

Oracle® Solaris Cluster 3.3 With StorEdge A1000 Array, Netra st A1000 Array, or StorEdge A3500 System Manual

For Solaris OS (SPARC Platform Edition)

License Restrictions Warranty/Consequential Damages Disclaimer

This software and related documentation are provided under a license agreement containing restrictions on use and disclosure and are protected by intellectual property laws. Except as expressly permitted in your license agreement or allowed by law, you may not use, copy, reproduce, translate, broadcast, modify, license, transmit, distribute, exhibit, perform, publish or display any part, in any form, or by any means. Reverse engineering, disassembly, or decompilation of this software, unless required by law for interoperability, is prohibited.

Warranty Disclaimer

The information contained herein is subject to change without notice and is not warranted to be error-free. If you find any errors, please report them to us in writing.

Restricted Rights Notice

If this is software or related documentation that is delivered to the U.S. Government or anyone licensing it on behalf of the U.S. Government, the following notice is applicable:

U.S. GOVERNMENT RIGHTS

Programs, software, databases, and related documentation and technical data delivered to U.S. Government customers are "commercial computer software" or "commercial technical data" pursuant to the applicable Federal Acquisition Regulation and agency-specific supplemental regulations. As such, the use, duplication, disclosure, modification, and adaptation shall be subject to the restrictions and license terms set forth in the applicable Government contract, and, to the extent applicable by the terms of the Government contract, the additional rights set forth in FAR 52.227-19, Commercial Computer Software License (December 2007). Oracle America, Inc., 500 Oracle Parkway, Redwood City, CA 94065.

Hazardous Applications Notice

This software or hardware is developed for general use in a variety of information management applications. It is not developed or intended for use in any inherently dangerous applications, including applications that may create a risk of personal injury. If you use this software or hardware in dangerous applications, then you shall be responsible to take all appropriate fail-safe, backup, redundancy, and other measures to ensure its safe use. Oracle Corporation and its affiliates disclaim any liability for any damages caused by use of this software or hardware in dangerous applications.

Trademark Notice

Oracle and Java are registered trademarks of Oracle and/or its affiliates. Other names may be trademarks of their respective owners.

Intel and Intel Xeon are trademarks or registered trademarks of Intel Corporation. All SPARC trademarks are used under license and are trademarks or registered trademarks of SPARC International, Inc. AMD, Opteron, the AMD logo, and the AMD Opteron logo are trademarks or registered trademarks of Advanced Micro Devices. UNIX is a registered trademark of The Open Group in the United States and other countries.

Third Party Content, Products, and Services Disclaimer

This software or hardware and documentation may provide access to or information on content, products, and services from third parties. Oracle Corporation and its affiliates are not responsible for and expressly disclaim all warranties of any kind with respect to third-party content, products, and services. Oracle Corporation and its affiliates will not be responsible for any loss, costs, or damages incurred due to your access to or use of third-party content, products, or services.

Ce logiciel et la documentation qui l'accompagne sont protégés par les lois sur la propriété intellectuelle. Ils sont concédés sous licence et soumis à des restrictions d'utilisation et de divulgation. Sauf disposition de votre contrat de licence ou de la loi, vous ne pouvez pas copier, reproduire, traduire, diffuser, modifier, breveter, transmettre, distribuer, exposer, exécuter, publier ou afficher le logiciel, même partiellement, sous quelque forme et par quelque procédé que ce soit. Par ailleurs, il est interdit de procéder à toute ingénierie inverse du logiciel, de le désassembler ou de le décompiler, excepté à des fins d'interopérabilité avec des logiciels tiers ou tel que prescrit par la loi.

Les informations fournies dans ce document sont susceptibles de modification sans préavis. Par ailleurs, Oracle Corporation ne garantit pas qu'elles soient exemptes d'erreurs et vous invite, le cas échéant, à lui en faire part par écrit.

Si ce logiciel, ou la documentation qui l'accompagne, est concédé sous licence au Gouvernement des Etats-Unis, ou à toute entité qui délivre la licence de ce logiciel ou l'utilise pour le compte du Gouvernement des Etats-Unis, la notice suivante s'applique :

U.S. GOVERNMENT RIGHTS. Programs, software, databases, and related documentation and technical data delivered to U.S. Government customers are "commercial computer software" or "commercial technical data" pursuant to the applicable Federal Acquisition Regulation and agency-specific supplemental regulations. As such, the use, duplication, disclosure, modification, and adaptation shall be subject to the restrictions and license terms set forth in the applicable Government contract, and, to the extent applicable by the terms of the Government contract, the additional rights set forth in FAR 52.227-19, Commercial Computer Software License (December 2007). Oracle America, Inc., 500 Oracle Parkway, Redwood City, CA 94065.

Ce logiciel ou matériel a été développé pour un usage général dans le cadre d'applications de gestion des informations. Ce logiciel ou matériel n'est pas conçu ni n'est destiné à être utilisé dans des applications à risque, notamment dans des applications pouvant causer des dommages corporels. Si vous utilisez ce logiciel ou matériel dans le cadre d'applications dangereuses, il est de votre responsabilité de prendre toutes les mesures de secours, de sauvegarde, de redondance et autres mesures nécessaires à son utilisation dans des conditions optimales de sécurité. Oracle Corporation et ses affiliés déclinent toute responsabilité quant aux dommages causés par l'utilisation de ce logiciel ou matériel pour ce type d'applications.

Oracle et Java sont des marques déposées d'Oracle Corporation et/ou de ses affiliés. Tout autre nom mentionné peut correspondre à des marques appartenant à d'autres propriétaires qu'Oracle.

AMD, Opteron, le logo AMD et le logo AMD Opteron sont des marques ou des marques déposées d'Advanced Micro Devices. Intel et Intel Xeon sont des marques ou des marques déposées d'Intel Corporation. Toutes les marques SPARC sont utilisées sous licence et sont des marques ou des marques déposées de SPARC International, Inc. UNIX est une marque déposée concédée sous licence par X/Open Company, Ltd.

Contents

Preface	5
1 Installing and Maintaining a SCSI RAID Storage Device	11
Restrictions and Requirements	11
Installing Storage Arrays	12
▼ How to Install a Storage Array in a New Cluster	12
▼ How to Add a Storage Array to an Existing Cluster	16
Configuring Storage Arrays	22
▼ How to Create a LUN	22
▼ How to Delete a LUN	25
▼ How to Reset the LUN Configuration	27
▼ How to Correct Mismatched Device ID Numbers	29
Maintaining Storage Arrays	31
FRUs That <i>Do Not</i> Require Oracle Solaris Cluster Maintenance Procedures	32
▼ How to Remove a Storage Array	34
▼ How to Replace a Failed Controller or Restore an Offline Controller	37
▼ How to Upgrade Controller Module Firmware	39
▼ How to Add a Disk Drive	40
▼ How to Replace a Disk Drive	41
▼ How to Remove a Disk Drive	43
How to Upgrade Disk Drive Firmware	44
▼ How to Replace a Host Adapter	44
A Cabling Diagrams	49
Sun StorEdge A1000 and Netra st A1000 Cabling Diagrams	49
Installing a Sun StorEdge A1000 or Netra st A1000 Storage Array	49
Sun StorEdge A3500 Cabling Diagrams	51

Installing a Sun StorEdge A3500 Storage System 51

Adding a Sun StorEdge A3500 Storage System 51

Index53

Preface

The Oracle Solaris Cluster 3.3 With StorEdge A1000 Array, Netra st A1000 Array, or StorEdge A3500 System Manual provides procedures specific to SCSI RAID storage devices that use Oracle's Sun StorEdge RAID Manager software in an Oracle Solaris Cluster environment.

