

Oracle® Solaris Cluster 3.3 With Sun StorEdge A3500FC System Manual

SPARC Platform Edition

Copyright © 2000, 2010, Oracle and/or its affiliates. All rights reserved.

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.

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.

If this is software or related software 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.

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 which 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.

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

AMD, Opteron, the AMD logo, and the AMD Opteron logo are trademarks or registered trademarks of Advanced Micro Devices. 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. UNIX is a registered trademark licensed through X/Open Company, Ltd.

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.

Contents

Preface	5
1 Installing and Maintaining a Sun StorEdge A3500FC System	11
Installing Storage Systems	11
▼ How to Install a Storage System in a New Cluster	12
▼ How to Add a Storage System to an Existing Cluster	15
Configuring Storage Systems	20
▼ How to Create a LUN	21
▼ How to Delete a LUN	23
▼ How to Reset the LUN Configuration	24
▼ How to Correct Mismatched Device ID Numbers	26
Maintaining Storage Systems	28
StorEdge A3500FC System FRUs	29
▼ How to Remove a Storage System	30
▼ How to Replace a Failed Controller or Restore an Offline Controller	32
▼ How to Upgrade Controller Module Firmware in a Running Cluster	34
▼ How to Add a Disk Drive in a Running Cluster	35
▼ How to Replace a Failed Disk Drive in a Running Cluster	36
▼ How to Remove a Disk Drive From a Running Cluster	36
How to Upgrade Disk Drive Firmware in a Running Cluster	37
▼ How to Replace a Host Adapter	37
Index	41

Preface

The *Oracle Solaris Cluster 3.3 With Sun StorEdge A3500FC System Manual* provides procedures specific to Sun StorEdge A3500FC systems that are placed in an Oracle Solaris Cluster environment.

Use this manual with the Oracle Solaris Cluster 3.3 software. See the [“Revision History”](#) on [page 6](#) for a list of changes to this manual.

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 one chapter that consists of three major sections.

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

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

Section 3 describes how to maintain Sun StorEdge 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 Oracle Solaris Cluster 3.3 With Sun StorEdge A3500FC System Manual

Revision Date	New Information
January 2009	Updated link in preface for multiple versions of Sun Cluster documentation.
September 2010	Updated product name, removed old CLI commands.

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://docs.sun.com>. Documentation that is not available at <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
<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 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

Command	Function
showrev -p	Reports which patches are installed
prtdiag -v	Displays system diagnostic information
/usr/cluster/bin/clnode show-rev	Displays Oracle Solaris Cluster release and package version information for each node

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

Documentation, Support, and Training

See the following web sites for additional resources:

- [Documentation \(http://docs.sun.com\)](http://docs.sun.com)
- [Support \(http://www.oracle.com/us/support/systems/index.html\)](http://www.oracle.com/us/support/systems/index.html)
- [Training \(http://education.oracle.com\)](http://education.oracle.com) – Click the Sun link in the left navigation bar.

Oracle Welcomes Your Comments

Oracle welcomes your comments and suggestions on the quality and usefulness of its documentation. If you find any errors or have any other suggestions for improvement, go to <http://docs.sun.com> and click Feedback. Indicate the title and part number of the documentation along with the chapter, section, and page number, if available. Please let us know if you want a reply.

[Oracle Technology Network \(http://www.oracle.com/technetwork/index.html\)](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 \(http://www.oracle.com/technology/obe/start/index.html\)](http://www.oracle.com/technology/obe/start/index.html).
- Download [Sample Code \(http://www.oracle.com/technology/sample_code/index.html\)](http://www.oracle.com/technology/sample_code/index.html).

Typographic Conventions

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

TABLE P-4 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-5 Shell Prompts

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

Installing and Maintaining a Sun StorEdge A3500FC System

This chapter describes the procedures about how to install, configure, and maintain a Sun StorEdge A3500FC (Fibre Channel-based) system in an Oracle Solaris Cluster environment. For information about the SCSI-based version of this device, see the *Oracle Solaris Cluster 3.3 With StorEdge A1000 Array, Netra st A1000 Array, or StorEdge A3500 System Manual* manual.

This chapter contains the following main topics:

- “Installing Storage Systems” on page 11
- “Configuring Storage Systems” on page 20
- “Maintaining Storage Systems” on page 28

For information about how to use storage arrays in a storage area network (SAN), see “SAN Solutions in an Oracle Solaris Cluster Environment” in *Oracle Solaris Cluster 3.3 Hardware Administration Manual*.

