

Oracle® Solaris Cluster 3.3 With Sun StorEdge 6130 Array Manual

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 Restrictions and Requirements	11
Requirements	11
Restrictions	11
2 Installing and Configuring a Sun StorEdge 6130 Array	13
Installing Storage Arrays	13
Storage Array Cabling Configurations	13
How to Install Storage Arrays in a New Cluster	14
▼ How to Add Storage Arrays to an Existing Cluster	16
Configuring Storage Arrays	17
▼ How to Create a Logical Volume	17
▼ How to Remove a Logical Volume	19
3 Maintaining a Sun StorEdge 6130 Array	23
FRUs That <i>Do Not</i> Require Oracle Solaris Cluster Maintenance Procedures	23
Maintaining Storage Arrays	24
▼ How to Upgrade Storage Array Firmware	24
▼ How to Remove a Storage Array	25
Replacing a Node-to-Switch Component	26
▼ How to Replace a Host Adapter	28
▼ How to Replace a Disk Drive	30
Index	33

Preface

The *Oracle Solaris Cluster 3.3 With Sun StorEdge 6130 Array Manual* provides procedures that are specific to Oracle's Sun StorEdge 6130 array that is placed in an Oracle Solaris Cluster environment.

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

This book assumes that you are performing one or more of the following tasks:

- You want to replace an array component to *prevent* a failure.
- You want to replace an array component because you have an *existing* failure.
- You want to add (to an *established* cluster) or install (to a *new* cluster) a storage array.

Note – This Oracle Solaris Cluster release supports systems that use the SPARC and x86 families of processor architectures: UltraSPARC, SPARC64, and AMD64. In this document, the label x86 refers to systems that use the AMD64 family of processor architectures.

The information in this document pertains to both platforms unless otherwise specified in a special chapter, section, note, bulleted item, figure, table, or example.

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 presales guide. You should have already determined your system requirements and purchased the appropriate equipment and software before reading this document.

How This Book Is Organized

This book contains the following chapters:

- [Chapter 1, “Restrictions and Requirements,”](#) lists limitations of your use of Sun StorEdge 6130 storage arrays in an Oracle Solaris Cluster environment.
- [Chapter 2, “Installing and Configuring a Sun StorEdge 6130 Array,”](#) discusses how to install Sun StorEdge 6130 storage arrays and how to configure logical units on them.
- [Chapter 3, “Maintaining a Sun StorEdge 6130 Array,”](#) describes how to maintain Sun StorEdge 6130 storage arrays in a running cluster.

Revision History

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

TABLE P-1 Oracle Solaris Cluster 3.3 With Sun StorEdge 6130 Array Manual

Revision Date	New Information
January 2009	Updated links in the Preface to lead to different versions of Sun Cluster documentation.
September 2010	Updated book for new product name and removed old CLI commands and procedures.

Related Documentation

The following books provide conceptual information or procedures to administer hardware and applications. If you plan to use this documentation in a 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>Sun StorEdge 6130 Array Getting Started Guide</i>	817-0032
<i>Sun StorEdge 6130 Array Site Preparation Guide</i>	819-0033

TABLE P-2 Hardware Documentation (Continued)

Title	Part Number
<i>Sun StorEdge 6130 Regulatory and Safety Compliance Manual</i>	819-0035
<i>Sun StorEdge 6130 Array Release Notes</i>	819-0034
<i>Sun StorEdge Expansion Cabinet Installation and Service Manual</i>	805-3067
<i>Sun Rack Installation Guide</i>	816-6386
<i>Sun Fire Cabinet Installation and Reference Manual</i>	806-2942
<i>Solaris Fibre Channel Storage Configuration and Multipathing Support Guide</i>	819-0139

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
<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, 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>

Restrictions and Requirements

This chapter includes only restrictions and requirements that have a direct impact on the procedures in this book. The *Oracle Solaris Cluster 3.x Configuration Guide* lists all restrictions and support information. This manual is not directly available to customers. However, your Oracle service provider has access to this manual. For general support information, contact your Oracle service provider.

Requirements

Storage devices that support logical unit number (LUN) masking must implement the LUN masking so that LUNs are *not* visible to the Sun StorEdge 6130 storage system.

Restrictions

When using storage-based replication, do not configure a replicated volume as a quorum device. Locate any quorum devices on an unreplicated volume. See “[Using Storage-Based Data Replication Within a Cluster](#)” in *Oracle Solaris Cluster System Administration Guide* for more information on storage-based replication.