Note – In this document, references to Oracle Real Application Clusters also apply to Oracle Parallel Server, unless otherwise stated.

Who Should Use This Book

This book is for Oracle representatives who are performing the initial installation of an Oracle Solaris 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 Oracle software and hardware. Do not use this document as a planning or a pre-sales 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 the following:

- [Chapter 1, “Installing and Maintaining a SCSI RAID Storage Device,”](#) contains procedures about installing, configuring, and maintaining SCSI RAID storage devices that use RAID Manager software.
- [Appendix A, “Cabling Diagrams,”](#) contains cabling diagrams for the various storage devices.

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 Oracle Solaris Cluster 3.3 With StorEdge A1000 Array, Netra st A1000 Array, or StorEdge A3500 System Manual

Revision Date	New Information
January 2009	Update links in the Preface to other books and to Sun Cluster documentation.
September 2010	Update product name throughout book and remove old CLI commands for 3.3 release.
May 2011	Changed links to point to Oracle sites for 3.3 5/11 release.

Related Documentation

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

The following books support the Oracle Solaris Cluster 3.3 release. You can also access the documentation for the Sun Cluster 3.1 and 3.2 releases. All Sun Cluster and Oracle Solaris Cluster documentation is available at <http://www.oracle.com/technetwork/indexes/documentation/index.html>.

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

- [Storage Array and Storage System patches](#)
- [Solaris I/O Multipathing and Sun StorEdge RAID Manager](#)
- [Sun StorEdge A1000 and Netra st A1000 Storage Array](#)
- [Sun StorEdge A3500 system](#)
- [Oracle Solaris Cluster and Sun Cluster Documentation](#)

TABLE P-2 Storage Array and Storage System Patches

Title
Sun StorEdge A1000/Sun StorEdgeA3x00/Sun StorEdge A3500FC/Sun StorEdge Software/Firmware Configuration Matrix 1.61
Log into My Oracle Support to view this Technical Instruction.

TABLE P-3 Solaris I/O multipathing and Sun StorEdge RAID Manager

Title	Part Number
<i>Sun StorEdge RAID Manager 6.22.1 Release Notes</i>	805-7758-14
<i>Sun StorEdge RAID Manager 6.22 User's Guide</i>	806-0478-10
<i>Solaris Fibre Channel Storage Configuration and Multipathing Support Guide</i>	819-0139

TABLE P-4 Sun StorEdge A1000 and Netra st A1000 Storage Array

Title	Part Number
<i>OpenBoot 2.x Command Reference Manual</i>	806-2906-10
<i>Netra st A1000/D1000 Installation and Maintenance Manual</i>	805-7147
<i>Sun StorEdge A1000 and D1000 Installation, Operations, and Service Manual</i>	805-2624
<i>Sun StorEdge D1000 Storage Guide</i>	805-4013

TABLE P-5 Sun StorEdge A3500 Storage System

Title	Part Number
<i>OpenBoot 2.x Command Reference Manual</i>	806-2906-10
<i>Sun StorEdge A3500/A3500FC Hardware Configuration Guide</i>	805-4981

TABLE P-6 Oracle Solaris Cluster and Sun Cluster Documentation

Documentation
Oracle Solaris Cluster 3.3
Sun Cluster 3.2
Sun Cluster 3.1

Using UNIX Commands

This document contains information about commands that are used to install, configure, or upgrade an Oracle Solaris Cluster configuration. This document might not contain complete information about basic UNIX commands and procedures such as shutting down the system, booting the system, and configuring devices.

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

- Online documentation for the Oracle Solaris Operating System (Oracle Solaris OS)
- Other software documentation that you received with your system

- Oracle Solaris Operating System man pages

Getting Help

If you have problems installing or using Oracle Solaris 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, Oracle Solaris 10)
- The release number of Oracle Solaris Cluster (for example, Oracle Solaris Cluster 3.3)

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

Command	Function
<code>prtconf -v</code>	Displays the size of the system memory and reports information about peripheral devices
<code>psrinfo -v</code>	Displays information about processors
<code>showrev -p</code>	Reports which patches are installed
<code>prtdiag -v</code>	Displays system diagnostic information
<code>/usr/cluster/bin/clnode show-rev</code>	Displays Oracle Solaris Cluster release and package version information for each node

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

Documentation and Support

See the following web sites for additional resources:

- [Documentation](http://www.oracle.com/technetwork/indexes/documentation/index.html) (<http://www.oracle.com/technetwork/indexes/documentation/index.html>)
- [Support](http://www.oracle.com/us/support/systems/index.html) (<http://www.oracle.com/us/support/systems/index.html>)

Oracle Software Resources

Oracle Technology Network (<http://www.oracle.com/technetwork/index.html>) offers a range of resources related to Oracle software:

- Discuss technical problems and solutions on the [Discussion Forums](http://forums.oracle.com) (<http://forums.oracle.com>).
- Get hands-on step-by-step tutorials with [Oracle By Example](http://www.oracle.com/technetwork/tutorials/index.html) (<http://www.oracle.com/technetwork/tutorials/index.html>).

Typographic Conventions

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

TABLE P-7 Typographic Conventions

Typeface	Meaning	Example
AaBbCc123	The names of commands, files, and directories, and onscreen computer output	Edit your <code>.login</code> file. Use <code>ls -a</code> to list all files. <code>machine_name%</code> you have mail.
AaBbCc123	What you type, contrasted with onscreen computer output	<code>machine_name%</code> su Password:
<i>aabbcc123</i>	Placeholder: replace with a real name or value	The command to remove a file is <i>rm filename</i> .
<i>AaBbCc123</i>	Book titles, new terms, and terms to be emphasized	Read Chapter 6 in the <i>User's Guide</i> . <i>A cache</i> is a copy that is stored locally. Do <i>not</i> save the file. Note: Some emphasized items appear bold online.

Shell Prompts in Command Examples

The following table shows the default UNIX system prompt and superuser prompt for shells that are included in the Oracle Solaris OS. Note that the default system prompt that is displayed in command examples varies, depending on the Oracle Solaris release.

TABLE P-8 Shell Prompts

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

Installing and Maintaining a SCSI RAID Storage Device

This chapter describes the procedures about how to install, configure, and maintain SCSI RAID storage devices that use Sun StorEdge RAID Manager software in an Oracle Solaris Cluster environment.

