Sun Cluster 3.1 - 3.2 With Sun StorEdge 6120 Array Manual for Solaris OS

Sun microsystems

Contents

Preface

The *Sun Cluster 3.1 - 3.2 With Sun StorEdge 6120 Array Manual for Solaris OS* provides procedures specific to Sun™ StorEdge™ 6120 arrays in dual-controller configurations and in single–controller configurations that are placed in a Sun Cluster environment.

Use this manual with any version of Sun Cluster 3.1 or 3.2 software on SPARC® based clusters and with any supported version of Sun Cluster 3.1 or 3.2 on x86 based clusters. Unless otherwise noted, procedures are the same for all Sun Cluster 3.1 or 3.2 versions.

Note – This Sun 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 for a list of changes to this manual.

Who Should Use This Book

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

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

How This Book Is Organized

This book contains one chapter 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 Sun Cluster 3.1 - 3.2 With Sun StorEdge 6120 Array Manual for Solaris OS

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.

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 Sun Cluster books support the Sun Cluster 3.1 and 3.2 releases. If you are maintaining a different version of Sun Cluster software, refer to the appropriate documentation. All Sun 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
Sun StorEdge Traffic Manager Software Release Notes	817-0385
Available on http://www.sun.com/storage/san/	
Sun StorEdge Traffic Manager Installation and Configuration Guide	816-1420
Available on http://www.sun.com/products-n-solutions/hardware/docs	

TABLE P–2 Hardware Documentation *(Continued)*

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

TABLE P–3 Sun Cluster Documentation

Documentation
Solaris Cluster 3.2
Sun Cluster 3.1

Using UNIX Commands

This document contains information about commands that are used to install, configure, or upgrade a Sun 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 Solaris™ Operating System (Solaris OS)
- Other software documentation that you received with your system
- Solaris Operating System man pages

Getting Help

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

- