

# **Oracle® Solaris Cluster 3.3 With Sun StorEdge 6120 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

---

<b>Preface</b> .....	5
<b>1 Installing and Maintaining a Sun StorEdge 6120 Array</b> .....	11
Installing Storage Arrays .....	11
Storage Array Cabling Configurations .....	12
▼ How to Install a Single-Controller Configuration in a New Cluster .....	16
▼ How to Install a Dual-Controller Configuration in a New Cluster .....	18
How to Add a Single-Controller Configuration to an Existing Cluster .....	21
How to Add a Dual-Controller Configuration to an Existing Cluster .....	26
Configuring Storage Arrays .....	31
▼ How to Create a Logical Volume .....	31
▼ How to Remove a Logical Volume .....	33
Maintaining Storage Arrays .....	36
StorEdge 6120 Array FRUs .....	37
▼ How to Upgrade Storage Array Firmware .....	37
▼ How to Remove a Single-Controller Configuration .....	38
▼ How to Remove a Dual-Controller Configuration .....	41
Replacing a Node-to-Switch Component .....	44
▼ How to Replace a Host Adapter .....	46
<b>Index</b> .....	49



# Preface

---

The *Oracle Solaris Cluster 3.3 With Sun StorEdge 6120 Array Manual* provides procedures specific to Oracle's Sun StorEdge 6120 arrays in dual-controller configurations and in single-controller configurations that are placed in an Oracle Solaris Cluster environment.

Use this manual with any version of Oracle Solaris Cluster 3.3 software on SPARC based clusters or x86 based clusters.

---

**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.

---

See the “[Revision History](#)” on page 6 for a list of changes to this manual.

## 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 one chapter with the following major sections.

Section 1 discusses how to install Sun StorEdge 6120 storage arrays.

Section 2 discusses how to configure logical units on Sun StorEdge 6120 storage arrays.

Section 3 describes how to maintain Sun StorEdge 6120 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 6120 Array Manual

Revision Date	New Information
January 2009	Updated link to different versions of Sun Cluster documentation in Preface. Added updated information for Solaris I/O multipathing (MPxIO) for the Solaris 10 OS.
September 2010	Update book for new product name and 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>Sun StorEdge Traffic Manager Software Release Notes</i> Available on <a href="http://www.sun.com/storage/san/">http://www.sun.com/storage/san/</a>	817-0385
<i>Sun StorEdge Traffic Manager Installation and Configuration Guide</i> Available on <a href="http://www.sun.com/products-n-solutions/hardware/docs">http://www.sun.com/products-n-solutions/hardware/docs</a>	816-1420

TABLE P-2 Hardware Documentation (Continued)

Title	Part Number
<i>Sun StorEdge 6120 Array Release Notes</i>	817-0201
Available on <a href="http://www.sun.com/products-n-solutions/hardware/docs">http://www.sun.com/products-n-solutions/hardware/docs</a>	
<i>Sun StorEdge 6020 and 6120 Array System Manual</i>	817-0200
Available on <a href="http://www.sun.com/products-n-solutions/hardware/docs">http://www.sun.com/products-n-solutions/hardware/docs</a>	
<i>Sun StorEdge 6120 Array Installation Guide</i>	817-0199
Available on <a href="http://www.sun.com/products-n-solutions/hardware/docs">http://www.sun.com/products-n-solutions/hardware/docs</a>	

TABLE P-3 Oracle Solaris Cluster and Sun Cluster Documentation

Documentation
<a href="#">Oracle Solaris Cluster 3.3</a>
<a href="#">Sun Cluster 3.2</a>
<a href="#">Sun Cluster 3.1</a>

## 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>) 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.
<b>AaBbCc123</b>	What you type, contrasted with onscreen computer output	<code>machine_name%</code> <b>su</b> 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. <b>Note:</b> 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 6120 Array

---

This chapter contains the procedures about how to install, configure, and maintain a Sun StorEdge 6120 array in dual-controller configurations and in single-controller configurations. These procedures are specific to a Sun StorEdge 6120 array in an Oracle Solaris Cluster environment.

This chapter contains the following main topics:

- [“Installing Storage Arrays” on page 11](#)
- [“Configuring Storage Arrays” on page 31](#)
- [“Maintaining Storage Arrays” on page 36](#)

## Installing Storage Arrays

This section contains the procedures for installing single-controller and dual-controller storage array configurations in new and existing Oracle Solaris Cluster configurations. The following table lists these procedures.

TABLE 1-1 Task Map: Installing a Storage Array

Task	Information
Installing arrays in a new cluster, using a single-controller configuration	<a href="#">“How to Install a Single-Controller Configuration in a New Cluster” on page 16</a>
Installing arrays in a new cluster, using a dual-controller configuration	<a href="#">“How to Install a Dual-Controller Configuration in a New Cluster” on page 18</a>
Adding arrays to an existing cluster, using a single-controller configuration.	<a href="#">“How to Add a Single-Controller Configuration to an Existing Cluster” on page 21</a>

TABLE 1-1 Task Map: Installing a Storage Array (Continued)

Task	Information
Adding arrays to an existing cluster, using a dual-controller configuration.	<a href="#">“How to Add a Dual-Controller Configuration to an Existing Cluster” on page 26</a>

## Storage Array Cabling Configurations

You can install your storage array in several different configurations. Use the *Sun StorEdge 6120 Array Installation Guide* to evaluate your needs and determine which configuration is best for your situation.

The following figures illustrate example configurations.

[Figure 1-1](#) shows two storage arrays, and two have controllers. The storage arrays connect to a two-node cluster through two switches. Dual-controller configurations require software RAID-1 (host-based mirroring).