The procedures in this chapter apply to the following SCSI RAID storage devices:

- Oracle's Netra st A1000 array
- Oracle's Sun StorEdge A1000 array
- Oracle's Sun StorEdge A3500 system

This chapter contains the following sections:

- [“Restrictions and Requirements” on page 11](#)
- [“Installing Storage Arrays” on page 12](#)
- [“Configuring Storage Arrays” on page 22](#)
- [“Maintaining Storage Arrays” on page 31](#)

Restrictions and Requirements

This section includes only restrictions and support information that have a direct impact on the procedures in this chapter. For general support information, contact your Oracle service provider.

- RAID Manager software 6.22.1 or later version is required.
- (*StorEdge A3500 Only*) Do not use LUNs as quorum devices.

Installing Storage Arrays

This section contains the instructions for installing storage arrays both in new clusters and existing clusters.

TABLE 1-1 Task Map: Installing Storage Arrays

Task	Information
Install an array in a new cluster, before the OS and Oracle Solaris Cluster software are installed.	“How to Install a Storage Array in a New Cluster” on page 12
Add an array to an operational cluster.	“How to Add a Storage Array to an Existing Cluster” on page 16

▼ How to Install a Storage Array in a New Cluster

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

This procedure uses an updated method for setting the `scsi-initiator-id`. The method that was 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 need to update your `nvrampc` script and set the `scsi-initiator-id` by following this procedure.

Before You Begin Before performing this procedure, ensure that you have met the following prerequisites. This procedure relies on the following prerequisites and assumptions.

- You have read the entire procedure.
- You can access necessary patches, drivers, software packages, and hardware.
- Your nodes are powered off or are at the OpenBoot PROM.
- Your arrays are powered off.
- Your interconnect hardware is connected to the nodes in your cluster.
- No software is installed.

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.

For cabling diagrams, see [Appendix A, “Cabling Diagrams.”](#)

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

To avoid SCSI-chain conflicts, the following steps instruct 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 steps instruct 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.

Note – If necessary, halt the nodes so that you can perform OpenBoot PROM (OBP) Monitor tasks at the ok prompt.

For the procedure about 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 encode-int " scsi-initiator-id" property
3: device-end
4: cd /pci@1f,4000/scsi@3
```

```

5: 7 encode-int " scsi-initiator-id" 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 the following command:**

```

{0} ok nvstore
{1} ok

```

■ **To discard the changes, type the following command:**

```

{0} ok nvquit
{1} ok

```

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

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

```

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

```

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

```

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

```

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 System, then apply any required Solaris patches.

For the most current list of patches, see [My Oracle Support](#).

6 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 7](#).
- If you are using a version of the Solaris Operating System earlier than Solaris 8 Update 4, proceed to [Step 7](#).
- If you are using RAID Manager 6.22 and the Solaris 8 Update 4 or later operating environment, reboot both nodes.

```
# reboot
```

7 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*.

For the required version of the RAID Manager software that Oracle Solaris Cluster software supports, see [“Restrictions and Requirements”](#) on page 11.

8 Install patches for the controller modules and RAID Manager software.

For the most current list of patches, see [My Oracle Support](#).

9 Check the NVSRAM file revision for the storage arrays. 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*.

10 Check the controller module firmware revision for the storage arrays. 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*.

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

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

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

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

13 Set up the storage arrays with logical unit numbers (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.

14 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.**15 Ensure that the new logical name for the LUN that you created in [Step 13](#) appears in the `/dev/rdisk` directory on both nodes.**

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

See Also To continue with Oracle Solaris Cluster software and data services installation tasks, see your Oracle Solaris Cluster software installation documentation and the Oracle Solaris Cluster data services developer's documentation. For a list of Oracle Solaris Cluster documentation, see [“Related Documentation” on page 6](#).

▼ How to Add a Storage Array to an Existing Cluster

Use this procedure to add a storage device to an existing cluster. If you need to install a storage device in a new cluster, use the procedure in [“How to Install a Storage Array in a New Cluster” on page 12](#).

You might want to perform this procedure in the following scenarios.

- You need to increase available storage.
- You need to upgrade to a higher-quality or larger storage array.
 - To upgrade storage arrays, remove the old storage array and then add the new storage array.
 - To replace a storage array with the same type of storage array, use this procedure.

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

- Your cluster is operational.
- This procedure defines Node A as the node with which you begin working. Node B is the remaining node.
- This procedure uses an updated method for setting the `scsi-initiator-id`. For this storage array, the method that was 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 need to update your `nvrsrc` script and set the `scsi-initiator-id` by using this procedure.

This procedure provides the long forms of the Oracle Solaris Cluster commands. Most commands also have short forms. Except for the forms of the command names, the commands are identical.

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

1 Install the RAID Manager software.

- For the required version of the RAID Manager software that Oracle Solaris Cluster software supports, see [“Restrictions and Requirements” on page 11](#).
- For the procedure about how to install RAID Manager software, see the *Sun StorEdge RAID Manager Installation and Support Guide*.
- For the most current list of software, firmware, and patches that your storage array or storage system requires, refer to the appropriate EarlyNotifier that is outlined in [“Related Documentation” on page 6](#). This document is available online to Oracle service providers and to customers with service contracts at [My Oracle Support](#).

2 Install the storage array or storage system patches.

For the location of patches and installation instructions, see your Oracle Solaris Cluster release notes documentation. For a list of Oracle Solaris Cluster documentation, see [“Related Documentation” on page 6](#).

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

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

4 Power on the storage array or storage system.

For the procedure about how to power on the storage array or storage system, see your storage documentation. For a list of storage documentation, see [“Related Documentation” on page 6](#).

5 If you are not installing new host adapters in your nodes, skip to [Step 8](#).

6 If you are installing new host adapters in your nodes, shut down and power off Node A.

For the procedure about how to shut down and power off a node, see your Oracle Solaris Cluster system administration documentation. For a list of Oracle Solaris Cluster documentation, see [“Related Documentation” on page 6](#).

7 Install the host adapters in Node A.

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 or storage system to Node A.

For cabling diagrams, see [Appendix A, “Cabling Diagrams.”](#)

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

To avoid SCSI-chain conflicts, the following steps instruct 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 steps instruct 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.

Note – If necessary, halt the nodes so that you can perform OpenBoot PROM (OBP) Monitor tasks at the ok prompt.

