.**
3. **Ensure that the new logical name for the LUN that you created appears in the `/dev/rdisk` directory on both nodes.**

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

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

```
# scgdevs
```

5. **Ensure that the DID numbers for the LUNs are the same on both nodes. In the sample output that follows, the DID numbers are different.**

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

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

- If the DID numbers are the same, proceed to Step 7.
- If the DID numbers are different, perform the procedure in “How to Correct Mismatched DID Numbers” on page 25 before you continue with Step 7 of this procedure.

7. If you want a volume manager to manage the new LUN, incorporate the new LUN into a diskset or disk group.

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



Caution – Do not configure LUNs as quorum devices. The use of LUNs as quorum devices is not supported.

▼ 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 system, determine the paths to the LUN that you are deleting.

```
# format
AVAILABLE DISK SELECTIONS:
0. c0t5d0 <SYMBIOS-StorEdgeA3500FCr-0301 cyl13 alt2 hd64 sec64>
   /pseudo/rdnexus@0/rdriver@5,0
1. c0t5d1 <SYMBIOS-StorEdgeA3500FCr-0301 cyl12025 alt2 hd64 sec64>
   /pseudo/rdnexus@0/rdriver@5,1
```

2. Does a volume manager manage the LUN that you are deleting?

- If no, proceed to Step 3.

- If yes, 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 LUN from any disk group, use the following commands.

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

3. From one node, delete the LUN.

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

4. 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*
```

5. Use the `lad` command to determine the *alternate* paths to the LUNs you are deleting.

The RAID Manager software creates two paths to the LUN in the `/dev/osa/dev/rdisk` directory. Substitute the `cNtXdY` number from the other controller module in the 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:

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

6. Remove the *alternate* paths to the LUNs you are deleting.

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

7. On both nodes, remove all obsolete device IDs (DID)s.

```
# scdidadm -C
```

8. Switch resources and device groups off the node.

```
# scswitch -Sh nodename
```

9. Shut down the node.

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

10. Perform a reconfiguration boot to create the new Solaris device files and links.
11. Repeat Step 3 through Step 10 on the other node that is attached to the storage system.

▼ How to Reset the LUN Configuration

Use this procedure to reset the LUN configuration.



Caution – If you reset a 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 system, determine the paths to the LUNs you are resetting.

```
# format
AVAILABLE DISK SELECTIONS:
0. c0t5d0 <SYMBIOS-StorEdgeA3500FCr-0301 cyl13 alt2 hd64 sec64>
/pseudo/rdnexus@0/rdriver@5,0
1. c0t5d1 <SYMBIOS-StorEdgeA3500FCr-0301 cyl12025 alt2 hd64 sec64>
/pseudo/rdnexus@0/rdriver@5,1
```

2. Does a volume manager manage the LUNs on the controller module you are resetting?

- If no, proceed to Step 3.
- If yes, 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
```

3. On one node, reset the LUN configuration.

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

4. Set the controller module back to active/active mode.

For more information on controller modes, see the *Sun StorEdge RAID Manager Installation and Support Guide* and the *Sun StorEdge RAID Manager User's Guide*.

5. By using the `format` command, label the new LUN 0.
6. Remove the paths to the old LUNs you reset.

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

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

7. Use the `lad` command to determine the *alternate* paths to the old LUNs you reset.

The RAID Manager software creates two paths to the LUN in the `/dev/osa/dev/rdisk` directory. Substitute the `cNtXdY` number from the other controller module in the 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*
```

8. Remove the *alternate* paths to the old LUNs you reset.

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

9. Update device namespaces on both nodes.

```
devfsadm -C
```

10. Remove all obsolete DID's on both nodes.

```
# scdidadm -C
```

11. Switch resources and device groups off the node.

```
# scswitch -Sh nodename
```

12. Shut down the node.

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

13. 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.
```

14. After the node reboots and joins the cluster, repeat Step 6 through Step 13 on the other node. This node is attached to the storage system.

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 A3500/A3500FC 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 19.

1. From one node that is connected to the storage system, 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 controller module in the 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. Switch resources and device groups off the node.

```
# scswitch -Sh nodename
```

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. This node is attached to the storage system.**
10. **Return to “How to Create a LUN” on page 19.**

Maintaining Storage Systems

This section contains the procedures about how to maintain a storage system 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 a storage system.