FIGURE 1-1 Installing a 1x1 Configuration With Software RAID-1

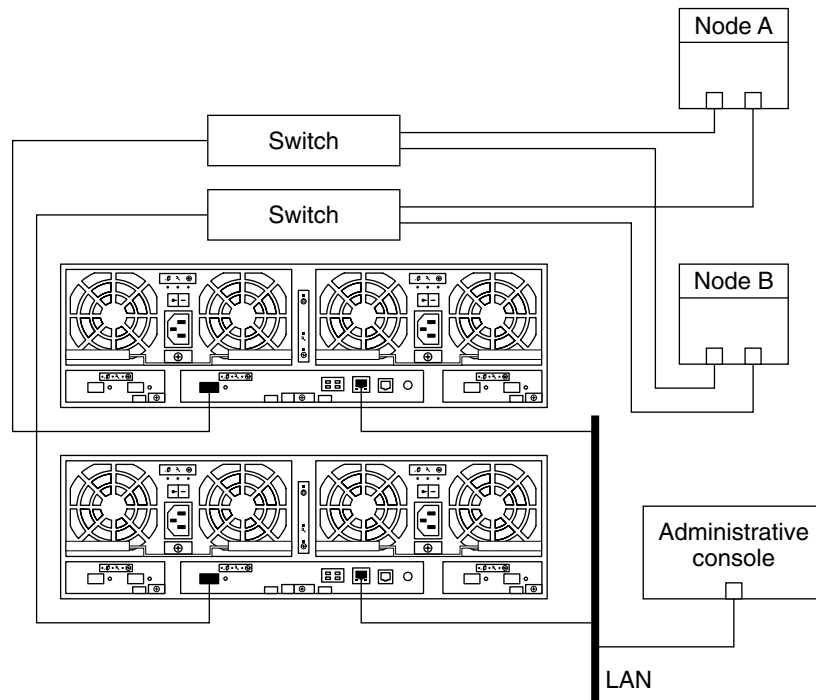


Figure 1-2 shows four storage arrays, and two have controllers. The first storage array without a controller connects to the second storage array, which has a controller. The third storage array without a controller connects to the fourth storage array, which has a controller. The two storage arrays with controllers connect to a two-node cluster through two switches. Dual-controller configurations require software RAID-1 (host-based mirroring).

FIGURE 1-2 Installing a 1x2 Configuration With Software RAID-1

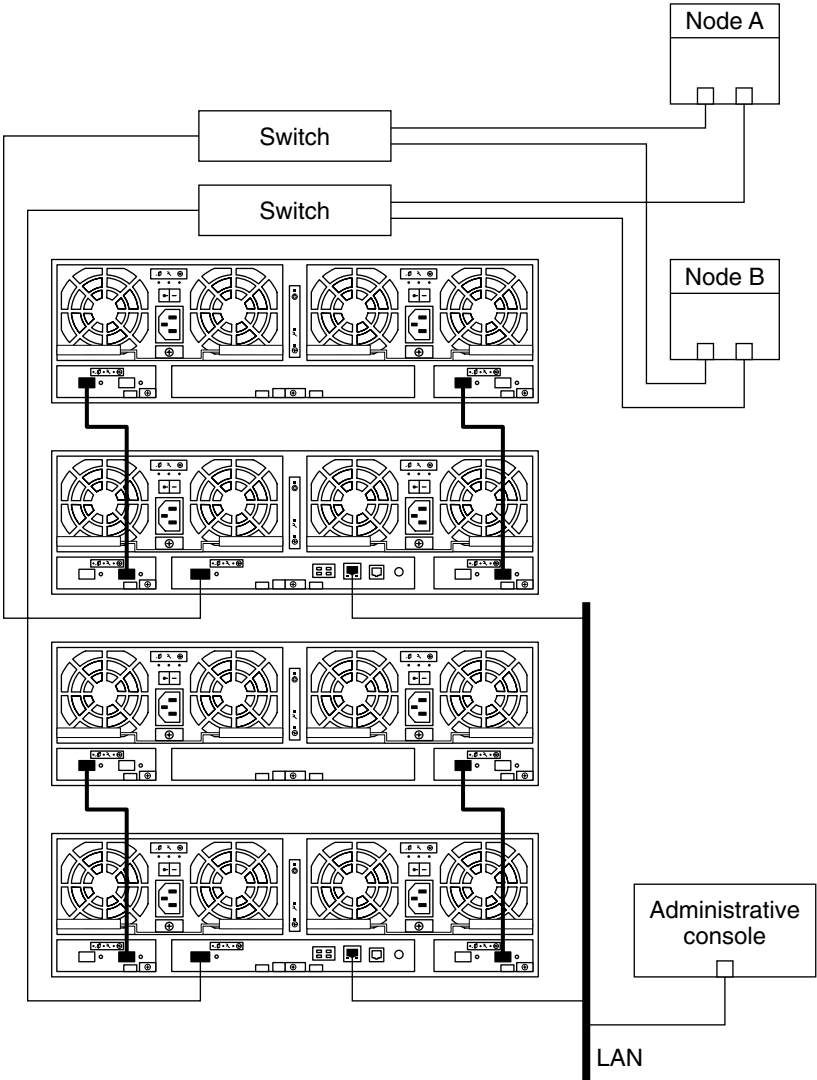


Figure 1-3 shows two storage arrays, and two have controllers. The two storage arrays are daisy-chained. The two storage arrays connect to a 2-node cluster through two switches.

FIGURE 1-3 Installing a 2x2 Configuration





Figure 1-4 shows four storage arrays, and two have controllers. All storage array are daisy-chained in the following order: alternate master, master, alternate master, and master. The two storage arrays with controllers connect to a two-node cluster through two switches.