For the procedure about 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 encode-int " scsi-initiator-id" property
3: device-end
4: cd /pci@1f,4000/scsi@3
5: 7 encode-int " scsi-initiator-id" 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 the following command:

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

- To discard the changes, type the following command:

```
{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 `nvrnrc` script, as shown in the following example.**

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

- 10 If you are not installing new host adapters in Node B, skip to [Step 13](#).**
- 11 If you are installing new host adapters in Node B to connect Node B to the storage array or storage system, shut down and power off the node.**

For the procedure about how to shut down and power off a node, see your Oracle Solaris Cluster system administration documentation. For a list of Oracle Solaris Cluster documentation, see “[Related Documentation](#)” on page 6.

- 12 Install the host adapters in Node B.**

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

- 13 Cable the storage array or storage system to Node B.**

For cabling diagrams, see “[Adding a Sun StorEdge A3500 Storage System](#)” on page 51.

- 14 If you did not power off Node B to install a host adapter, skip to [Step 17](#).**

- 15 If you powered off Node B to install a host adapter, power on Node B and the storage array or storage system.**

Do not enable the node to boot. If necessary, halt the system to continue with OpenBoot PROM (OBP) Monitor tasks.

- 16 Verify that Node B recognizes the new host adapters and disk drives.**

If the node does not recognize the new hardware, check all hardware connections and repeat the installation steps you performed in [Step 12](#).

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

- 17 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
...
```

- 18 If you did not power off Node B to install a host adapter, skip to [Step 23](#).**
- 19 If you powered off Node B to install a host adapter, perform a reconfiguration boot to create the new Oracle Solaris device files and links.**
- 20 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*.
- 21 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” on page 39](#).
- 22 One node at a time, boot each node into cluster mode.**
`# reboot`
- 23 On one node, verify that the device IDs have been assigned to the LUNs for all nodes. These nodes are attached to the storage array or storage system.**
`# cldevice show`
- 24 (StorEdge A3500 Only) Verify that the controller module is set to active/active mode.**
For more information about controller modes, see the *Sun StorEdge RAID Manager Installation and Support Guide* and the *Sun StorEdge RAID Manager User's Guide*.

See Also To create a LUN from disk drives that are unassigned, see [“How to Create a LUN” on page 22](#).

To upgrade controller module firmware, see [“How to Upgrade Controller Module Firmware” on page 39](#).

Configuring Storage Arrays

This section contains the procedures about how to configure a storage array or storage system *after* you install Oracle Solaris Cluster software. [Table 1–2](#) lists these procedures.

To configure a storage array or storage system *before* you install Oracle Solaris 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 Oracle Solaris 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 ID might change unexpectedly. When you perform a check of the device ID configuration by running the `cldevice check` command, the following error message appears on your console if the device ID 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 device IDs that report this error, run the `cldevice repair` command for each affected device.

TABLE 1–2 Task Map: Configuring Disk Drives

Task	Information
Create a logical unit number (LUN).	“How to Create a LUN” on page 22
Remove a LUN.	“How to Delete a LUN” on page 25
Reset the LUN configuration.	“How to Reset the LUN Configuration” on page 27

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

- Rebalancing running LUNs
- Creating a hot spare
- Deleting a hot spare
- Increasing the size of a drive group

▼ How to Create a LUN

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

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 RAC data stored on the storage array. Do not place RAC data under volume management control. You must place all non-RAC 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-RAC data.

Hardware RAID uses the storage array's or storage 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 without volume management disruptions. If you use hardware RAID level 0, some maintenance procedures require additional volume management administration because the availability of the LUNs is impacted.

Before You Begin

This procedure provides the long forms of the Oracle Solaris Cluster commands. Most commands also have short forms. Except for the forms of the command names, the commands are identical.

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

- 1 **With all nodes booted and attached to the storage array or 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 array or storage system.

If the following SCSI warning is displayed, ignore the message. Continue with the next step.

```
...
corrupt label - wrong magic number
```

For the procedure about how to create a LUN, refer to your storage device's documentation. 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/rdsk` directory on both nodes.

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

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

```
# cldevice populate
```

- 5 Ensure that the device ID numbers for the LUNs are the same on both nodes. In the sample output that follows, the device ID numbers are different.

For an example of output indicating that the ID numbers have not been properly updated, see [Example 1-1](#).

Run the following command:

```
# cldevice show
```

- 6 If the device ID numbers that you received from running the `cldevice show` or `scdidadm` command in [Step 5](#) are the same for both nodes, proceed to [Step 8](#).

- 7 If the device ID numbers that you received from running the `cldevice show` or `scdidadm` command are different, perform the procedure in “[How to Correct Mismatched Device ID Numbers](#)” on page 29 before you proceed to [Step 8](#) of this procedure.

- 8 (A1000 Only) 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 Solaris Volume Manager or Veritas Volume Manager documentation.

Note – The StorEdge A3500 system does not support using LUNs as quorum devices.

Example 1-1 Verifying the Device IDs

[Step 5](#) in the preceding procedure directs you to verify that the device ID numbers for the LUNs are the same on both nodes. In the sample output that follows, the device ID numbers are different.

```
# cldevice show
...
DID Device Name:                /dev/did/rdsk/d3
Full Device Path:                phys-schost-2:/dev/rdsk/clt3d0
Full Device Path:                phys-schost-1:/dev/rdsk/clt3d1
Replication:                     none
default_fencing:                 global
```


▼ How to Delete a LUN

Use this procedure to delete one or more LUNs. You might need to delete a LUN to free up or reallocate resources, or to use the disks for other purposes. See the *Sun StorEdge RAID Manager Release Notes* for the latest information about LUN administration.

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

- All data on the LUN that you delete will be removed.
- You are not deleting LUN 0.

This procedure provides the long forms of the Oracle Solaris Cluster commands. Most commands also have short forms. Except for the forms of the command names, the commands are identical.

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

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

format

For example:

```
phys-schost-1# format
Searching for disks...done
AVAILABLE DISK SELECTIONS:
0. c0t5d0 <SYMBIOS-StorEdgeA3500FCr-0301 cyl3 alt2 hd64 sec64>
   /pseudo/rdnexus@0/rdriver@5,0
1. c0t5d1 <SYMBIOS-StorEdgeA3500FCr-0301 cyl2025 alt2 hd64 sec64>
   /pseudo/rdnexus@0/rdriver@5,1
```

- 2 (A1000 Only) If the LUN that you are removing is not a quorum device, proceed to [Step 4](#).**

Note – Your storage array or storage system might not support LUNs as quorum devices. To determine if this restriction applies to your storage array or storage system, see [“Restrictions and Requirements” on page 11](#).

To determine whether this LUN is a quorum device, use the following command.

clquorum show

- 3 (A1000 Only) If the LUN that you are removing is a quorum device, relocate that quorum device to another suitable storage array.**

For procedures about how to add and remove quorum devices, see your Oracle Solaris Cluster system administration documentation. For a list of Oracle Solaris Cluster documentation, see [“Related Documentation” on page 6](#).

- 4 If a volume manager does not manage the LUN that you are deleting, proceed to [Step 6](#).

- 5 If a volume manager manages the LUN that you are deleting, remove the LUN from any diskset or disk group. For more information, see your 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
```

- 6 Delete the LUN.

For the procedure about how to delete a LUN, refer to your storage device's documentation.

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

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

- 8 Complete the removal of the paths by issuing the following RAID Manager commands.

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

- 9 (*StorEdge A3500 Only*) 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:

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

Therefore, the alternate paths are as follows:

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

- 10 (*StorEdge A3500 Only*) Remove the alternate paths to the LUNs you are deleting.

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

- 11 On both nodes, remove all obsolete device IDs.

```
# cldevice clear
```

- 12 Move all resource groups and device groups off the node.