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’s device ID (DID) might change unexpectedly. When you perform a check of the DID configuration by running the `sddidadm -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.
```

Run the `sddidadm -R` command for each affected device.

TABLE 1–2 Tasks: Maintaining a Storage System

Task	Information
Storage system procedures:	
Add a storage system.	“How to Add a Storage System” on page 30
Remove a storage system.	“How to Remove a Storage System” on page 37

TABLE 1-2 Tasks: Maintaining a Storage System (Continued)

Task	Information
Replace a failed controller module. Restore an offline controller module.	"How to Replace a Failed Controller or Restore an Offline Controller" on page 39
Upgrade controller module firmware and NVSRAM file.	"How to Upgrade Controller Module Firmware in a Running Cluster" on page 41
Replace a controller module power cord. Shut down the cluster, then follow the same procedure that you use in a noncluster environment.	Sun Cluster system administration documentation for procedures on shutting down a cluster <i>Sun StorEdge A3500/A3500FC Controller Module Guide</i> for replacement procedures
Replace a controller module fan canister. Follow the same procedure that you use in a noncluster environment.	<i>Sun StorEdge A3500/A3500FC Controller Module Guide</i>
Replace a controller module card cage. Shut down the cluster, then follow the same procedure that you use in a noncluster environment.	Sun Cluster system administration documentation for procedures on shutting down a cluster <i>Sun StorEdge A3500/A3500FC Controller Module Guide</i> for replacement procedures
Replace a controller module assembly. Shut down the cluster, then follow the same procedure that you use in a noncluster environment.	Sun Cluster system administration documentation for procedures on shutting down a cluster <i>Sun StorEdge A3500/A3500FC Controller Module Guide</i> for replacement procedures
Cable, hub, and connector procedures:	
Replace a SCSI cable from the controller module to the storage array. Follow the same procedure that you use in a noncluster environment. You might encounter I/O errors when you replace this cable. These errors are temporary. These errors disappear when the new cable is securely in place. Use your volume management recovery procedure to recover from these I/O errors.	<i>Sun StorEdge A3500/A3500FC Controller Module Guide</i> <i>Sun StorEdge RAID Manager User's Guide</i> <i>Sun StorEdge RAID Manager Release Notes</i>

TABLE 1–2 Tasks: Maintaining a Storage System (Continued)

Task	Information
<p>Replace a storage array-to-host SCSI cable.</p> <p>Follow the same procedure that you use in a noncluster environment.</p>	<p><i>Sun StorEdge A3500/A3500FC Controller Module Guide</i></p> <p><i>Sun StorEdge RAID Manager User's Guide</i></p> <p><i>Sun StorEdge RAID Manager Release Notes</i></p>
<p>Replace a storage array-to-host or storage array-to-hub fiber-optic cable.</p> <p>Follow the same procedure that you use in a noncluster environment.</p>	<p><i>Sun StorEdge A3500/A3500FC Controller Module Guide</i></p> <p><i>Sun StorEdge RAID Manager User's Guide</i></p> <p><i>Sun StorEdge RAID Manager Release Notes</i></p>
<p>Replace an FC hub.</p> <p>Follow the same procedure that you use in a noncluster environment.</p>	<p><i>Sun StorEdge A3500/A3500FC Controller Module Guide</i></p> <p><i>Sun StorEdge FC-100 Hub Installation and Service Manual</i></p>
<p>Replace an FC switch.</p> <p>This procedure applies to SAN-configured clusters only.</p>	<p>See "SAN Considerations" on page 47 for SAN information.</p>
<p>Replace an FC hub gigabit interface converter (GBIC) or Small Form Factor Pluggable (SFP) that connects cables to the host or hub.</p> <p>Follow the same procedure that you use in a noncluster environment.</p>	<p><i>Sun StorEdge A3500/A3500FC Controller Module Guide</i></p> <p>See the documentation that shipped with your FC hub or FC switch.</p>
<p>Replace a GBIC or an SFP on a node.</p> <p>Follow the same procedure that you use in a noncluster environment.</p>	<p><i>Sun StorEdge A3500/A3500FC Controller Module Guide</i></p>
Cabinet and power subassembly procedures:	
<p>Replace the power supply fan canister.</p>	<p><i>Sun StorEdge A3500/A3500FC Controller Module Guide</i></p>
<p>Replace a DC-power or battery harness.</p> <p>Shut down the cluster, then follow the same procedure that you use in a noncluster environment.</p>	<p>Sun Cluster system administration documentation for procedures on shutting down a cluster</p> <p><i>Sun StorEdge A3500/A3500FC Controller Module Guide</i> for replacement procedures</p>

TABLE 1-2 Tasks: Maintaining a Storage System (Continued)

Task	Information
Replace the battery unit. Shut down the cluster, then follow the same procedure that you use in a noncluster environment.	Sun Cluster system administration documentation for procedures on shutting down a cluster <i>Sun StorEdge A3500/A3500FC Controller Module Guide</i> for replacement procedures
Replace the power supply housing. Shut down the cluster, then follow the same procedure that you use in a noncluster environment.	Sun Cluster system administration documentation for procedures on shutting down a cluster <i>Sun StorEdge A3500/A3500FC Controller Module Guide</i>
Replace a power cord to the cabinet power distribution unit. Shut down the cluster, then follow the same procedure that you use in a noncluster environment.	Sun Cluster system administration documentation for procedures on shutting down a cluster <i>Sun StorEdge Expansion Cabinet Installation and Service Manual</i> for replacement procedures.
Storage array procedures:	
Add a disk drive.	"How to Add a Disk Drive in a Running Cluster" on page 42
Replace a disk drive.	"How to Replace a Failed Disk Drive in a Running Cluster" on page 43
Remove a disk drive.	"How to Remove a Disk Drive From a Running Cluster" on page 44
Upgrade disk drive firmware.	"How to Upgrade Disk Drive Firmware in a Running Cluster" on page 44
Replace a power cord to a storage array. Shut down the cluster, then follow the same procedure that you use in a noncluster environment.	Sun Cluster system administration documentation for procedures on shutting down a cluster <i>Sun StorEdge A1000 and D1000 Installation, Operations, and Service Manual</i> for replacement procedures
Node and host adapter procedures:	

TABLE 1-2 Tasks: Maintaining a Storage System (Continued)

Task	Information
Replace a host adapter in a node.	"How to Replace a Host Adapter in a Node (StorEdge A3500 Storage System Only)" on page 44 "How to Replace a Host Adapter in a Node (StorEdge A3500FC Storage System Only)" on page 46

▼ How to Add a Storage System

Use this procedure to add a storage system to a running cluster.

This procedure uses an updated method for setting the `scsi-initiator-id`. For this storage array, 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.

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 software.

Note – For the most current list of software, firmware, and patches that are required for the storage system, refer to *EarlyNotifier 20029, A1000/A3x00/A3500FC 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>.

2. Install the storage system patches.

For the location of patches and installation instructions, see your Sun Cluster release notes documentation.

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

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

4. Power on the storage system.

For the procedure about how to power on the storage system, see the *Sun StorEdge A3500/A3500FC Controller Module Guide*.

5. Which type of system you are adding?

- If you are adding a StorEdge A3500 system, proceed to Step 6.
- If you are adding a StorEdge A3500FC system, set the loop ID of the controller module. Set the loop ID by installing jumpers to the appropriate pins on the rear of the controller module.

For diagrams and information about how to set FC-AL ID settings, see the *Sun StorEdge A3500/A3500FC Controller Module Guide*.

6. Are you installing new host adapters in your nodes?

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

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 system to Node 1. The method you choose to cable your system depends on the type of system you are adding.

- If you are adding a StorEdge A3500 system, connect the differential SCSI cable between the node and the controller module, as shown in Figure 1–3. Make sure that the *entire* SCSI bus length to each enclosure 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, and host adapter.
- If you are installing a StorEdge A3500FC storage system, see Figure 1–4 for an example about how to cable a StorEdge A3500FC storage system. In this example, Node 1 connects a StorEdge A3500FC controller module.

For more sample configurations, see the *Sun StorEdge A3500/A3500FC Hardware Configuration Guide*.

For the procedure about how to install the cables, see the *Sun StorEdge A3500/A3500FC Controller Module Guide*.

Note – If you want to create a SAN by using two FC switches and Sun SAN software, see “SAN Considerations” on page 47 for more information. Sun SAN software does not support storage array.




FIGURE 1-3 Sample Storage System Cabling

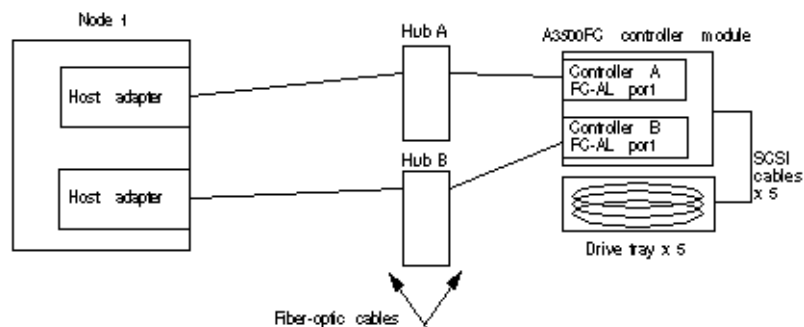


FIGURE 1-4 Sample Storage System Cabling (1st Node Attached)

9. Did you power off the Node 1 to install a host adapter?

- If no, proceed to Step 10.
- If yes, power on Node 1 and the storage system. *Do not enable the node to boot.* If necessary, halt the system to continue with OpenBoot PROM (OBP) Monitor tasks.

10. Which type of controller module are you adding?

- If a StorEdge A3500FC controller module, skip to Step 12.
- If a StorEdge A3500 controller module, proceed to Step 11.

11. 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 `nvrnrc` 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 nvramrc script that you created, as shown in the following example.**

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

```
{0} ok printenv nvramrc
nvramrc =
    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
```

12. Did you power off the Node 1 to install a host adapter?

- If no, skip to Step 18.
- If yes, perform a reconfiguration boot to create the new Solaris device files and links.

13. Are you installing new host adapters in Node 2 to connect Node 2 to the storage system?

- If no, skip to Step 18.
- If yes, shut down and power off the node.

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 system to your node. The method that you choose to cable your storage system depends on the type of controller module you are adding.

- If you are adding a StorEdge A3500 controller module, connect the differential SCSI cable between the node and the controller module, as shown in Figure 1–3. Make sure that the *entire* SCSI bus length to each enclosure 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, and host adapter.
- If you are installing a StorEdge A3500FC controller module, see Figure 1–5 for an example about how to cable a StorEdge A3500FC storage system. In this example, two nodes connect to a StorEdge A3500FC controller module.

For more sample configurations, see the *Sun StorEdge A3500/A3500FC Hardware Configuration Guide*.

For the procedure about how to install the cables, see the *Sun StorEdge A3500/A3500FC Controller Module Guide*.

Note – If you want to create a SAN by using two FC switches and Sun SAN software, see “SAN Considerations” on page 47 for more information. Sun SAN software does not support storage array.




FIGURE 1–5 Cabling a Storage System: An Example

16. Did you power off Node 2 to install a host adapter?

- If no, skip to Step 18.

- If yes, power on the node and the storage system. *Do not enable the node to boot.* If necessary, halt the system to continue with OpenBoot PROM (OBP) Monitor tasks.

17. Verify that Node 2 recognizes the new host adapters and disk drives.

If the node does not recognize the new hardware, check all hardware connections and repeat installation steps you performed in Step 14.

```
{0} ok show-disks
...
b) /sbus@6,0/QLGC,isp@2,10000/sd...
d) /sbus@2,0/QLGC,isp@2,10000/sd...{0} ok
```

18. Which type of controller module are you adding?

- If a StorEdge A3500FC controller module, skip to Step 24.
- If a StorEdge A3500 controller module, Step 19.

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

- 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.**
- 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
...
```

20. Did you power off Node 2 to install a host adapter?

- If no, skip to Step 24.
- If yes, perform a reconfiguration boot to create the new Solaris device files and links.

21. Check the controller module NVSRAM file revision. If necessary, install the most recent revision.

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

22. Verify the controller module firmware revision. If necessary, install the most recent firmware revision.

For the revision number and boot level of the controller module firmware, see the *Sun StorEdge RAID Manager Release Notes*. For the procedure about how to upgrade the controller firmware, see "How to Upgrade Controller Module Firmware in a Running

Cluster” on page 41.

23. One node at a time, boot each node into cluster mode.

```
# reboot
```

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

```
# scdidadm -L
```

25. Verify that the controller module is set to active/active mode.

For more information on controller modes, see the *Sun StorEdge RAID Manager Installation and Support User’s Guide* and the *Sun StorEdge RAID Manager User’s Guide*.

Where to Go From Here

To create a LUN from disk drives that are unassigned, see “How to Create a LUN” on page 19.