FIGURE 1-4 Installing a 2x4 Configuration



## ▼ How to Install a Single-Controller Configuration in a New Cluster

Use this procedure to install a storage array in a single-controller configuration *before* you install the Solaris operating environment and Oracle Solaris Clustersoftware on your nodes. The following procedures contain instructions for other array-installation situations:

- “How to Install a Dual-Controller Configuration in a New Cluster” on page 18
- “How to Add a Single-Controller Configuration to an Existing Cluster” on page 21
- “How to Add a Dual-Controller Configuration to an Existing Cluster” on page 26

### 1 Install the host adapters in the nodes that are to be connected to the storage array.

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

### 2 Install the Fibre Channel (FC) switches.

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

### 3 Set up a Reverse Address Resolution Protocol (RARP) server on the network on which the new storage array is to reside.

Use the RARP server to set up the following network settings.

- IP address
- gateway (if necessary)
- netmask (if necessary)
- hostname

For the procedure about how to set up a RARP server, see the *Sun StorEdge 6120 Array Installation Guide*.

### 4 Cable the storage arrays.

For the procedures on how to connect your storage array, see the *Sun StorEdge 6120 Array Installation Guide*.

- a. Connect the storage arrays to the FC switches by using fiber-optic cables.
- b. Connect the Ethernet cables from each storage array to the Local Area Network (LAN).
- c. If necessary, install the interconnect cables between storage arrays.
- d. Connect the power cords to each storage array.



---

**5 Power on the storage array.**

Verify that all components are powered on and functional.

---

**Note** – The storage array might require a few minutes to boot.

---

For the procedure about how to power on a storage array, see the *Sun StorEdge 6120 Array Installation Guide*.

**6 Install any required controller firmware for the storage arrays.**

For the required controller firmware for the storage array, see the *Sun StorEdge 6120 Array Release Notes*.

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.

For required firmware, see the *Sun System Handbook*.

**7 Ensure that the cache and mirror settings for each storage array are set to auto.**

For more information, see the *Sun StorEdge 6020 and 6120 Array System Manual*.

**8 Ensure that the mp\_support parameter for each storage array is set to none.**

For more information, see the *Sun StorEdge 6020 and 6120 Array System Manual*.

**9 Ensure that all storage array controllers are ONLINE.**

For more information about how to bring controllers online, see the *Sun StorEdge 6020 and 6120 Array System Manual*.

**10 Reset the storage array to update the network settings and system settings that you changed.**

For the procedure about how to reboot or reset a storage array, see the *Sun StorEdge 6020 and 6120 Array System Manual*.

- 11 On all nodes, install the Oracle Solaris operating environment. Apply any required Oracle Solaris patches for Oracle Solaris Cluster software and storage array support.**

For the procedure about how to install the Oracle Solaris operating environment, see your Oracle Solaris Cluster software installation documentation.

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.

For required firmware, see the *Sun System Handbook*.

- 12 On each node, ensure that the `mpxio-disable` parameter is set to `yes` in the `/kernel/drv/scsi_vhci.conf` file.**

**See Also** To continue with Oracle Solaris Cluster software installation tasks, see your Oracle Solaris Cluster software installation documentation.

## ▼ How to Install a Dual-Controller Configuration in a New Cluster

Use this procedure to install a storage array in a dual-controller configuration *before* you install the Oracle Solaris operating environment and Oracle Solaris Cluster software on your nodes. The following procedures contain instructions for other array-installation situations:

- “How to Install a Single-Controller Configuration in a New Cluster” on page 16
- “How to Add a Single-Controller Configuration to an Existing Cluster” on page 21
- “How to Add a Dual-Controller Configuration to an Existing Cluster” on page 26

- 1 Install the host adapters in the nodes to be connected to the storage arrays.**

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

**2 Install the Fibre Channel (FC) switches.**

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

**3 If necessary, install GBICs or SFPs in the FC switches.**

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

**4 Set up a Reverse Address Resolution Protocol (RARP) server on the network on which the new storage array is to reside.**

Use the RARP server to set up the following network settings.

- IP address
- gateway (if necessary)
- netmask (if necessary)
- hostname

For the procedure about how to set up a RARP server, see the *Sun StorEdge 6120 Array Installation Guide*.

**5 Cable the storage arrays.**

For the procedures on how to connect your storage array, see the *Sun StorEdge 6120 Array Installation Guide*.

- a. **Connect the storage arrays to the FC switches by using fiber-optic cables.**
- b. **Connect the Ethernet cables from each storage array to the LAN.**
- c. **If necessary, install the interconnect cables between storage arrays.**
- d. **Connect the power cords to each storage array.**

For the procedure about how to install fiber-optic, Ethernet, and interconnect cables, see the *Sun StorEdge 6120 Array Installation Guide*.

**6 Power on the storage arrays.**

Verify that all components are powered on and functional.

For the procedure about how to power on the storage arrays, see the *Sun StorEdge 6120 Array Installation Guide*.

**7 Install any required controller firmware for the storage arrays.**

Access the *master* controller unit and administer the storage arrays. The master controller unit is the storage array that has the interconnect cables attached to the second port of each interconnect card, when viewed from the rear of the storage arrays.

For the required controller firmware for the storage array, see the *Sun StorEdge 6120 Array Release Notes*.

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.

For required firmware, see the *Sun System Handbook*.

**8 Ensure that each storage array has a unique target address.**

For the procedure about how to assign a target address to a storage array, see the *Sun StorEdge 6020 and 6120 Array System Manual*.

**9 Ensure that the cache and mirror settings for each storage array are set to auto.**

For more information, see the *Sun StorEdge 6020 and 6120 Array System Manual*.

**10 On each node, ensure that the mp\_support parameter for each storage array is set to mp\_xio.**

For more information, see the *Sun StorEdge 6020 and 6120 Array System Manual*.

**11 Ensure that all storage array controllers are ONLINE.**

For more information about how to bring controllers online, see the *Sun StorEdge 6020 and 6120 Array System Manual*.

**12 Reset the storage array to update the network settings and system settings that you changed.**

For the procedure about how to reboot or reset a storage array, see the *Sun StorEdge 6020 and 6120 Array System Manual*.

**13 On all nodes, install the Oracle Solaris operating system and apply the required Oracle 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*.

- 14 **Confirm that all storage arrays that you installed are visible to all nodes.**

```
# luxadm probe
```

**See Also** To continue with Oracle Solaris Cluster software installation tasks, see your Oracle Solaris Cluster software installation documentation.

## How to Add a Single-Controller Configuration to an Existing Cluster

Use this procedure to add a single-controller configuration to a running cluster. The following procedures contain instructions for other array-installation situations:

- “How to Install a Single-Controller Configuration in a New Cluster” on page 16
- “How to Install a Dual-Controller Configuration in a New Cluster” on page 18
- “How to Add a Dual-Controller Configuration to an Existing Cluster” on page 26

This procedure defines Node N as the node with which you begin working.

### ▼ How to Perform Initial Configuration Tasks on the Storage Array

- 1 **Power on the storage array.**

---

**Note** – The storage array might require a few minutes to boot.

---

For the procedure about how to power on a storage array, see the *Sun StorEdge 6120 Array Installation Guide*.

- 2 **Administer the storage array's network settings. Network settings include the following settings.**

- IP address
- gateway (if necessary)
- netmask (if necessary)
- hostname

For the procedure about how to set up an IP address, gateway, netmask, and hostname on a storage array, see the *Sun StorEdge 6020 and 6120 Array System Manual*.

- 3 **Ensure that the cache and mirror settings for each storage array are set to auto.**

For more information, see the *Sun StorEdge 6020 and 6120 Array System Manual*.

- 4 **Ensure that the mp\_support parameter for each storage array is set to none.**

For more information about, see the *Sun StorEdge 6020 and 6120 Array System Manual*.

## 5 Install any required controller firmware for the storage arrays you are adding.

For the required controller firmware for the storage array, see the *Sun StorEdge 6120 Array Release Notes*.

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.

For required firmware, see the *Sun System Handbook*.

## 6 Reset the storage array to update the network settings and system settings that you changed.

For the procedure about how to reboot or reset a storage array, see the *Sun StorEdge 6020 and 6120 Array System Manual*.

## 7 Confirm that all storage arrays that you upgraded are visible to all nodes.

```
# luxadm probe
```

## ▼ How to Connect the Storage Array to FC Switches

### 1 Install the GBICs or SFPs in the storage array that you plan to add.

For the procedure about how to install a GBICs or SFPs, see the *Sun StorEdge 6120 Array Installation Guide*.

### 2 If necessary, install GBICs or SFPs in the FC switches.

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

### 3 Install the Ethernet cable between the storage array and the Local Area Network (LAN).

### 4 If necessary, daisy-chain or interconnect the storage arrays.

For the procedure about how to install interconnect cables, see the *Sun StorEdge 6120 Array Installation Guide*.

## 5 Install a fiber-optic cable between the FC switch and the storage array.

For the procedure about how to install a fiber-optic cable, see the *Sun StorEdge 6120 Array Installation Guide*.

## ▼ How to Connect the Node to the FC Switches or the Storage Array

### 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` role-based access control (RBAC) authorization.

### 1 Determine the resource groups and device groups that are running on all nodes.

Record this information because you use this information in [Step 19](#) and [Step 20](#) of this procedure to return resource groups and device groups to these nodes.

Use the following command:

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

### 2 Move all resource groups and device groups off Node N.

```
# clnode evacuate nodename
```

### 3 If you need to install a host adapter in Node N, proceed to [Step 4](#).

If you do not need to install host adapters, skip to [Step 10](#).

### 4 If the host adapter that you are installing is the first FC host adapter on Node N, determine whether the required drivers for the host adapter are already installed on this node.

For the required packages, see the documentation that shipped with your host adapters. If the host adapter that you are installing is *not* the first FC host adapter on Node N, skip to [Step 6](#).

### 5 If the Fibre Channel support packages are not installed, install them.

The storage array packages are located in the Product directory of the Oracle Solaris CD-ROM. Add any necessary packages.

### 6 Shut down and power off Node N.

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

### 7 Install the host adapter in Node N.

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

**8 Power on and boot Node N into noncluster mode by adding -x to your boot instruction.**

For the procedure about how to boot a node in noncluster mode, see [Chapter 3, “Shutting Down and Booting a Cluster,” in \*Oracle Solaris Cluster System Administration Guide\*](#)Chapter 3, “Shutting Down and Booting a Cluster,” in *Oracle Solaris Cluster System Administration Guide*.

**9 If necessary, upgrade the host adapter firmware on Node N.**

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.

For required firmware, see the *Sun System Handbook*.

**10 If necessary, install a GBIC or an SFP in the FC switch or the storage array.**

For the procedure about how to install a GBIC or an SFP, see the documentation that shipped with your FC switch hardware. For the procedure on how to install a GBIC or an SFP, see the *Sun StorEdge 6120 Array Installation Guide*.

**11 Connect a fiber-optic cable between the FC switch and Node N.**

For the procedure about how to install a fiber-optic cable, see the *Sun StorEdge 6120 Array Installation Guide*.

**12 If necessary, install the required Oracle Solaris patches for storage array support on Node N.**

For a list of required Oracle Solaris patches for storage array support, see the *Sun StorEdge 6120 Array Release Notes*.

**13 On the node, update the /devices and /dev entries.**

```
# devfsadm -C
```

**14 Boot the node into cluster mode.**

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



- 15 On the node, update the paths to the DID instances.**  
`# cldevice populate`
- 16 If necessary, label the new logical volume.**  
 For the procedure about how to label a logical volume, see the *Sun StorEdge 6020 and 6120 Array System Manual*.
- 17 (Optional) On Node N, verify that the device IDs (DIDs) are assigned to the new LUNs.**  
`# cldevice clear`  
`# cldevice list -v`
- 18 Repeat Step 2 through Step 17 for each remaining node that you plan to connect to the storage array.**
- 19 (Optional) Restore the device groups to the original node.**  
 Perform the following step for each device group you want to return to the original node.  
`# cldevicegroup switch -n nodename devicegroup1[ devicegroup2 ...]`  
`-n nodename` The node to which you are restoring device groups.  
`devicegroup1[ devicegroup2 ...]` The device group or groups that you are restoring to the node.
- 20 (Optional) Restore the resource groups to the original node.**  
 Perform the following step for each resource group you want to return to the original node.  
`# clresourcegroup switch -n nodename resourcegroup1[ resourcegroup2 ...]`  
`nodename` For failover resource groups, the node to which the groups are returned. For scalable resource groups, the node list to which the groups are returned.  
`resourcegroup1[ resourcegroup2 ...]` The resource group or groups that you are returning to the node or nodes.
- 21 Perform volume management administration to incorporate the new logical volumes into the cluster.**  
 For more information, see your Solaris Volume Manager or Veritas Volume Manager documentation.

## How to Add a Dual-Controller Configuration to an Existing Cluster

Use this procedure to add a dual-controller configuration to a running cluster. The following procedures contain instructions for other array-installation situations:

- “How to Install a Single-Controller Configuration in a New Cluster” on page 16
- “How to Install a Dual-Controller Configuration in a New Cluster” on page 18
- “How to Add a Single-Controller Configuration to an Existing Cluster” on page 21

This procedure defines Node N as the node with which you begin working.

### ▼ How to Perform Initial Configuration Tasks on the Storage Array

#### 1 Power on the storage arrays.

---

**Note** – The storage arrays might require several minutes to boot.

---

For the procedure about how to power on storage arrays, see the *Sun StorEdge 6120 Array Installation Guide*.

#### 2 Administer the storage array's network settings. Network settings include the following settings.

- IP address
- gateway (if necessary)
- netmask (if necessary)
- hostname

Assign an IP address to the *master* controller unit only. The master controller unit is the storage array that has the interconnect cables attached to the second port of each interconnect card.

For the procedure about how to set up an IP address, gateway, netmask, and hostname on a storage array, see the *Sun StorEdge 6020 and 6120 Array System Manual*.

#### 3 Ensure that each storage array has a unique target address.

For the procedure about how to assign a target address to a storage array, see the *Sun StorEdge 6020 and 6120 Array System Manual*.

#### 4 Ensure that the cache and mirror settings for each storage array are set to auto.

For more information, see the *Sun StorEdge 6020 and 6120 Array System Manual*.

#### 5 Ensure that the mp\_support parameter for each storage array is set to mpzio.

For more information, see the *Sun StorEdge 6020 and 6120 Array System Manual*.

**6 Install any required controller firmware for the storage arrays you are adding.**

Access the *master* controller unit and administer the storage arrays. The master controller unit is the storage array that has the interconnect cables attached to the second port of each interconnect card, when viewed from the rear of the storage arrays.

For the required controller firmware for the storage array, see the *Sun StorEdge 6120 Array Release Notes*.

**7 Reset the storage array to update the network settings and system settings that you changed.**

For the procedure about how to reboot or reset a storage array, see the *Sun StorEdge 6020 and 6120 Array System Manual*.

**▼ How to Connect the Storage Array to FC Switches****1 Install the GBICs or SFPs in the storage array that you plan to add.**

For the procedure about how to install a GBICs or SFPs, see the *Sun StorEdge 6120 Array Installation Guide*.

**2 If necessary, install GBICs or SFPs in the FC switches.**

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

**3 Install the Ethernet cable between the storage arrays and the local area network (LAN).****4 If necessary, daisy-chain or interconnect the storage arrays.**

For the procedure about how to install interconnect cables, see the *Sun StorEdge 6120 Array Installation Guide*.

**5 Install a fiber-optic cable between each FC switch and both new storage arrays of the partner group.**

For the procedure about how to install a fiber-optic cable, see the *Sun StorEdge 6120 Array Installation Guide*.

**▼ How to Connect the Node to the FC Switches or the Storage Array****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` role-based access control (RBAC) authorization.

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

Record this information because you use this information in [Step 18](#) and [Step 19](#) of this procedure to return resource groups and device groups to these nodes.

Use the following command:

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

**2 Move all resource groups and device groups off Node N.**

```
# clnode evacuate nodename
```

**3 If you need to install host adapters in Node N, proceed to [Step 4](#).**

If you do not need to install host adapters, skip to [Step 10](#).

**4 If the host adapter that you that are installing is the first host adapter on Node N, determine whether the required drivers for the host adapter are already installed on this node.**

For the required packages, see the documentation that shipped with your host adapters.

If the host adapter that you are installing is *not* the first host adapter on Node N, skip to [Step 6](#).

**5 If the required support packages are not already installed, install them.**

The support packages are located in the Product directory of the Solaris CD-ROM.

**6 Shut down and power off Node N.**

For the 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*.

**7 Install the host adapters in Node N.**

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

**8 Power on and boot Node N into noncluster mode by adding `-x` to your boot instruction.**

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

**9 If necessary, upgrade the host adapter firmware on Node N.**

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.

For required firmware, see the *Sun System Handbook*.

**10 If necessary, install a GBIC or an SFP in the FC switch or the storage array.**

For the procedure about how to install a GBIC or an SFP, see the documentation that shipped with your FC switch hardware. For the procedure on how to install a GBIC or an SFP, see the *Sun StorEdge 6120 Array Installation Guide*.