```
# clnode evacuate from-node
```

13 Shut down the node.

For the procedure about how to shut down and power off a node, see your Oracle Solaris Cluster system administration documentation. For a list of Oracle Solaris Cluster documentation, see [“Related Documentation” on page 6](#).

14 Perform a reconfiguration boot to create the new Solaris device files and links.**15 Repeat [Step 6](#) through [Step 14](#) on the other node that is attached to the storage array or storage system.**

▼ How to Reset the LUN Configuration

Use this procedure to completely remove and reset the LUN configuration.



Caution – If you reset a LUN configuration, a new device ID number is assigned to LUN 0. This change occurs because the software assigns a new world wide name (WWN) to the new LUN.

Before You Begin

This procedure provides the long forms of the Oracle Solaris Cluster commands. Most commands also have short forms. Except for the forms of the command names, the commands are identical.

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

1 From one node that is connected to the storage array or storage system, determine the paths to the LUNs you are resetting.

format

For example:

```
phys-schost-1# format
Searching for disks...done
AVAILABLE DISK SELECTIONS:
0. c0t5d0 <SYMBIOS-StorEdgeA3500FCr-0301 cyl3 alt2 hd64 sec64>
   /pseudo/rdnexus@0/rdriver@5,0
1. c0t5d1 <SYMBIOS-StorEdgeA3500FCr-0301 cyl2025 alt2 hd64 sec64>
   /pseudo/rdnexus@0/rdriver@5,1
```

2 (A1000 Only) If the LUN that you are resetting is not a quorum device, proceed to [Step 4](#).

To determine whether this LUN is a quorum device, use the following command.

clquorum show

- 3 **(A1000 Only) If the LUN that you are resetting is a quorum device, relocate that quorum device to another suitable storage array.**

For procedures about how to add and remove quorum devices, see your Oracle Solaris Cluster system administration documentation. For a list of Oracle Solaris Cluster documentation, see [“Related Documentation” on page 6](#).

- 4 **If a volume manager does not manage the LUNs on the controller module that you are resetting, proceed to [Step 6](#).**

- 5 **If a volume manager manages the LUNs on the controller module that you are resetting, remove the LUN from any diskset or disk group.**

For more information, see your Solaris Volume Manager or Veritas Volume Manager documentation.

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

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

- 6 **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*.

- 7 **(StorEdge A3500 Only) Set the controller module back to active/active mode.**

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

- 8 **Use the `format` command to label the new LUN 0.**

- 9 **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*
```

- 10 **(StorEdge A3500 Only) Determine the alternate paths to the old LUNs you reset. Use the `lad` command.**

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.

Example:

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

Therefore, the alternate paths are as follows:

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

- 11 *(StorEdge A3500 Only)* Remove the alternate paths to the old LUNs you reset.

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

- 12 On both nodes, update device namespaces.

```
# devfsadm -C
```

- 13 On both nodes, remove all obsolete device IDs.

```
# cldevice clear
```

- 14 Move all resource groups and device groups off the node.

```
# clnode evacuate from-node
```

- 15 Shut down the node.

For the procedure about how to shut down and power off a node, see your Oracle Solaris Cluster system administration documentation. For a list of Solaris Cluster documentation, see [“Related Documentation” on page 6](#).

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

If an error message like the following appears, ignore it. Continue with the next step.

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

- 17 After the node reboots and joins the cluster, repeat [Step 9](#) through [Step 16](#) on the other node. This node is attached to the storage array or storage system.

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

▼ How to Correct Mismatched Device ID Numbers

Use this section to correct mismatched device ID numbers that might appear during the creation of LUNs. You correct the mismatch by deleting Solaris and Oracle Solaris Cluster paths to the LUNs that have device ID 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 22](#).

Before You Begin This procedure provides the long forms of the Oracle Solaris Cluster commands. Most commands also have short forms. Except for the forms of the command names, the commands are identical.

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

- 1 **From one node that is connected to the storage array or storage system, determine the paths to the LUNs. These LUNs have different device ID numbers.**

format

- 2 **Remove the paths to the LUNs that have different device ID numbers.**

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

- 3 **(StorEdge A3500 Only) Use the `lad` command to determine the *alternate* paths to the LUNs that have different device ID 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:

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

Therefore, the alternate paths are as follows:

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

- 4 **(StorEdge A3500 Only) Remove the *alternate* paths to the LUNs that have different device ID numbers.**

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

- 5 **On both nodes, remove all obsolete device IDs.**

cldevice clear

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

clnode evacuate *from-node*

7 Shut down the node.

For the procedure about how to shut down and power off a node, see your Oracle Solaris Cluster system administration documentation. For a list of Oracle Solaris Cluster documentation, see [“Related Documentation” on page 6](#).

8 To create the new Oracle Solaris device files and links, perform a reconfiguration boot by adding -r to your boot instruction.**9 Repeat [Step 1](#) through [Step 8](#) on the other node. This node is attached to the storage array or storage system.****10 Return to [“How to Create a LUN” on page 22](#).**

Maintaining Storage Arrays

The maintenance procedures in [“FRUs That Do Not Require Oracle Solaris Cluster Maintenance Procedures” on page 32](#) are performed the same as in a noncluster environment. [Table 1–3](#) lists the procedures that require cluster-specific steps.

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 might change unexpectedly. When you perform a check of the device ID configuration by running the `cldevice check` command, the following error message appears on your console if the device ID 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 device IDs that report this error, run the `cldevice repair` command for each affected device.

TABLE 1–3 Task Map: Maintaining a Storage Array or Storage System

Task	Information
Remove a storage array or storage system	“How to Remove a Storage Array” on page 34
Replace a storage array or storage system	“How to Add a Storage Array to an Existing Cluster” on page 16
Replacing a storage array or storage system, requires first removing the storage array or storage system, then adding a new storage array or storage system to the configuration.	“How to Remove a Storage Array” on page 34
Replace a failed controller module or restore an offline controller module	“How to Replace a Failed Controller or Restore an Offline Controller” on page 37

TABLE 1-3 Task Map: Maintaining a Storage Array or Storage System (Continued)

Task	Information
Upgrade controller module firmware and NVSRAM file	“How to Upgrade Controller Module Firmware” on page 39
Add a disk drive	“How to Add a Disk Drive” on page 40
Replace a disk drive	“How to Replace a Disk Drive” on page 41
Remove a disk drive	“How to Remove a Disk Drive” on page 43
Upgrade disk drive firmware	“How to Upgrade Disk Drive Firmware” on page 44
Replace a host adapter	“How to Replace a Host Adapter” on page 44

FRUs That *Do Not* Require Oracle Solaris Cluster Maintenance Procedures

Each storage device has a different set of FRUs that do not require cluster-specific procedures. Choose among the following storage devices:

- [Netra st A1000 array](#)
- [Sun StorEdge A1000 array](#)
- [Sun StorEdge A3500 system](#)

Sun StorEdge A1000 Array and Netra st A1000 Array FRUs

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* and the *Netra st A1000/D1000 Installation and Maintenance Manual* for these procedures.

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

Replacing a storage array-to-host SCSI cable requires no cluster-specific procedures. See the *Sun StorEdge RAID Manager User’s Guide* and the *Sun StorEdge RAID Manager Release Notes* for these procedures.