Installing Storage Systems

This section contains procedures both for installing a storage system in a new cluster and for adding storage systems to an existing cluster.

TABLE 1-1 Task Map: Installing Storage Systems

Task	Information
Install a storage system to a new cluster, before the OS and Oracle Solaris Cluster software are installed.	“How to Install a Storage System in a New Cluster” on page 12
Add a storage system to an operational cluster.	“How to Add a Storage System to an Existing Cluster” on page 15

▼ How to Install a Storage System in a New Cluster

Use this procedure for an initial installation and an initial configuration, *before* you install the Solaris Operating System and Oracle Solaris Cluster software. If you need to add a storage system to an operational cluster, use the procedure in [“How to Add a Storage System to an Existing Cluster”](#) on page 15.

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.

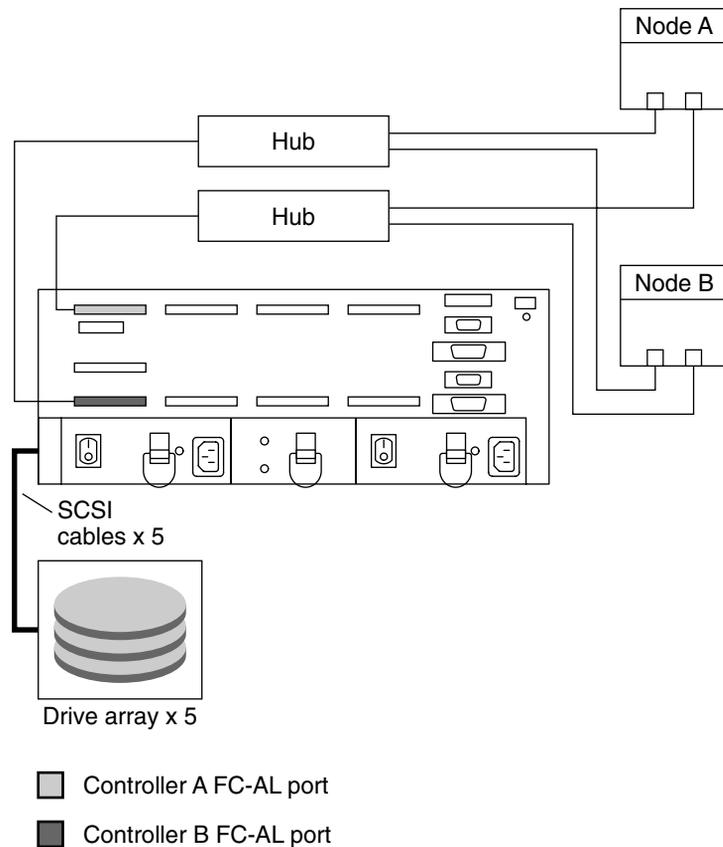
See [Figure 1–1](#) for an example about how to cable a StorEdge A3500FC system.

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 Solutions in an Oracle Solaris Cluster Environment”](#) in *Oracle Solaris Cluster 3.3 Hardware Administration Manual* for more information.

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



- 3 **Set the loop ID of the controller module 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.**

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

- 5 **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 Oracle Solaris 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 Oracle service providers and to customers with SunSolve service contracts at the SunSolve site: <http://sunsolve.sun.com>.

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

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

7 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*.

8 Check the storage system 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*.

9 Set the Rdac parameters in the `/etc/osa/rmparms` file.

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

10 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*.

11 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 Oracle Solaris logical device names.

12 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.

- 13** Ensure that the new logical name for the LUN that you created in [Step 11](#) appears in the `/dev/rdsk` 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 Oracle Solaris Cluster data services collection.

▼ How to Add a Storage System to an Existing Cluster

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

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.

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

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 Oracle Solaris 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 Oracle 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 Oracle Solaris Cluster release notes documentation.

3 Set the `Rdac` parameters in the `/etc/osa/rmparms` 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 If you are installing new host adapters in your nodes, shut down and power off Node A.

If you are not installing host adapters, skip to [Step 7](#).

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

6 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.

7 Cable the storage system to Node A. The method you choose to cable your system depends on the type of system you are adding.

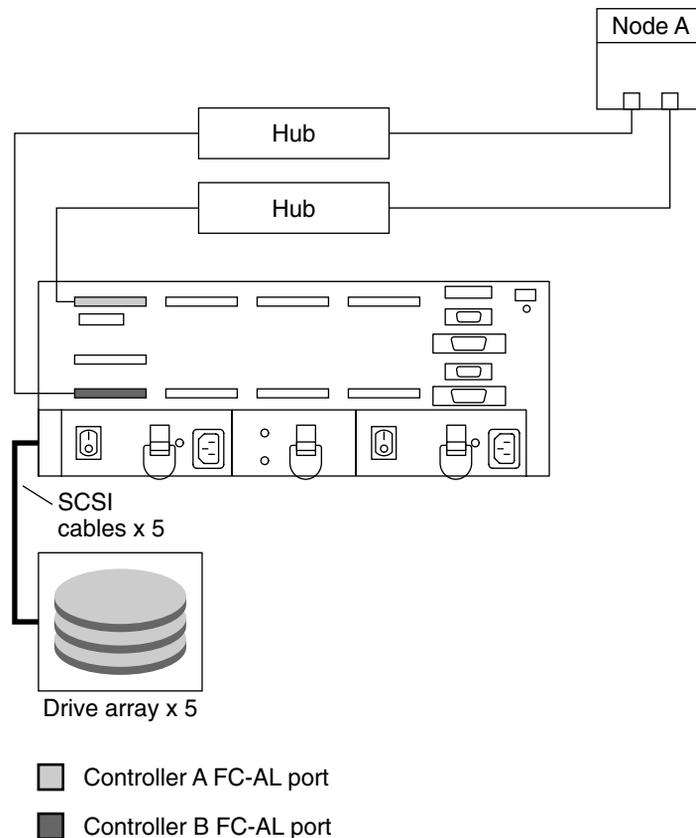
See [Figure 1–2](#) for an example about how to cable a StorEdge A3500FC storage system. In this example, Node A connects to a StorEdge A3500FC controller module.

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 Solutions in an Oracle Solaris Cluster Environment](#)” in *Oracle Solaris Cluster 3.3 Hardware Administration Manual* for more information. Sun SAN software does not support storage array.

FIGURE 1-2 Cabling a StorEdge A3500FC Storage System: An Example, Part 1



- 8 If you powered off Node A to install a host adapter, perform a reconfiguration boot to create the new Oracle Solaris device files and links.**

If you did not power off Node A, skip to [Step 9](#).

- 9 If you are installing new host adapters in Node B to connect Node B to the storage system, shut down and power off the node.**

If you are not installing host adapters, skip to [Step 11](#).

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

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

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

11 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.

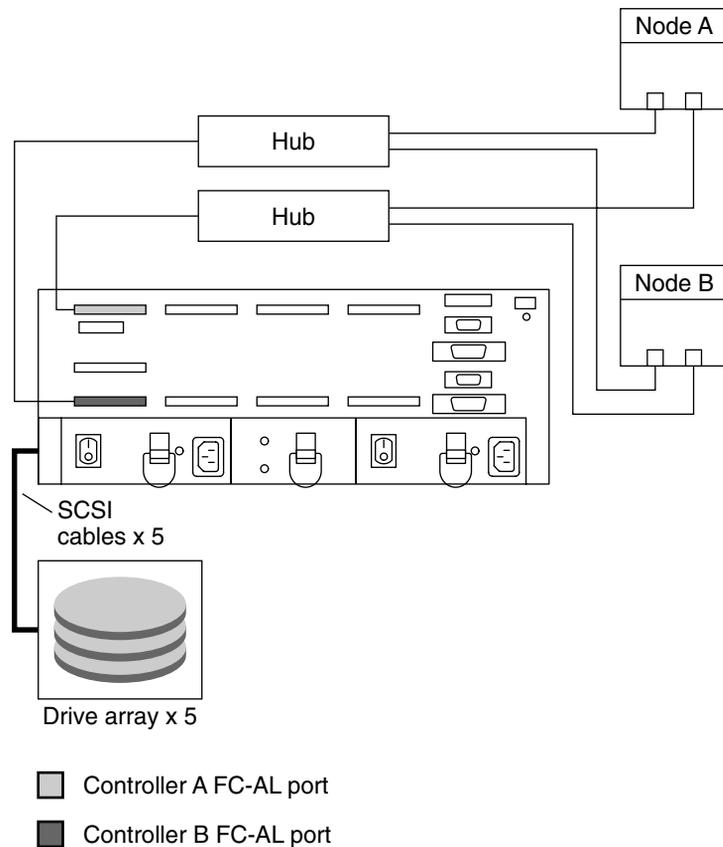
See [Figure 1-3](#) for an example about how to cable a StorEdge A3500FC storage system. In this example, two nodes connect to a StorEdge A3500FC controller module.