Installing and Configuring a Sun StorEdge 6130 Array

This chapter contains the procedures about how to install and configure a Sun StorEdge 6130 array. These procedures are specific to an Oracle Solaris Cluster environment.

This chapter contains the following main topics:

- “Installing Storage Arrays” on page 13
- “Configuring Storage Arrays” on page 17

For detailed information about storage array architecture, features, configuration utilities, and installation, see the *Sun StorEdge 6130 Array Getting Started Guide*.

For a URL to this storage documentation, see “[Related Documentation](#)” on page 6.

Installing Storage Arrays

This section contains the procedures listed in [Table 2-1](#).

TABLE 2-1 Task Map: Installing Storage Arrays

Task	Information
Install a storage array in a new cluster, before the OS and Oracle Solaris Cluster software are installed.	“ How to Install Storage Arrays in a New Cluster ” on page 14
Add a storage array to an existing cluster.	“ How to Add Storage Arrays to an Existing Cluster ” on page 16

Storage Array Cabling Configurations

You can install your storage array in several different configurations: [Figure 2-1](#) and [Figure 2-2](#) are two examples.

FIGURE 2-1 Direct-Connect Configuration

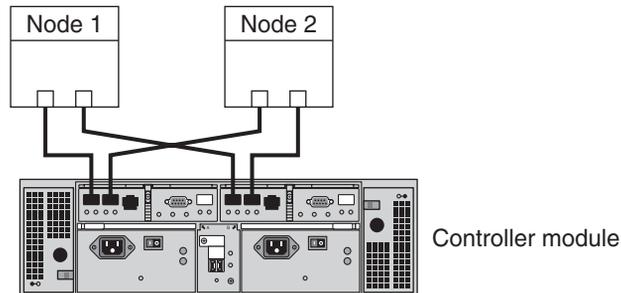
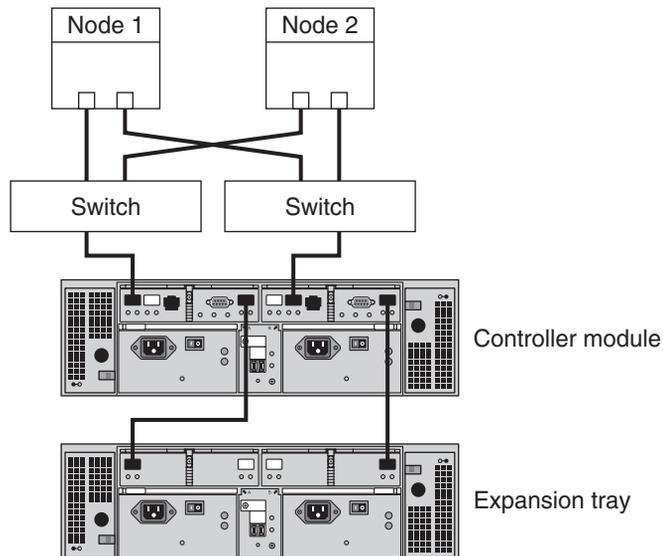


FIGURE 2-2 Switch Configuration



How to Install Storage Arrays in a New Cluster

Use this procedure to install a storage array in a new cluster. To add a storage array to an existing cluster, use the procedure in [“How to Add Storage Arrays to an Existing Cluster”](#) on page 16.

This procedure relies on the following assumptions:

- You have *not* installed the Oracle Solaris Operating System.
- You have *not* installed the Oracle Solaris Cluster software.

- Your cluster has *not* been established.
- You have enough host adapters to connect the nodes and the storage array to the service processor panel.
If you need to install host adapters, see [“How to Replace a Host Adapter”](#) on page 28. When this procedure asks you to replace the failed host adapter, simply install the new host adapter instead.

▼ Install and Cable the Hardware

1 Unpack, place, and level the storage array.

For instructions, see the *Sun StorEdge 6130 Array Getting Started Guide*. For a list of StorEdge documentation, see [Table P-2](#).

2 (Optional) Install the Fibre Channel (FC) switch for the storage array if you do not have a switch installed.

For the procedure about how to install an FC switch, see the documentation that shipped with your FC switch hardware.

3 Connect the nodes to the service processor panel.

- *(SAN Configuration)* Connect the FC switches to the service processor panel.
- *(Direct-Attached Configuration)* Connect each node to the host ports on the array.

For instructions, see the *Sun StorEdge 6130 Array Getting Started Guide*.