Sun StorEdge A3500 System FRUs

With the exception of one item, the following is a list of administrative tasks that require no cluster-specific procedures. Shut down the cluster, and then see the *Sun StorEdge A3500/A3500FC Controller Module Guide*, the *Sun StorEdge A1000 and D1000 Installation, Operations, and Service Manual*, and the *Sun StorEdge Expansion Cabinet Installation and Service Manual* for the following procedures. See the Oracle Solaris Cluster system administration documentation for procedures about how to shut down a cluster. For a list of Oracle Solaris Cluster documentation, see [“Related Documentation”](#) on page 6.

- Replacing a power cord that connects to the cabinet power distribution unit (see the *Sun StorEdge Expansion Cabinet Installation and Service Manual*).
- Replacing a power cord to a storage array (see the *Sun StorEdge A1000 and D1000 Installation, Operations, and Service Manual*).
- Replacing the power-supply housing.
- Replacing the battery unit.
- Replacing a DC-power or battery harness.
- Replacing a controller module assembly.
- Replacing a controller module power cord.
- Replacing a controller module card cage.

The following is a list of administrative tasks that require no cluster-specific procedures. See the *Sun StorEdge A3500/A3500FC Controller Module Guide*, the *Sun StorEdge RAID Manager User's Guide*, the *Sun StorEdge RAID Manager Release Notes*, the *Sun StorEdge FC-100 Hub Installation and Service Manual*, and the documentation that shipped with your FC hub or FC switch for the following procedures.

- Replacing a storage array-to-host SCSI cable.
- Replacing a controller module fan canister.
- Replacing the power supply fan canister.
- Replacing a SCSI cable from the controller module to the storage array.
- Replacing a storage array-to-host or storage array-to-hub fiber-optic cable.
- Replacing an FC hub (see the *Sun StorEdge FC-100 Hub Installation and Service Manual*).
- Replacing an FC hub gigabit interface converter (GB(C) or Small Form-Factor Pluggable (SFP) that connects cables to the host or hub.
- Replacing a GBIC or an SFP on a node.
- Replacing an FC switch.

▼ How to Remove a Storage Array



Caution – This procedure removes all data that is on the storage array or storage system you are removing.

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

- Your cluster is operational.
- You no longer need the data that is stored on the storage array or storage system you are removing.

This procedure provides the long forms of the Oracle Solaris Cluster commands. Most commands also have short forms. Except for the forms of the command names, the commands are identical.

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

- 1 Migrate any Oracle Real Application Clusters tables, data services, or volumes off the storage array or storage system.**
- 2 If no LUNS in the storage array that you are removing are quorum devices, proceed to [Step 4](#).**

Note – Your storage array or storage system might not support LUNs as quorum devices. To determine if this restriction applies to your storage array or storage system, see [“Restrictions and Requirements” on page 11](#).

To determine whether any LUNs in the storage array are quorum devices, use the following command.

```
# clquorum show
```

- 3 If one of the LUNs in the storage array is a quorum device, relocate that quorum device to another suitable storage array.**

For procedures about how to add and remove quorum devices, see your Oracle Solaris Cluster system administration documentation. For a list of Oracle Solaris Cluster documentation, see [“Related Documentation” on page 6](#).

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

For instructions, see your storage device documentation and your operating system documentation.

- 5 If a volume manager does not manage any of the LUNs on the controller module you are removing, proceed to [Step 12](#).

- 6 If a volume manager manages any LUNs on the controller module that you are removing, remove the LUN from any diskset or disk group.

For more information, see your Solaris Volume Manager or Veritas Volume Manager documentation.

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

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

- 7 Delete the LUN.

For the procedure about how to delete a LUN, see your storage device's documentation.

- 8 Remove the paths to the LUNs you deleted in [Step 7](#).

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

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

- 9 On all nodes, remove references to the storage array.

```
# cldevice clear
```

- 10 *(StorEdge A3500 Only)* Use the `lad` command to determine the *alternate* paths to the LUN you are deleting.

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

For example:

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

Therefore, the alternate paths are as follows:

```
/dev/osa/dev/dsk/c1t4d1*
/dev/osa/dev/rdsk/c1t4d1*
```

- 11 *(StorEdge A3500 Only)* Remove the *alternate* paths to the LUNs you deleted in [Step 7](#).

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

- 12 Disconnect all cables from the storage array and storage system. Remove the hardware from your cluster.

- 13 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 no other parallel SCSI devices are connected to the nodes, you can delete the contents of the `nvrामrc` script. At the OpenBoot PROM, set `setenv use-nvrामrc?` to `false`.

- 14 Remove any unused host adapter from nodes that were attached to the storage array or storage system.**

- a. Shut down and power off Node A, from which you are removing a host adapter.**

For the procedure about how to shut down and power off a node, see your Oracle Solaris Cluster system administration documentation. For a list of Oracle Solaris Cluster documentation, see [“Related Documentation” on page 6](#).

- b. Remove the host adapter from Node A.**

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 B that was attached to the storage array or storage system.**

- 15 Restore resource groups to their primary nodes.**

Use the following command for each resource group you want to return to the original node.

```
# clresourcegroup switch -n nodename resourcegroup1[ resourcegroup2 ...]
```

nodename

For failover resource groups, the node to which the groups are returned. For scalable resource groups, the node list to which the groups are returned.

resourcegroup1[resourcegroup2 ...]

The resource group or groups that you are returning to the node or nodes.

- 16 If this is not the *last* storage array or storage system in your cluster, you are finished with this procedure.**

- 17 If this is the *last* storage array or storage system in the cluster, 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 or storage system.

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

This procedure assumes that your cluster is operational. For conceptual information about SCSI reservations and failure fencing, see your Oracle Solaris Cluster concepts documentation. For a list of Oracle Solaris Cluster documentation, see [“Related Documentation” on page 6](#).

Before You Begin This procedure provides the long forms of the Oracle Solaris Cluster commands. Most commands also have short forms. Except for the forms of the command names, the commands are identical.

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

- 1 *(StorEdge A1000 Only)* **If none of the LUNs in the storage array is a quorum device, proceed to [Step 3](#).**

Note – Your storage array or storage system might not support LUNs as quorum devices. To determine if this restriction applies to your storage array or storage system, see [“Restrictions and Requirements” on page 11](#).

To determine whether any LUNs in the storage array are quorum devices, use the following command.

```
# clquorum show
```

- 2 *(StorEdge A1000 Only)* **If any of the LUNs in the storage array is a quorum device, relocate that quorum device to another suitable storage array.**

For procedures about how to add and remove quorum devices, see your Oracle Solaris Cluster system administration documentation. For a list of Oracle Solaris Cluster documentation, see [“Related Documentation” on page 6](#).

- 3 *(StorEdge A3500 Only)* 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 8](#) 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*.

- 4 **Restart the RAID Manager daemon.**

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

- 5 **If your controller module is offline, but does not have a failed controller, proceed to [Step 7](#).**
- 6 **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.

- 7 **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 about 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*.