**11 Connect a fiber-optic cable between the FC switch and Node N.**

For the procedure about how to install a fiber-optic cable, see the *Sun StorEdge 6120 Array Installation Guide*.

**12 Install the required Solaris patches for storage array support on Node N.**

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.

For required firmware, see the *Sun System Handbook*.

- 13 Perform a reconfiguration boot on Node N by adding -r to your boot instruction, to create the new Solaris device files and links.**

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

- 14 On Node N, update the paths to the DID instances.**

```
# cldevice populate
```

- 15 If necessary, label the new storage array logical volume.**

For the procedure about how to label a logical volume, see the *Sun StorEdge 6020 and 6120 Array System Manual*.

- 16 (Optional) On Node N, verify that the device IDs (DIDs) are assigned to the new LUNs.**

```
# cldevice clear
# cldevice list -v
```

- 17 Repeat Step 2 through Step 16 for each remaining node that you plan to connect to the storage array.**

- 18 (Optional) Restore the device groups to the original node.**

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

```
# cldevicegroup switch -n nodename devicegroup1[ devicegroup2 ...]
```

*-n nodename*

The node to which you are restoring device groups.

*devicegroup1[ devicegroup2 ...]*

The device group or groups that you are restoring to the node.

- 19 (Optional) Restore the resource groups to the original node.**

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

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

*nodename*

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

*resourcegroup1[ resourcegroup2 ...]*

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

- 20 Perform volume management administration to incorporate the new logical volumes into the cluster.**

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

# Configuring Storage Arrays

This section contains the procedures about how to configure a storage array in a running cluster. [Table 1–2](#) lists these procedures.

TABLE 1–2 Task Map: Configuring a Storage Array

Task	Information
Create a LUN	<a href="#">“How to Create a Logical Volume” on page 31</a>
Remove a LUN	<a href="#">“How to Remove a Logical Volume” on page 33</a>

The following is a list of administrative tasks that require no cluster-specific procedures. See the *Sun StorEdge 6020 and 6120 Array System Manual* for the following procedures.

- Removing a logical volume.
- Creating a logical volume.

## ▼ 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
```

**13 For each additional node that is connected to the shared storage that hosted the logical volume, repeat [Step 9](#) to [Step 12](#).****14 (Optional) Restore the device groups to the original node.**

Do the following for each device group that you want to return to the original node.

```
# cldevicegroup switch -n nodename devicegroup1 [ devicegroup2 ...]
```

*-n nodename*

The node to which you are restoring device groups.

*devicegroup1* [ *devicegroup2* ...]

The device group or groups that you are restoring to the node.

**15 (Optional) Restore the resource groups to the original node.**

Do the following for each resource group that you want to return to the original node.

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

*nodename*

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

*resourcegroup1* [ *resourcegroup2* ...]

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

## Maintaining Storage Arrays

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

**Note** – When you upgrade firmware on a storage device or on an enclosure, redefine the stripe size of a LUN, or perform other LUN operations, a device ID might change unexpectedly. When you perform a check of the device ID configuration by running the `cldevice check` command, the following error message appears on your console if the device ID changed unexpectedly.

```
device id for nodename:/dev/rdisk/cXtYdZsN does not match physical
device's id for ddecimalnumber, device may have been replaced.
```

To fix device IDs that report this error, run the `cldevice repair` command for each affected device.

**TABLE 1–3** Task Map: Maintaining a Storage Array

Task	Information
Upgrade storage array firmware.	<a href="#">“How to Upgrade Storage Array Firmware” on page 37</a>
Remove a storage array or partner group.	<a href="#">“How to Remove a Single-Controller Configuration” on page 38</a> <a href="#">“How to Remove a Dual-Controller Configuration” on page 41</a>
Replace a node-to-switch component. <ul style="list-style-type: none"> <li>▪ Node-to-switch fiber-optic cable</li> <li>▪ FC host adapter</li> <li>▪ FC switch</li> <li>▪ GBIC or SFP</li> </ul>	<a href="#">“Replacing a Node-to-Switch Component” on page 44</a>
Replace a node's host adapter.	<a href="#">“How to Replace a Host Adapter” on page 46</a>
Add a node to the storage array.	Oracle Solaris Cluster system administration documentation
Remove a node from the storage device.	Oracle Solaris Cluster system administration documentation

## StorEdge 6120 Array FRUs

The following is a list of administrative tasks that require no cluster-specific procedures. See the *Sun StorEdge 6020 and 6120 Array System Manual* for the following procedures.

- Adding an expansion unit.
- Removing an expansion unit.
- Replacing an expansion unit.
- Adding a disk drive.
- Replacing a failed disk drive.
- Removing a disk drive.
- Replacing an Ethernet cable.
- Replacing a power cable on the storage array.
- Replacing a power and cooling unit (PCU).
- Replacing a chassis.
- Replacing a controller.
- Replacing a storage array-to-switch component in a dual-controller configuration.

A storage array-to-switch component includes all of the following components.

- Fiber-optic cable that connects an FC switch to a storage array
- FC switch GBIC or an SFP that connects an FC switch to a storage array
- FC switch that connects an FC switch to a storage array
- FC switch power cord
- Interconnect cables that connect two storage arrays
- Interconnect cards

## ▼ 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 the list of required patches, see the *Sun StorEdge 6120 Array Release Notes*. For the procedure about how to apply firmware patches, see the firmware patch README file. For the procedure about how to verify the firmware level, see the *Sun StorEdge 6020 and 6120 Array System Manual*.  
  
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 Single-Controller Configuration

Use this procedure to permanently remove a storage array from a running cluster. This storage array resides in a single-controller configuration. This procedure provides the flexibility to remove the host adapters from the nodes for the storage array that you are removing.

This procedure defines Node A as the node with which you begin working. Node B is another node in the cluster.



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

---

### 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` role-based access control (RBAC) authorization.

- 1 Back up all database tables, data services, and volumes that are associated with the storage array that you are removing.**
- 2 Detach the submirrors from the storage array that you are removing. Detach the submirrors to stop all I/O activity to the storage array.**

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

- 3 Remove the references to the LUN(s) from any diskset or disk group.**

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

- 4 Determine the resource groups and device groups that are running on Node A.**

You will need this information to restore resource groups and device groups to the original node in [Step 18](#) and [Step 19](#) of this procedure.

Use the following command:

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

- 5 Shut down Node A.**

For the 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 If the storage array that you are removing is the last storage array that is connected to Node A, disconnect the fiber-optic cable between Node A and the FC switch that is connected to this storage array. Then disconnect the fiber-optic cable between the FC switch and this storage array.**

For the procedure about how to remove a fiber-optic cable, see the *Sun StorEdge 6020 and 6120 Array System Manual*.

- 7 If you want to remove the host adapter from Node A, power off Node A . . .**

If you do not want to remove the host adapter, skip to [Step 10](#).

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

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

**9 Without enabling the node to boot, power on Node A.**

For more information, see your Oracle Solaris Cluster system administration documentation.

**10 Boot Node A into cluster mode.**

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

**11 Shut down Node B.**

For the 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*.

**12 If the storage array that you are removing is the last storage array that is connected to the FC switch, disconnect the fiber-optic cable that connects this FC switch and Node B.**

For the procedure about how to remove a fiber-optic cable, see the *Sun StorEdge 6020 and 6120 Array System Manual*.

**13 If you want to remove the host adapter from Node B, power off Node B.**

If you do not want to remove the host adapter, skip to [Step 16](#).

**14 Remove the host adapter from Node B.**

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

**15 Without enabling the node to boot, power on Node B.**

For more information, see your Oracle Solaris Cluster system administration documentation.

**16 Boot Node B into cluster mode.**

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

**17 On all nodes, update the /devices and /dev entries.**

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

**18 (Optional) Restore the device groups to the original node.**

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