4 Install the storage array.

For instructions, see the *Sun StorEdge 6130 Array Getting Started Guide*.

5 Power on the storage array and the nodes.

For instructions, see the *Sun StorEdge 6130 Array Getting Started Guide*.

▼ Install the Oracle Solaris Operating System and Configure Multipathing

- On all nodes, install the Oracle Solaris operating system and apply the required Solaris patches for Oracle Solaris Cluster software and storage array support.

For the procedure about how to install the Oracle Solaris operating environment, see [“How to Install Solaris Software”](#) in *Oracle Solaris Cluster Software Installation Guide*.

See Also ▪ To create a logical volume, see [“How to Create a Logical Volume”](#) on page 17.

- To continue with Oracle Solaris Cluster software installation tasks, see your Oracle Solaris Cluster software installation documentation.

▼ How to Add Storage Arrays to an Existing Cluster

Use this procedure to add a new storage array to a *running* cluster. To install a new storage array in a Oracle Solaris Cluster configuration that is not running, use the procedure in “[How to Install Storage Arrays in a New Cluster](#)” on page 14.

Before You Begin This procedure relies on the following assumptions:

- (*Veritas Volume Manager Only*) You have a version of Veritas Volume Manager that includes Array Support Library (ASL). For more information on this requirement, see “[Requirements](#)” on page 11.

- You have enough host adapters to connect the nodes and the storage array to the service processor panel.

If you need to install host adapters, see “[How to Replace a Host Adapter](#)” on page 28. When this procedure asks you to replace the failed host adapter, simply install the new host adapter instead.

- You do *not* need to install a new node in your cluster.

If you need to add a node to your cluster, see your Oracle Solaris Cluster system administration documentation. Ensure that you install the required Oracle Solaris patches for storage array support.

1 Unpack, place, and level the storage array.

For instructions, see the *Sun StorEdge 6130 Array Getting Started Guide*. For a list of StorEdge documentation, see [Table P-2](#).

2 (Optional) Install the Fibre Channel (FC) switch for the storage array if you do not have a switch installed.

For the procedure about how to install an FC switch, see the documentation that shipped with your FC switch hardware.

3 Connect the nodes to the service processor panel.

- (*SAN Configuration*) **Connect the FC switches to the service processor panel.**
- (*Direct-Attached Configuration*) **Connect each node to the host ports on the array.**

For instructions, see the *Sun StorEdge 6130 Array Getting Started Guide*.

4 Install the storage array.

For instructions, see the *Sun StorEdge 6130 Array Getting Started Guide*.

5 Power on the storage array and the nodes.

For instructions, see the *Sun StorEdge 6130 Array Getting Started Guide*.

- See Also**
- To create a logical volume, see [“How to Create a Logical Volume”](#) on page 17.
 - If you added a new node, perform the steps in [“Install the Oracle Solaris Operating System and Configure Multipathing”](#) on page 15 as appropriate. Apply the correct patches for storage array support enable multipathing.

Configuring Storage Arrays

This section contains the procedures about how to configure a storage array in a running cluster. [Table 2–2](#) lists these procedures.

TABLE 2–2 Task Map: Configuring a Storage Array

Task	Information
Create a logical volume.	“How to Create a Logical Volume” on page 17
Remove a logical volume.	“How to Remove a Logical Volume” on page 19

The following is a list of administrative tasks that require no cluster-specific procedures. See the storage array's online help for the following procedures.

- Creating a storage pool
- Removing a storage pool
- Creating a volume group
- Removing a volume group
- Creating an initiator group
- Adding an initiator group
- Removing an initiator group

▼ How to Create a Logical Volume

Use this procedure to create a logical volume from unassigned storage capacity.

Note – Oracle's Sun storage documentation uses the following terms:

- Logical volume
- Logical device
- Logical unit number (LUN)

This manual uses *logical volume* to refer to all such logical constructs.

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

- All nodes are booted in cluster mode and attached to the storage device.
- The storage device is installed and configured. If you are using multipathing, the storage device is configured as described in the installation procedure.
- If you are using Solaris I/O multipathing (MPxIO) for the Oracle Solaris 10 OS, previously called Sun StorEdge Traffic Manager in the Solaris 9 OS, verify that the paths to the storage device are functioning. To configure multipathing, see the *Solaris Fibre Channel Storage Configuration and Multipathing Support Guide*.