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 Solutions in an Oracle Solaris Cluster Environment](#)” in *Oracle Solaris Cluster 3.3 Hardware Administration Manual* for more information.

FIGURE 1-3 Cabling a StorEdge A3500FC Storage System: An Example, Part 2



- 12 If you powered off Node B to install a host adapter, power on the node and the storage system.

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

If you did not power off Node B, skip to [Step 15](#).

- 13 Verify that Node B recognizes the new host adapters and disk drives.
- 14 If the node does not recognize the new hardware, check all hardware connections and repeat installation steps you performed in [Step 10](#).

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

- 15 On one node, verify that the device IDs have been assigned to the LUNs for all nodes. These nodes are attached to the storage system.

```
# cldevice show
```

- 16 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 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 21.

To upgrade controller module firmware, see “[How to Upgrade Controller Module Firmware in a Running Cluster](#)” on page 34.

Configuring Storage Systems

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

To configure a 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 a Storage System

Task	Information
Create a logical unit number (LUN).	“How to Create a LUN” on page 21
Remove a LUN.	“How to Delete a LUN” on page 23
Reset the LUN configuration.	“How to Reset the LUN Configuration” on page 24

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

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

▼ How to Create a LUN

Use this procedure to create a logical unit 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 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.



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

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.

Before You Begin 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 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 Oracle 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.

```
# 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.

```
# cldevice show
```

6 If the device ID numbers that you received in [Step 5](#) are not the same for both nodes, perform the procedure in “[How to Correct Mismatched Device ID Numbers](#)” on [page 26](#) before you continue with [Step 7](#) of this procedure.

If the device ID numbers are the same, proceed to [Step 7](#).

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 Solaris Volume Manager or Veritas Volume Manager documentation.

▼ 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.

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.

Before You Begin 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 system, determine the paths to the LUN that you are deleting.**

```
# format
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 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
```

- 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.

```
# cldevice clear
```

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

```
# clnode evacuate nodename
```

9 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.

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

For more information about how to boot nodes, see [Chapter 3, “Shutting Down and Booting a Cluster,”](#) in *Oracle Solaris Cluster System Administration Guide*.

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 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 If a volume manager does not manage the LUNs on the controller module that you are resetting, proceed to [Step 4](#).**
- 3 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
```

- 4 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*.

- 5 Use the `format` command to 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 On both nodes, update device namespaces.

```
# devfsadm -C
```

8 On both nodes, remove all obsolete device IDs.

```
# cldevice clear
```

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

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

10 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.

11 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.
```

12 After the node reboots and joins the cluster, repeat [Step 6](#) through [Step 11](#) on the other node. This node is attached to the storage array 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 A3500FC LUNs. You correct the mismatch by deleting the Oracle Solaris OS and Oracle Solaris 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 21](#).

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.

Before You Begin

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 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 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, 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 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 nodename
```

- 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.

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

For more information about how to boot nodes, see [Chapter 3, “Shutting Down and Booting a Cluster,”](#) in *Oracle Solaris Cluster System Administration Guide*.

- 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 21.

Maintaining Storage Systems

This section contains the procedures about how to maintain a storage system in an Oracle Solaris Cluster environment.