- 8 **On one node that is connected to the storage array or storage system, verify that the controller has the correct SCSI reservation state.**

Use the following command on LUN 0 of the controller you want to bring online.

In the following command, *devicename* is the full UNIX path name of the device, for example `/dev/dsk/c1tXdY`

```
# cldevice repair devicename
```

- 9 *(StorEdge A3500 Only)* **Set the controller to active/active mode. Assign LUNs to the controller.**
For more information about controller modes, see the *Sun StorEdge RAID Manager Installation and Support Guide* and the *Sun StorEdge RAID Manager User's Guide*.
- 10 *(StorEdge A3500 Only)* **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*.
- 11 *(StorEdge A3500 Only)* **Restart the RAID Manager daemon.**

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

▼ How to Upgrade Controller Module Firmware

Use either the *online* or the *offline* method to upgrade your NVSRAM firmware. The method that you choose depends on your firmware.

Before You Begin This procedure assumes that your cluster is operational

- 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 have mirrored data.**
 - a. **Halt all activity to the controller module.**
For more information, see your Solaris Volume Manager or Veritas Volume Manager documentation.
 - b. **Update the firmware files by using the offline method, as described in the *Sun StorEdge RAID Manager User's Guide*.**

c. Restore all activity to the controller module.

For more information, see your 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 Oracle Solaris 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 Oracle Solaris Cluster system administration documentation. For a list of Oracle Solaris Cluster documentation, see [“Related Documentation” on page 6](#).

c. Update the firmware files by using the offline method, as described in the *Sun StorEdge RAID Manager User’s Guide*.

d. Boot both nodes into cluster mode.

For more information about how to boot nodes, see your Oracle Solaris Cluster system administration documentation. For a list of Oracle Solaris Cluster documentation, see [“Related Documentation” on page 6](#).

This step completes the firmware upgrade.

▼ How to Add a Disk Drive

Adding a disk drive enables you to increase your storage space after a storage array has been added to your cluster.



Caution – If the disk drive that you are adding was previously owned by another controller module, reformat the disk drive to wipe clean the old DacStore information before adding the disk drive to this storage array. If any old DacStore information remains, it can cause aberrant behavior including the appearance of ghost disks or LUNs in the RAID Manager interfaces.

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

- Your cluster is operational.
- Your storage array contains an empty disk slot.
- Your nodes are not configured with dynamic reconfiguration functionality.

If your nodes are configured for dynamic reconfiguration, see the Oracle Solaris Cluster system administration documentation and skip steps that instruct you to shut down the node. For a list of Oracle Solaris Cluster documentation, see [“Related Documentation” on page 6](#).

1 Verify that the new disk drive is formatted.

For information about how to move drives between storage arrays, 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 your storage documentation. For a list of storage documentation, see [“Related Documentation” on page 6](#).

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

4 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*.

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 1](#) through [Step 5](#) for each disk drive you are adding.

See Also To create LUNs for the new drives, see [“How to Create a LUN” on page 22](#) for more information.

▼ How to Replace a Disk Drive

Removing a disk drive enables you to reduce or reallocate your existing storage pool. You might want to perform this procedure if a disk has failed or is behaving in an unreliable manner.

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

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

- Your cluster is operational.
- Your nodes are not configured with dynamic reconfiguration functionality.
If your nodes are configured for dynamic reconfiguration, see the Oracle Solaris Cluster system administration documentation and skip steps that instruct you to shut down the node. For a list of Oracle Solaris Cluster documentation, see [“Related Documentation” on page 6](#).

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

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 Solaris Volume Manager or Veritas Volume Manager documentation.

▼ How to Remove a Disk Drive

Removing a disk drive enables you to reduce or reallocate your existing storage pool. You might want to perform this procedure in the following scenarios.

- You no longer need to make data accessible to a particular node.
- You want to migrate a portion of your storage to another storage array.

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

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

- Your cluster is operational.
- You do not need to remove the entire storage array.
If you need to remove the storage array, see [“How to Remove a Storage Array” on page 34](#).
- You do not need to replace the storage array's chassis.
If you need to replace your storage array's chassis, see [“FRUs That Do Not Require Oracle Solaris Cluster Maintenance Procedures” on page 32](#).
- Your nodes are not configured with dynamic reconfiguration functionality.
If your nodes are configured for dynamic reconfiguration, see the Oracle Solaris Cluster system administration documentation and skip steps that instruct you to shut down the node. For a list of Oracle Solaris Cluster documentation, see [“Related Documentation” on page 6](#).

This procedure provides the long forms of the Oracle Solaris Cluster commands. Most commands also have short forms. Except for the forms of the command names, the commands are identical.

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

- 1 If the LUN that is associated with the disk drive you are removing is not a quorum device, proceed to [Step 3](#).**

Note – Your storage array or storage system might not support LUNs as quorum devices. To determine if this restriction applies to your storage array or storage system, see [“Restrictions and Requirements” on page 11](#).

To determine whether any LUNs in the storage array are quorum devices, use the following command.

```
# clquorum show
```

- 2 **If the LUN that is associated with the disk drive you are removing is a quorum device, relocate that quorum device to another suitable storage array.**

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

- 3 **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 25](#).

- 4 **Remove the disk drive from the storage array.**

For the procedure about how to remove a disk drive, see your storage documentation. For a list of storage documentation, see [“Related Documentation” on page 6](#).



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 Oracle service provider to perform disk drive firmware updates. If you need to upgrade drive firmware, contact your Oracle service provider.

▼ How to Replace a Host Adapter

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*.

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

- Your cluster is operational.
- The node on which the host adapter resides is attached to a *SCSI-based* storage array or storage system.

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

This procedure provides the long forms of the Oracle Solaris Cluster commands. Most commands also have short forms. Except for the forms of the command names, the commands are identical.

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

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

```
# clresourcegroup status -n nodename
# cldevicegroup status -n nodename
```

Note the device groups, the resource groups, and the node list for the resource groups. You will need this information to restore the cluster to its original configuration in [Step 25](#) of this procedure.

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

```
# clnode evacuate fromnode
```

3 Without powering off the node, shut down Node A.

For the procedure about how to shut down and power off a node, see your Oracle Solaris Cluster system administration documentation. For a list of Oracle Solaris Cluster documentation, see [“Related Documentation” on page 6](#).

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

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

5 From the controller module end of the SCSI cable, disconnect the SCSI bus A cable. This cable connects the controller module to Node A. Afterward, replace this cable with a differential SCSI terminator.