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.modify` role-based access control (RBAC) authorization.**
- 2 Follow the instructions in your storage device's documentation to create and map the logical volume. For a URL to this storage documentation, see [“Related Documentation” on page 6](#).**
 - Completely set up the logical volume. When you are finished, the volume must be created, mapped, mounted, and initialized.
 - If necessary, partition the volume.
 - To allow multiple clusters and nonclustered nodes to access the storage device, create initiator groups by using LUN masking.
- 3 If you are not using multipathing, skip to [Step 5](#).**
- 4 If you are using multipathing, and if any devices that are associated with the volume you created are at an unconfigured state, configure the multipathing paths on each node that is connected to the storage device.**

To determine whether any devices that are associated with the volume you created are at an unconfigured state, use the following command.

```
# cfgadm -al | grep disk
```

Note – To configure the Oracle Solaris I/O multipathing paths on each node that is connected to the storage device, use the following command.

```
# cfgadm -o force_update -c configure controllerinstance
```

To configure multipathing, see the *Solaris Fibre Channel Storage Configuration and Multipathing Support Guide*.

- 5 **On one node that is connected to the storage device, use the `format` command to label the new logical volume.**
- 6 **From any node in the cluster, update the global device namespace.**

```
# cldevice populate
```

Note – You might have a volume management daemon such as `vold` running on your node, and have a DVD drive connected to the node. Under these conditions, a device busy error might be returned even if no disk is inserted in the drive. This error is expected behavior. You can safely ignore this error message.

- 7 **To manage this volume with volume management software, use Solaris Volume Manager or Veritas Volume Manager commands to update the list of devices on all nodes that are attached to the new volume that you created.**

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

- See Also**
- To configure a logical volume as a quorum device, see Chapter 6, “Administering Quorum,” in *Oracle Solaris Cluster System Administration Guide*.
 - To create a new resource or configure a running resource to use the new logical volume, see Chapter 2, “Administering Data Service Resources,” in *Oracle Solaris Cluster Data Services Planning and Administration Guide*.

▼ How to Remove a Logical Volume

Use this procedure to remove a logical volume. This procedure defines Node A as the node with which you begin working.

Note – Sun storage documentation uses the following terms:

- Logical volume
- Logical device
- Logical unit number (LUN)

This manual uses *logical volume* to refer to all such logical constructs.

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

- All nodes are booted in cluster mode and attached to the storage device.
- The logical volume and the path between the nodes and the storage device are both operational.

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 Identify the logical volume that you are removing.

Refer to your Solaris Volume Manager or Veritas Volume Manager documentation for more information.

3 (Optional) Migrate all data off the logical volume that you are removing. Alternatively, back up that data.

4 If the LUN that you are removing is configured as a quorum device, choose and configure another device as the quorum device. Then remove the old quorum device.

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

```
# clquorum show
```

For procedures about how to add and remove quorum devices, see [Chapter 6, “Administering Quorum,”](#) in *Oracle Solaris Cluster System Administration Guide*.

5 If you are using volume management software, use that software to update the list of devices on all nodes that are attached to the logical volume that you are removing.

For instructions about how to update the list of devices, see your Solaris Volume Manager or Veritas Volume Manager documentation.

- 6 If you are using volume management software, run the appropriate Solaris Volume Manager or Veritas Volume Manager commands to remove the logical volume from any diskset or disk group.**

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

Note – Volumes that were managed by Veritas Volume Manager must be completely removed from Veritas Volume Manager control before you can delete them from the Oracle Solaris Cluster environment. After you delete the volume from any disk group, use the following commands on both nodes to remove the volume from Veritas Volume Manager control.

```
# vxdisk offline Accessname
# vxdisk rm Accessname

Accessname    Disk access name
```

- 7 If you are using multipathing, unconfigure the volume in Solaris I/O multipathing.**

```
# cfgadm -o force_update -c unconfigure Logical_Volume
```

- 8 Access the storage device and remove the logical volume.**

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

- 9 Determine the resource groups and device groups that are running on all nodes.**

Record this information because you use it in [Step 14](#) and [Step 15](#) of this procedure to return resource groups and device groups to these nodes.

Use the following command:

```
# clresourcegroup status +
# cldevicegroup status +
```

- 10 Move all resource groups and device groups off Node A.**

```
# clnode evacuate nodename
```

- 11 Shut down and reboot Node A.**

To shut down and boot a node, see [Chapter 3, “Shutting Down and Booting a Cluster,”](#) in *Oracle Solaris Cluster System Administration Guide*.