To upgrade controller module firmware, see “How to Upgrade Controller Module Firmware in a Running Cluster” on page 41.

▼ How to Remove a Storage System

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



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

1. Migrate any Oracle Parallel Server (OPS) tables, data services, or volumes off the storage system.

2. Halt all activity to the controller module.

For instructions, see the *Sun StorEdge RAID Manager User’s Guide* and your operating system documentation.

3. Does a volume manager manage any of the LUNs on the controller module you are removing?

- If no, proceed to Step 4.
- If yes, 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. **Disconnect all cables from the storage system. Remove the hardware from your cluster.**

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. **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*
```

7. **Use the `lad` command to determine the *alternate* paths to the LUNs you are deleting.**

The RAID Manager software creates two paths to the LUN in the `/dev/osa/dev/rdisk` directory. Substitute the `cNtXdY` number from the other controller module in the 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 the following.

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

8. **Remove the *alternate* paths to the LUNs you are deleting.**

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

9. **On all nodes, remove references to the storage system.**

```
# scdidadm -C
```

10. **Are you are removing the *last* StorEdge A3500FC controller module from a hub or FC switch in your cluster?**

- If no, proceed to Step 12.
- If yes, remove the hub or FC switch hardware and cables from your cluster.

Note – If you are using your StorEdge A3500FC storage array in a SAN-configured cluster, you must keep two FC switches configured in parallel. This configuration maintains cluster availability. See “SAN Considerations” on page 47 for more information.

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

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

12. Remove any unused host adapter from nodes that were attached to the storage system.
 - a. Shut down and power off Node 1 from which you are removing a host adapter.
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 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 system.
13. Switch the cluster back online.