6 Restart I/O activity on SCSI bus A.

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

7 If servicing the failed host adapter affects SCSI bus B, proceed to [Step 9](#).

8 If servicing the failed host adapter does not affect SCSI bus B, skip to [Step 12](#).

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

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

- 10 From the controller module end of the SCSI cable, disconnect the SCSI bus B cable. This cable connects the controller module to Node A. Afterward, replace this cable with a differential SCSI terminator.**
- 11 Restart I/O activity on SCSI bus B.**
For instructions, see the *Sun StorEdge RAID Manager User's Guide*.
- 12 Power off Node A.**
- 13 Replace Node A's host adapter.**
For the procedure about how to replace a host adapter, see the documentation that shipped with your node hardware.
- 14 Power on Node A. Do not enable the node to boot. If necessary, halt the system.**
- 15 From Node B, halt I/O activity to the controller module on SCSI bus A.**
For instructions, see the *Sun StorEdge RAID Manager User's Guide*.
- 16 Remove the differential SCSI terminator from SCSI bus A. Afterward, reinstall the SCSI cable to connect the controller module to Node A.**
- 17 Restart I/O activity on SCSI bus A.**
For instructions, see the *Sun StorEdge RAID Manager User's Guide*.
- 18 Did you install a differential SCSI terminator to SCSI bus B in [Step 10](#)?**
 - If no, skip to [Step 21](#).
 - If yes, halt I/O activity on SCSI bus B, then continue with [Step 19](#).
- 19 Remove the differential SCSI terminator from SCSI bus B. Afterward, reinstall the SCSI cable to connect the controller module to Node A.**
- 20 Restart I/O activity on SCSI bus B.**
For instructions, see the *Sun StorEdge RAID Manager User's Guide*.
- 21 Bring the controller module online.**
For instructions, see the *Sun StorEdge RAID Manager User's Guide*.
- 22 Rebalance all logical unit numbers (LUNs).**
For instructions, see the *Sun StorEdge RAID Manager User's Guide*.

- Perform the following step for each device group you want to return to the original node.

<code>-n <i>nodename</i></code>	The node to which you are restoring device groups.
---------------------------------	--

In these commands, *devicegroup* is one or more device groups that are returned to the node.

- Perform the following step for each resource group you want to return to the original node.

<i>nodename</i>	For failover resource groups, the node to which the groups are returned. For scalable resource groups, the node list to which the groups are returned.
-----------------	--

Cabling Diagrams

- Sun StorEdge A1000 and Netra st A1000 arrays
- Sun StorEdge A3500 system

Sun StorEdge A1000 and Netra st A1000 Cabling Diagrams

Installing a Sun StorEdge A1000 or Netra st A1000 Storage Array

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

To cable the storage arrays, connect the differential SCSI cable between the node and the storage array. Ensure 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 storage array, node, and host adapter. [Figure A-1](#) and [Figure A-2](#) illustrate a storage array that is cabled in a Oracle Solaris Cluster environment.

FIGURE A-1 Installing a Storage Array With One Pair





Figure A-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 A-2 Storage Array With Two Pairs and Daisy Chaining



Sun StorEdge A3500 Cabling Diagrams

Installing a Sun StorEdge A3500 Storage System

FIGURE A-3 Installing a StorEdge A3500 Storage System: An Example




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*.

Adding a Sun StorEdge A3500 Storage System

Connect the differential SCSI cable between the node and the controller module, as shown in [Figure A-4](#). Ensure that the *entire* SCSI bus length to each storage array is less than 25 m. This measurement includes the cables to both nodes, as well as the bus length that is internal to each storage array, node, and host adapter.

FIGURE A-4 Cabling a StorEdge A3500 Storage System: An Example, Part 1







-  Differential SCSI terminators
-  Controller A
-  Controller B

FIGURE A-5 Cabling a StorEdge A3500 Storage System: An Example, Part 2



Index

A

A1000

- FRUs, 32
- procedures, 11–47
- related documents, 7

A3500

- FRUs, 33
- procedures, 11–47
- related documents, 7

adding

- See also* installing
- disk drives, 40–41
- storage arrays, 16–21
- storage systems, 16–21

B

battery harness, replacing, 33

battery units

- replacing (for A1000), 32
- replacing (for A3500), 33

C

cables

- replacing fiber-optic, 33
- replacing SCSI, 32, 33

cabling diagrams, 49–52

configuring

- storage arrays, 22–31

configuring (*Continued*)

- storage systems, 22–31

controller module fan canisters, replacing, 33

controller modules

- replacing assemblies, 33
- replacing card cages, 33
- replacing power cords, 33
- upgrading firmware, 39–40

controllers, replacing and restoring, 37–39

cooling canisters, replacing, 32

correcting, DID numbers, 29–31

creating

- hot spares, 22
- LUNs, 22–24

D

deleting

- See also* removing
- hot spares, 22
- LUNs, 25–27

diagrams, cabling, 49–52

DID numbers, correcting, 29–31

disk drive firmware, upgrading, 44

disk drives

- adding, 40–41
- removing, 43–44
- replacing, 41–42

drive groups, expanding, 22

E

expanding, drive groups, 22

F

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

field replaceable units (FRUs), *See* FRUs

FRUs

A1000, 32

A3500, 33

G

GBICs, replacing, 33

gigabit interface converters, *See* GBICs

H

help, 8

host adapters, replacing, 44–47

hot spares, creating and deleting, 22

hubs, replacing, 33

I

installing

See also adding

storage arrays, 12–16

storage system, 12–16

L

logical unit numbers, *See* LUNs

logical units, LUNs, 22–24

logical volumes, *See* LUNs

LUNs

creating, 22–24

deleting, 25–27

rebalancing, 22

LUNs (*Continued*)

resetting configuration, 27–29

M

maintaining

storage arrays, 31–47

storage systems, 31–47

N

Netra st A1000, procedures, 11–47

O

Oracle Parallel Server, 5

P

power cables, replacing storage array cables, 32

power cords

replacing, 33

replacing cabinet cords, 33

replacing storage array cords, 33

power supplies, replacing, 32

power supply housing, replacing, 33

R

RAID

hardware, 22, 23

software, 22, 23

rebalancing, LUNs, 22

removing

See also deleting

disk drives, 43–44

storage arrays, 34–37

storage systems, 34–37

replacing

battery harness, 33

replacing (*Continued*)

- battery units (for A1000), 32
- battery units (for A3500), 33
- controller module assemblies, 33
- controller module card cages, 33
- controller module fan canisters, 33
- controller module power cords, 33
- controllers, 37–39
- cooling canisters, 32
- disk drives, 41–42
- fiber-optic cables, 33
- GBICs, 33
- host adapters, 44–47
- hubs, 33
- power cables, 32
- power cords, 33
- power supplies, 32
- power supply fan canisters, 33
- power supply housing, 33
- SCSI cables, 32, 33
- SFPs, 33
- storage array power cords, 33
- switches, 33

requirements, 11

resetting, LUN configurations, 27–29

restoring, controllers, 37–39

restrictions, 11

S

SFPs, replacing, 33

Small Form-Factor Pluggable, *See* SFPs

storage array procedures, 11–47

storage arrays

- adding, 16–21
- configuring, 22–31
- installing, 12–16
- maintaining, 31–47
- removing, 34–37

storage systems

- adding, 16–21
- cabling, 49–52
- configuring, 22–31
- installing, 12–16

storage systems (*Continued*)

- maintaining, 31–47

- removing, 34–37

Sun StorEdge A1000 procedures, 11–47

Sun StorEdge A3500 procedures, 11–47

switches, replacing, 33

T

technical support, 8

U

upgrading

- controller module firmware, 39–40

- disk drive firmware, 44