- 12 On Node A, remove the paths to the logical volume that you removed. Remove obsolete device IDs.**

```
# devfsadm -C
# cldevice clear
```


Maintaining a Sun StorEdge 6130 Array

This chapter contains the procedures about how to maintain a Sun StorEdge 6130 array. These procedures are specific to an Oracle Solaris Cluster environment.

This chapter contains the following procedures:

- “How to Upgrade Storage Array Firmware” on page 24
- “How to Remove a Storage Array” on page 25
- “Replacing a Node-to-Switch Component” on page 26
- “How to Replace a Host Adapter” on page 28
- “How to Replace a Disk Drive” on page 30

For detailed information about storage array architecture, features, and configuration utilities, see the *Sun StorEdge 6130 Array Getting Started Guide*.

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

In general, the following is a list of administrative tasks that require no cluster-specific procedures. See the base-product documentation for these procedures.

Note – Depending on your configuration type and the state of your cluster, a few of the following FRUs might require cluster-specific steps. Such FRUs include the DSP and service processor panel.

- Adding a disk drive.
- Replacing a storage array's chassis.
- Replacing an Ethernet cable.
- Replacing an Ethernet hub in a cabinet.
- Replacing an expansion cabinet.
- Replacing a fan assembly.

- Replacing the power cable on the storage array.
- Replacing a power and cooling unit (PCU).
- Replacing a controller.
- Replacing a battery.

Before removing a disk drive, stop all I/O to the storage array. There are no other cluster-specific steps.

Maintaining Storage Arrays

This section contains the procedures about how to maintain a storage system in a running cluster. [Table 3–1](#) lists these procedures.

TABLE 3–1 Task Map: Maintaining a Storage System

Task	Information
Remove a storage array.	“How to Remove a Storage Array” on page 25
Upgrade storage array firmware.	“How to Upgrade Storage Array Firmware” on page 24
Replace a node-to-switch component.	“How to Replace a Node-to-Switch Component in a Cluster Without Multipathing” on page 27
Replace a node's host adapter.	“How to Replace a Host Adapter” on page 28
Replace a disk drive.	“How to Replace a Disk Drive” on page 30
Add a node to the storage array.	Oracle Solaris Cluster system administration documentation
Remove a node from the storage array.	Oracle Solaris Cluster system administration documentation

▼ How to Upgrade Storage Array Firmware

Use this procedure to upgrade storage array firmware in a running cluster. Storage array firmware includes controller firmware, unit interconnect card (UIC) firmware, EPROM firmware, and disk drive firmware.

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.

- 1 **Stop all I/O to the storage arrays you are upgrading.**
- 2 **Apply the controller, disk drive, and loop-card firmware patches by using the arrays' GUI tools.**
For specific instructions, see your storage array's documentation.
- 3 **Confirm that all storage arrays that you upgraded are visible to all nodes.**
`# luxadm probe`
- 4 **Restart all I/O to the storage arrays.**
You stopped I/O to these storage arrays in [Step 1](#).

▼ How to Remove a Storage Array

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

This procedure defines Node N as the node that is connected to the storage array you are removing and the node with which you begin working.



Caution – During this procedure, you lose access to the data that resides on the storage array that 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.

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

- 1 **If necessary, back up all database tables, data services, and volumes that are associated with each partner group that is affected.**

- 2 **Remove references to the volumes that reside on the storage array that you are removing.**
For more information, see your Solaris Volume Manager or Veritas Volume Manager documentation.
- 3 **Disconnect the cables that connected Node N to the FC switches in your storage array.**
- 4 **On all nodes, remove the obsolete Oracle Solaris links and device IDs.**

```
# devfsadm -C  
# cldevice cclear
```
- 5 **Repeat [Step 3](#) through [Step 4](#) for each node that is connected to the storage array.**

Replacing a Node-to-Switch Component

Use this procedure to replace a node-to-switch component that has failed or that you suspect might be contributing to a problem.

Note – Node-to-switch components that are covered by this procedure include the following components:

- Node-to-switch fiber-optic cables
- Gigabit interface converters (GBICs) or small form-factor pluggables (SFPs) on an FC switch
- FC switches

To replace a host adapter, see [“How to Replace a Host Adapter”](#) on page 28.

This procedure defines Node A as the node that is connected to the node-to-switch component that you are replacing. This procedure assumes that, except for the component you are replacing, your cluster is operational.

Ensure that you are following the appropriate instructions:

- If your cluster uses multipathing, see [“How to Replace a Node-to-Switch Component in a Cluster That Uses Multipathing”](#) on page 27.
- If your cluster does *not* use multipathing, see [“How to Replace a Node-to-Switch Component in a Cluster Without Multipathing”](#) on page 27.