Some maintenance procedures in [Table 1–3](#) are performed in the same way as in a noncluster environment. This section references these procedures, but this section does not contain these procedures. [Table 1–3](#) 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 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 Tasks: Maintaining a Storage System

Task	Information
Remove a storage system.	“How to Remove a Storage System” on page 30
Replace a failed controller module. Restore an offline controller module.	“How to Replace a Failed Controller or Restore an Offline Controller” on page 32
Upgrade controller module firmware and NVSRAM file.	“How to Upgrade Controller Module Firmware in a Running Cluster” on page 34
Add a disk drive.	“How to Add a Disk Drive in a Running Cluster” on page 35
Replace a disk drive.	“How to Replace a Failed Disk Drive in a Running Cluster” on page 36
Remove a disk drive.	“How to Remove a Disk Drive From a Running Cluster” on page 36
Upgrade disk drive firmware.	“How to Upgrade Disk Drive Firmware in a Running Cluster” on page 37

TABLE 1-3 Tasks: Maintaining a Storage System (Continued)

Task	Information
Replace a host adapter in a node.	“How to Replace a Host Adapter” on page 37

StorEdge A3500FC System FRUs

With the exception of one instruction, 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.

- 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 converters (GBIC) 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 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.

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.

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

1 Migrate any Oracle Real Application Clusters 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 If a volume manager manages any of the LUNs on the controller module 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 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.

```
# cldevice clear
```

10 If you are removing the *last* StorEdge A3500FC controller module from a hub or FC switch in your cluster, 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 Solutions in an Oracle Solaris Cluster Environment](#)” in *Oracle Solaris Cluster 3.3 Hardware Administration Manual* for more information.

If you are not removing the last controller module, skip to [Step 12](#).

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

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

12 Remove any unused host adapter from nodes that were attached to the 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.

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 system.****13 Switch the cluster back online.**

```
# clresourcegroup online +
```

14 If you are removing the *last* storage system from your cluster, remove the 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 Oracle Solaris Cluster concepts documentation.

Note – If you want to create a SAN by using two FC switches and Sun SAN software, see “[SAN Solutions in an Oracle Solaris Cluster Environment](#)” in *Oracle Solaris Cluster 3.3 Hardware Administration Manual* for more information.

- 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/rmparms` file to `false`.**



Caution – You must set the `System_LunReDistribution` parameter in the `/etc/raid/rmparms` 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 `rmparms` 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 If you have a failed controller, replace the failed controller with a new controller.

Note – *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.

If your controller module is offline, but does not have a failed controller, proceed to [Step 4](#).

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 repair device command on LUN 0 of the controller you want to bring online.

```
# cldevice repair
```

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 NVSRAM firmware. The method that you choose depends on your firmware.

1 Determine the correct procedure for your upgrade.

■ 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.

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 Oracle Solaris 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.

See Also To create LUNs for the new drives, see [“How to Create a LUN”](#) on page 21 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 If replacing the disk drive affects any LUN's availability, 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 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 If the failed drive does *not* belong to a device group, skip to [Step 6](#).**

- 5 If the failed drive belongs to 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 7](#).

- 6 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*.

- 7 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 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 23.

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 an 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

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

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

- Except for the failed host adapter, your cluster is operational and all nodes are powered on.
- Your nodes are not configured with dynamic reconfiguration functionality.

If your nodes are configured for dynamic reconfiguration **and** you are using two entirely separate hardware paths to your shared data, see the *Oracle Solaris Cluster 3.3 Hardware Administration Manual* and skip steps that instruct you to shut down the cluster.

You cannot replace a single, dual-port HBA that has quorum configured on that storage path by using DR. Follow all steps in the procedure. For the details on the risks and limitations of this configuration, see “Configuring Cluster Nodes With a Single, Dual-Port HBA” in *Oracle Solaris Cluster 3.3 Hardware Administration Manual*.

Exceptions to this restriction include three-node or larger cluster configurations where no storage device has a quorum device configured.

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.

1 Become superuser or assume a role that provides `solaris.cluster.read` and `solaris.cluster.modify` RBAC authorization.

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

Record this information because you use this information in [Step 13](#) and [Step 14](#) of this procedure to return resource groups and device groups to Node A.

```
# clresourcegroup status -n NodeA
# cldevicegroup status -n NodeA
```

-n *NodeA* The node for which you are determining resource groups and device groups.

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

```
# clnode evacuate nodename
```

4 Halt all I/O activity on the affected controller module.

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

5 Shut down Node A.

For the full procedure about how to shut down and power off a node, see [Chapter 3, “Shutting Down and Booting a Cluster,”](#) in *Oracle Solaris Cluster System Administration Guide*.

6 Power off Node A.**7 Replace the failed host adapter.**

To remove and add host adapters, see the documentation that shipped with your nodes.