```
# scswitch -Z
```
14. Are you removing the *last* the storage system from your cluster?
 - If no, you are finished with this procedure.
 - If yes, remove software packages. For the procedure about how to remove software packages, see the documentation that shipped with your storage system.

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

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

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

Note – If you want to create a SAN by using two FC switches and Sun SAN software, see “SAN Considerations” on page 47 for more information. Sun SAN software does not support storage array.

1. On both nodes, to prevent LUNs from automatic assignment to the controller that is being brought online, set the `System_LunReDistribution` parameter in the `/etc/raid/rmparams` file to `false`.



Caution – You must set the `System_LunReDistribution` parameter in the `/etc/raid/rmparams` file to `false` so that no LUNs are assigned to the controller being brought online. After you verify in Step 5 that the controller has the correct SCSI reservation state, you can balance LUNs between both controllers.

For the procedure about how to modify the `rmparams` file, see the *Sun StorEdge RAID Manager Installation and Support Guide*.

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 controller module 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 controllers, see the *Sun StorEdge A3500/A3500FC Controller Module Guide* and the *Sun StorEdge RAID Manager Installation and Support Guide* for additional considerations.

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. Do not use the Redundant Disk Array Controller Utility (`rdacutil`) because this utility ignores the value of the `System_LunReDistribution` parameter in the `/etc/raid/rmparams` file.

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 Installation and Support Guide*.

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

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

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


6. Set the controller to active/active mode. Assign LUNs to the controller.

For more information on controller modes, see the *Sun StorEdge RAID Manager Installation and Support Guide* and the *Sun StorEdge RAID Manager User's Guide*.

7. Reset the `System_LunReDistribution` parameter in the `/etc/raid/rmparams` file to `true`.

For the procedure about how to change the `rmparams` file, see the *Sun StorEdge RAID Manager Installation and Support Guide*.

8. Restart the RAID Manager daemon.

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

▼ How to Upgrade Controller Module Firmware in a Running Cluster

Use this procedure to upgrade firmware in a controller module, in a running cluster. Use either the *online* or the *offline* method to upgrade your firmware. The method that you choose depends on your firmware.

1. Are you upgrading the NVSRAM firmware file?

- If you are not upgrading the NVSRAM file, you can use the online method. Upgrade the firmware by using the online method, as described in the *Sun StorEdge RAID Manager User's Guide*. No special steps are required for a cluster environment.
- If you are upgrading the NVSRAM file, you must use an offline method. Use one of the following procedures.
 - If the data on your controller module is mirrored on another controller module, use the procedure in Step 2.
 - If the data on your controller module is *not* mirrored on another controller module, use the procedure in Step 3.

2. Use this step if you are upgrading the NVSRAM and other firmware files on a controller module. This controller module must not have mirrored data.

- a. **Halt all activity to the controller module.**

For more information, see your Solstice DiskSuite/Solaris Volume Manager or VERITAS Volume Manager documentation.
- b. **Update the firmware files by using the offline method, as described in the *RAID Manager User's Guide*.**
- c. **Restore all activity to the controller module.**

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

This step completes the firmware upgrade.

3. Use this step if you are upgrading the NVSRAM and other firmware files on a controller module. This controller module must not have mirrored data.

a. Shut down the entire cluster.

For the procedure about how to shut down a cluster, see your Sun Cluster system administration documentation.

b. Boot one node that is attached to the controller module into noncluster mode.

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

c. Update the firmware files using the offline method, as described in the *RAID Manager User's Guide*.

d. Boot both nodes into cluster mode.

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

This step completes the firmware upgrade.

▼ How to Add a Disk Drive in a Running Cluster

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



Caution – If the disk drive that you are adding was previously owned by another controller module, reformat the disk drive. Reformat the disk drive to wipe clean the old DacStore information before adding the disk drive to this storage array.

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

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

2. Enable the disk drive to spin up approximately 30 seconds.

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

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

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

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

5. Repeat Step 1 through Step 4 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 19 for more information.

▼ How to Replace a Failed Disk Drive in a Running Cluster

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

1. **Does replacing the disk drive affects 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 the drive does *not* belong to a device group, proceed to Step 5.
 - If the drive is part of a device group, reconstruction is started 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 From a Running Cluster

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

1. **Remove the logical unit number (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 21.

2. **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 in a Running Cluster



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 Host Adapter in a Node (StorEdge A3500 Storage System Only)

Note – This section describes the procedure about how to replace a failed host adapter in a running node. This node is attached to a *SCSI-based* storage system. For the same procedure for a node that is attached to a *fiber-optic-based* storage system, see “How to Replace a Host Adapter in a Node (StorEdge A3500FC Storage System Only)” on page 46.

This procedure defines Node 1 as the node with the host adapter on SCSI bus A. This host adapter is the host adapter that you are replacing. Node 2 is the node that remains in service.