```
# cldevicegroup switch -n nodename devicegroup1[ devicegroup2 ...]
```

-n nodename                                  The node to which you are restoring device groups.

devicegroup1[ devicegroup2 ...]          The device group or groups that you are restoring to the node.



**19 (Optional) Restore the resource groups to the original node.**

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

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

*nodename*

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

*resourcegroup1[ resourcegroup2 ...]*

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

## ▼ How to Remove a Dual-Controller Configuration

Use this procedure to remove a dual-controller configuration from a running cluster. This procedure defines Node A as the node with which you begin working. Node B is another node in the cluster.




---

**Caution** – During this procedure, you lose access to the data that resides on each partner group that you are removing.

---

### 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` 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 you are removing.**
- 2 If necessary, detach the submirrors from each storage array or partner group that you are removing. Detach the submirrors to stop all I/O activity to the storage array or partner group.**  
For more information, see your Solaris Volume Manager or Veritas Volume Manager documentation.
- 3 Remove references to each LUN. This LUN belongs to the storage array or partner group that you are removing.**  
For more information, see your Solaris Volume Manager or Veritas Volume Manager documentation.

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

You will need this information to restore resource groups and device groups to the original node in [Step 19](#) and [Step 20](#) of this procedure.

Use the following command:

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

**5 Shut down Node A.**

For the 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 Disconnect from both storage arrays the fiber-optic cables that connect to the FC switches. Then disconnect the Ethernet cables.****7 If any storage array that you are removing is the last storage array that is connected to an FC switch that is on Node A, disconnect the fiber-optic cable between Node A and the FC switch that was connected to this storage array.**

If no storage array that you are removing is the last array connected to any switch, skip to [Step 11](#).

For the procedure about how to remove a fiber-optic cable, see the *Sun StorEdge 6020 and 6120 Array System Manual*.

**8 If you want to remove the host adapters from Node A, power off the node.**

If you do *not* want to remove host adapters from the node, skip to [Step 11](#).

**9 Remove the host adapters from Node A.**

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

**10 Without enabling the node to boot, power on Node A.**

For more information, see your Oracle Solaris Cluster system administration documentation.

**11 Boot Node A into cluster mode.**

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

**12 Shut down Node B.**

For the 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*.

- 13 If any storage array that you are removing is the last storage array that is connected to an FC switch that is on Node B, disconnect the fiber-optic cable that connects this FC switch to Node B.**  
For the procedure about how to remove a fiber-optic cable, see the *Sun StorEdge 6020 and 6120 Array System Manual*.
- 14 If you want to remove the host adapters from Node B, power off the node.**  
If you do not want to remove host adapters, skip to [Step 17](#).
- 15 Remove the host adapters from Node B.**  
For the procedure about how to remove host adapters, see the documentation that shipped with your nodes.
- 16 Without enabling the node to boot, power on Node B.**  
For more information, see your Oracle Solaris Cluster system administration documentation.
- 17 Boot Node B into cluster mode.**  
For more information about how to boot nodes, see [Chapter 3, “Shutting Down and Booting a Cluster,”](#) in *Oracle Solaris Cluster System Administration Guide*.
- 18 On all nodes, update the /devices and /dev entries.**
- ```
# devfsadm -C
# cldevice clear
```
- 19 (Optional) Restore the device groups to the original node.**  
Perform the following step for each device group you want to return to the original node.

```
# cldevicegroup switch -n nodename devicegroup1[ devicegroup2 ...]
```

`-n nodename`                              The node to which you are restoring device groups.

`devicegroup1[ devicegroup2 ...]`      The device group or groups that you are restoring to the node.

**20 (Optional) Restore the resource groups to the original node.**  
Perform the following step for each resource group you want to return to the original node.

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

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

`resourcegroup1[ resourcegroup2 ...]`      The resource group or groups that you are returning to the node or nodes.

## 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 46](#).

---

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 44](#).
- If your cluster does *not* use multipathing, see [“How to Replace a Node-to-Switch Component in a Cluster Without Multipathing” on page 45](#).

### ▼ **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.

- 6 **(Optional) Restore the device groups to the original node.**

Do the following for each device group that you want to return to the original node.