8 If you need to upgrade the node's host adapter firmware, boot Node A into noncluster mode by adding -x to your boot instruction. Proceed to [Step 9](#).

If you do not need to upgrade firmware, skip to [Step 10](#).

9 Upgrade the host adapter firmware on Node A.

The Oracle Enterprise Manager Ops Center 2.5 software helps you patch and monitor your data center assets. Oracle Enterprise Manager Ops Center 2.5 helps improve operational efficiency and ensures that you have the latest software patches for your software. Contact your Oracle representative to purchase Oracle Enterprise Manager Ops Center 2.5.

Additional information for using the Oracle patch management tools is provided in *Oracle Solaris Administration Guide: Basic Administration* at <http://docs.sun.com> (<http://docs.sun.com>). Refer to the version of this manual for the Oracle Solaris OS release that you have installed.

If you must apply a patch when a node is in noncluster mode, you can apply it in a rolling fashion, one node at a time, unless instructions for a patch require that you shut down the entire cluster. Follow the procedures in “[How to Apply a Rebooting Patch \(Node\)](#)” in *Oracle Solaris Cluster System Administration Guide* to prepare the node and to boot it in noncluster mode. For ease of installation, consider applying all patches at the same time. That is, apply all patches to the node that you place in noncluster mode.

Index

A

- adding
 - See also* installing
 - disk drives, 35
 - systems, 15–20
- arrays, *See* storage arrays

B

- battery harness, replacing, 29
- battery units, replacing, 29

C

- cables
 - replacing, 29
 - fiber-optic cables, 29
- configuring, storage systems, 20–28
- controller module assemblies, replacing, 29
- controller module card cages, replacing, 29
- controller module fan canisters, replacing, 29
- controller module firmware, upgrading, 34–35
- controller module power cords, replacing, 29
- controllers
 - replacing, 32–33
 - restoring, 32–33
- correcting, DID numbers, 26–28
- creating
 - hot spares, 21
 - LUNs, 21–22

D

- deleting
 - See also* removing
 - hot spares, 21
 - LUNs, 23–24
- DID numbers, correcting, 26–28
- disk drive firmware, upgrading, 37
- disk drives
 - adding, 35
 - removing, 36–37
 - replacing, 36
- drive groups, expanding, 21

E

- expanding, drive groups, 21

F

- fan canisters
 - See* controller module fan canisters
 - See* power supply fan canisters
- firmware, upgrading, 34–35
- FRUs, 29–30

G

- GBICs
 - replacing, 29
- gigabit interface converters, *See* GBIC

H

HBAs, replacing, 37–39
help, 7–8
host adapters, replacing, 37–39
hot spares
 creating, 21
 deleting, 21
hubs, replacing, 29

I

installing
 See also adding
 storage systems, 11–20

L

logical unit numbers, *See* LUNs
logical units, LUNs, 21–22
logical volumes, *See* LUNs
LUNs
 creating, 21–22
 deleting, 23–24
 rebalancing, 21
 resetting configuration, 24–26

M

maintaining, storage systems, 28–39

O

Oracle Parallel Server, 5

P

power cords
 replacing, 29
 storage array
 replacing, 29

power supply fan canisters, replacing, 29
power supply housing, replacing, 29

R**RAID**

 hardware, 21
 software, 21
rebalancing, LUNs, 21
removing
 See also deleting
 disk drives, 36–37
 storage systems, 30–32

replacing

 battery harness, 29
 battery units, 29
 cables, 29
 controller module assemblies, 29
 controller module card cages, 29
 controller module fan canisters, 29
 controller module power cords, 29
 controllers, 32–33
 disk drives, 36
 fiber-optic cables, 29
 GBICs, 29
 host adapters, 37–39
 hubs, 29
 power cords, 29
 power supply fan canisters, 29
 power supply housing, 29
 SFPs, 29
 storage array power cords, 29
 switches, 30

resetting, LUN configurations, 24–26

restoring, controllers, 32–33

S**SFPs**

 replacing, 29
Small Form-Factor Pluggable, *See* SFP
storage systems
 configuring, 20–28

storage systems (*Continued*)

installing, 11–20

maintaining, 28–39

removing, 30–32

switches, replacing, 30

systems*See* storage systems

adding, 15–20

T

technical support, 7–8

U

upgrading

controller module firmware, 34–35

disk drive firmware, 37