Note – Several steps in this procedure require you to halt I/O activity. To halt I/O activity, take the controller module 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 controller module end of the SCSI cable, disconnect the SCSI bus A cable. This cable connects the controller module to Node 1. Afterward, 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 SCSI bus B is not affected, skip to Step 9.
- If SCSI bus B is affected, proceed to Step 6.

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

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

7. From the controller module end of the SCSI cable, disconnect the SCSI bus B cable. This cable connects the controller module to Node 1. Afterward, 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 system.

12. From Node 2, halt I/O activity to the controller module 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. Afterward, reinstall the SCSI cable to connect the controller module 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. Afterward, reinstall the SCSI cable to connect the controller module 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 controller module online.**
For instructions, see the *Sun StorEdge RAID Manager User's Guide*.
19. **Rebalance all logical unit numbers (LUNs).**
For instructions, see the *Sun StorEdge RAID Manager User's Guide*.
20. **Boot node 1 into cluster mode.**

▼ How to Replace a Host Adapter in a Node (StorEdge A3500FC Storage System Only)

Note – This section describes the procedure about how to replace a failed host adapter in a node that is attached to a *fiber-optic-based* storage system. For the same procedure for a node that is attached to a *SCSI-based* storage system, see “How to Replace a Host Adapter in a Node (StorEdge A3500 Storage System Only)” on page 44.

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

1. **Move all Sun Cluster data services off of the node in which you are replacing a host adapter.**
For instructions, see the *Sun StorEdge RAID Manager User's Guide*.
2. **Halt all I/O activity on the affected controller module.**
For instructions, see the *Sun StorEdge RAID Manager User's Guide*.

- 3. Shut down and power off the node in which you are replacing a host adapter.**
For the procedure about how to shut down and power off a node, see your Sun Cluster system administration documentation.
- 4. Disconnect the fiber-optic cable from the host adapter that you are replacing.**
- 5. Replace the host adapter in the node.**
For the procedure about how to replace a host adapter, see the documentation that shipped with your node hardware.
- 6. Connect the fiber-optic cable to the new host adapter that you just installed.**
- 7. Boot the node into cluster mode.**
For more information about how to boot nodes, see your Sun Cluster system administration documentation.
- 8. Restart I/O activity on the affected controller module.**
For instructions, see the *Sun StorEdge RAID Manager User's Guide* and your operating system documentation.
- 9. Check the status of the affected controller module.**
Observe the front panel LEDs. Use *Sun StorEdge A3500/A3500FC Controller Module Guide* to interpret the LEDs.
- 10. Rebalance LUNs that are running on the affected controller module.**
For more information, see the *Sun StorEdge RAID Manager User's Guide*.
- 11. Return the Sun Cluster data services to the node in which you replaced the host adapter.**
For instructions, see the Sun Cluster data services collection.

SAN Considerations

This section contains information about how to use storage array in a SAN. This information is specific to a SAN in a Sun Cluster environment. Use the cluster-specific procedures in this chapter to install and maintain a storage array in your cluster.

For instructions about how to create and maintain a SAN, see the documentation that shipped with your switch hardware. For information on switch ports, zoning, and required software and firmware, also see the documentation that shipped with your switch hardware.

SAN hardware includes the following components.

- FC switches

- FC host adapters
- Storage devices
- Enclosures

SAN software includes the following components.

- Drivers that are packaged with the operating system
- Firmware for the switches
- Management tools for the switches and storage devices
- Volume management software
- Other administration tools

SAN Clustering Considerations

If you are replacing an FC switch and you intend to save the switch IP configuration for restoration to the replacement switch, wait to connect the cables to the replacement switch. Connect the cables to the replacement switch *after* you recall the Fabric configuration to the replacement switch. For more information about how to save and recall switch configurations, see the documentation that shipped with your switch hardware.

Index

A

- adding
 - See also* installing
 - disk drives, 42
 - systems, 30
- arrays
 - See* storage arrays

B

- battery units, replacing, 29

C

- cables
 - replacing, 27
 - fiber-optic cables, 28
 - SCSI, 28
- controller module assemblies, replacing, 27
- controller module card cages, replacing, 27
- controller module fan canisters, replacing, 27
- controller module firmware, upgrading, 41
- controller module power cords, replacing, 27
- controllers
 - replacing, 39
 - restoring, 39
- correcting, DID numbers, 25
- creating
 - hot spares, 19
 - LUNs, 19

D

- deleting
 - See also* removing
 - hot spares, 19
 - LUNs, 21
- DID numbers, correcting, 25
- disk drive firmware, upgrading, 44
- disk drives
 - adding, 42
 - removing, 44
 - replacing, 43
- drive groups, expanding, 19

E

- expanding, drive groups, 19

F

- fan canisters
 - See* controller module fan canisters
 - See* power supply fan canisters

G

- GBICs
 - replacing, 28

gigabit interface converters
 See GBIC

H

help, 8
host adapters, replacing, 46
hot spares
 creating, 19
 deleting, 19
hubs, replacing, 28

I

installing
 See also adding
 storage systems, 12

L

logical unit numbers
 See LUNs
logical units, LUNs, 19
logical volumes
 See LUNs
LUN configurations, resetting, 23
LUNs
 creating, 19
 deleting, 21
 rebalancing, 19

P

power cords
 replacing, 27
 storage array
 replacing, 29
power supply fan canisters, replacing, 28
power supply housing, replacing, 29

R

RAID
 hardware, 19
 software, 19
rebalancing, LUNs, 19
removing
 See also deleting
 disk drives, 44
 storage systems, 37
replacing
 battery units, 29
 cables, 27
 controller module assemblies, 27
 controller module card cages, 27
 controller module fan canisters, 27
 controller module power cords, 27
 controllers, 39
 disk drives, 43
 fiber-optic cables, 28
 GBICs, 28
 host adapters, 46
 hubs, 28
 power supply fan canisters, 28
 power supply housing, 29
 SCSI cables, 28
 SFPs, 28
 storage array power cords, 29
 switches, 28
resetting, LUN configurations, 23
restoring, controllers, 39

S

SAN
 clustering considerations, 48
 described, 47
SFPs
 replacing, 28
Small Form Factor Plugable
 See SFP
storage area network
 See SAN
storage systems
 installing, 12
 removing, 37
switches, replacing, 28
systems, *See* storage systems, 11

systems, *See* storage systems (Continued)
adding, 30

T

technical support, 8

U

upgrading
controller module firmware, 41
disk drive firmware, 44