▼ How to Replace a Node-to-Switch Component in a Cluster That Uses Multipathing

- 1 If your configuration is active-passive, and if the active path is the path that needs a component replaced, make that path passive.
- 2 Replace the component.
Refer to your hardware documentation for any component-specific instructions.
- 3 (Optional) If your configuration is active-passive and you changed your configuration in [Step 1](#), switch your original data path back to active.

▼ How to Replace a Node-to-Switch Component in a Cluster Without Multipathing

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.

- 1 Become superuser or assume a role that provides `solaris.cluster.read` and `solaris.cluster.modify` RBAC authorization.
- 2 If the physical data path has failed, do the following:
 - a. Replace the component.
 - b. Fix the volume manager error that was caused by the failed data path.
 - c. (Optional) If necessary, return resource groups and device groups to this node.

You have completed this procedure.

- 3 If the physical data path has not failed, determine the resource groups and device groups that are running on Node A.

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

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

- 4 Move all resource groups and device groups to another node.

```
# clnode evacuate nodename
```

- 5 Replace the node-to-switch component.

Refer to your hardware documentation for any component-specific instructions.

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 10](#) and [Step 11](#) 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 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*.

5 Power off Node A.

6 Replace the failed host adapter.

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

7 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 8](#).

If you do not need to upgrade firmware, skip to [Step 9](#).

8 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.

- 2 If the failed disk drive affect the storage array logical volume's availability, If yes, use volume manager commands to detach the submirror or plex.**

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

- 3 If the logical volume (in [Step 1](#)) is configured as a quorum device, choose another volume to configure as the quorum device. Then remove the old quorum device.**

To determine whether the LUN is configured as a quorum device, use one of the following commands.

```
# clquorum show
```

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

- 4 Replace the failed disk drive.**
 - 5 (Optional) If the new disk drive is part of a logical volume that you want to be a quorum device, add the quorum device.**
- To add a quorum device, see your Oracle Solaris Cluster system administration documentation.
- 6 If you detached a submirror or plex in [Step 1](#), use volume manager commands to reattach the submirror or plex.**

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

Index

A

adding

See also installing

disk drives, 23

initiator groups, 17

storage arrays, 16–17

arrays, *See* storage arrays

B

batteries, replacing, 24

C

cables, replacing, 23

chassis, replacing, 23

controllers, replacing, 24

cooling units, replacing, 24

creating

initiator groups, 17

logical volumes, 17–19

storage pools, 17

volume groups, 17

D

deleting logical volumes, 19–22

disk drives

adding, 23

removing, 24

disk drives (*Continued*)

replacing, 30–31

E

Ethernet cables, replacing, 23

Ethernet hubs, replacing, 23

expansion cabinets, replacing, 23

F

fan assemblies, replacing, 23

firmware, upgrade storage array firmware, 24–25

FRUs, 23

H

HBAs, replacing, 28–30

help, 7–8

host adapters, replacing, 28–30

I

initiator groups, 17

installing

See also adding

storage arrays, 14–16

L

logical devices, *See* logical volumes
logical unit numbers, *See* logical volumes
logical volumes
 creating, 17–19
 removing, 19–22
LUN masking, 11
LUNs, *See* logical volumes

M

midplane, replacing, 23
modifying, initiator groups, 17

N

node-to-switch components, replacing, 26–28
nodes, adding and removing, 24

P

pools, 17
power cables, replacing, 24
power units, replacing, 24

R

removing
 disk drives, 24
 initiator groups, 17
 logical volumes, 19–22
 nodes, 24
 storage arrays, 25–26
 storage pools, 17
 volume groups, 17
replacing
 batteries, 24
 chassis, 23
 controllers, 24
 cooling units, 24
 disk drives, 30–31

replacing (*Continued*)

 Ethernet cables, 23
 Ethernet hubs, 23
 expansion cabinets, 23
 fan assemblies, 23
 host adapters, 28–30
 midplane, 23
 node-to-switch components, 26–28
 power cables, 24
 power units, 24

S

storage array firmware, upgrading, 24–25
storage arrays
 adding, 16–17
 installing, 14–16
 removing, 25–26
storage-based replication, 11
storage pools, 17
systems, *See* storage arrays

T

technical support, 7–8

U

upgrading, storage array firmware, 24–25

V

volume groups, 17