```
# cldevicegroup switch -n nodename devicegroup1 [ devicegroup2 ... ]
```

-n *nodename*     The node to which you are restoring device groups.

*devicegroup1* [ *devicegroup2* ... ]     The device group or groups that you are restoring to the node.

- 7 **(Optional) Restore the resource groups to the original node.**

Do the following for each resource group that you want to return to the original node.

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

|                                                    |                                                                                                                                                        |
|----------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------|
| <i>nodename</i>                                    | For failover resource groups, the node to which the groups are returned. For scalable resource groups, the node list to which the groups are returned. |
| <i>resourcegroup1</i> [ <i>resourcegroup2</i> ...] | The resource group or groups that you are returning to the node or nodes.                                                                              |

## ▼ 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 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.

For required firmware, see the *Sun System Handbook*.

**9 Boot Node A into cluster mode.**

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

**10 (Optional) Restore the device groups to the original node.**

Do the following for each device group that you want to return to the original node.

```
# cldevicegroup switch -n nodename devicegroup1[ devicegroup2 ...]
```

|                                        |                                                                |
|----------------------------------------|----------------------------------------------------------------|
| <i>-n nodename</i>                     | The node to which you are restoring device groups.             |
| <i>devicegroup1[ devicegroup2 ...]</i> | The device group or groups that you are restoring to the node. |

**11 (Optional) Restore the resource groups to the original node.**

Do the following for each resource group that you want to return to the original node.

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

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

|                                            |                                                                           |
|--------------------------------------------|---------------------------------------------------------------------------|
| <i>resourcegroup1[ resourcegroup2 ...]</i> | The resource group or groups that you are returning to the node or nodes. |
|--------------------------------------------|---------------------------------------------------------------------------|



# Index

---

## Numbers and Symbols

- 1x1 configuration, 12
- 1x2 configuration, 13
- 2x2 configuration, 14
- 2x4 configuration, 14
- 6120 storage array
  - cabling, 12–14
  - FRUs, 37
  - product documentation, 6

## A

- adding
  - See also* installing
  - disk drives, 37
  - dual-controller configurations, 26–30
  - expansion units, 37
  - nodes, 36
  - partner groups, 26–30
  - single-controller configurations, 21–25
  - storage arrays, 21–25
- arrays, *See* storage arrays

## C

- cables
  - replacing, 37
- cabling configurations
  - 1x1 configuration, 12
  - 1x2 configuration, 13

- cabling configurations (*Continued*)
  - 2x2 configuration, 14
  - 2x4 configuration, 14
- chassis, replacing, 37
- configuring, storage arrays, 31–35
- controllers, replacing, 37
- cooling units, replacing, 37
- creating, logical volumes, 31–33

## D

- deleting logical volumes, 33–35
- disk drives
  - adding, 37
  - removing, 37
  - replacing, 37
- dual-controller configurations
  - adding, 26–30
  - installing, 18–21
  - removing, 41–43

## E

- Ethernet cables, replacing, 37
- expansion units
  - adding, 37
  - removing, 37
  - replacing, 37

**F**

firmware, upgrade storage array firmware, 37–38  
FRUs, Sun StorEdge 6120 storage array, 37

**H**

HBAs, replacing, 46–48  
help, 7–8  
host adapters, replacing, 46–48

**I**

installing  
    *See also* adding  
    storage arrays, 11–30, 18–21  
interconnect cables, replacing, 37  
interconnect cards, replacing, 37

**L**

logical devices, *See* logical volumes  
logical unit numbers, *See* logical volumes  
logical volumes  
    creating, 31–33  
    removing, 33–35  
loop cables, *See* interconnect cables  
loop cards, *See* interconnect cards  
LUNs, *See* logical volumes

**M**

maintaining, storage arrays, 36–48  
midplane, replacing, 37

**N**

node-to-switch components, replacing, 44–46  
nodes  
    adding, 36  
    removing, 36

**P**

partner groups  
    adding, 26–30  
    removing, 41–43  
power cables, replacing, 37  
power units, replacing, 37

**R**

RAID-1 configurations  
    1x1 configuration, 12  
    1x2 configuration, 13  
removing  
    disk drives, 37  
    dual-controller configurations, 41–43  
    expansion units, 37  
    logical volumes, 33–35  
    nodes, 36  
    partner groups, 41–43  
    single-controller configurations, 38–41  
    storage arrays, 38–41  
replacing  
    chassis, 37  
    controllers, 37  
    cooling units, 37  
    disk drives, 37  
    Ethernet cables, 37  
    expansion units, 37  
    host adapters, 46–48  
    midplane, 37  
    node-to-switch components, 44–46  
    power cables, 37  
    power units, 37  
    switch-to-storage array components, 37

**S**

single-controller configurations  
    adding, 21–25  
    installing, 16–18  
    removing, 38–41  
software RAID-1 configurations  
    1x1 configuration, 12

software RAID-1 configurations (*Continued*)

1x2 configuration, 13

storage array firmware, upgrading, 37–38

storage arrays, 11–48

adding single-controller configurations, 21–25

configuring, 31–35

installing, 11–30

installing dual-controller configurations, 18–21

installing single-controller configurations, 16–18

maintaining, 36–48

removing, 38–41

Sun StorEdge 6120 storage array

cabling, 12–14

FRUs, 37

product documentation, 6

switch-to-storage array components, replacing, 37

## **T**

technical support, 7–8

## **U**

upgrading, storage array firmware, 37–38

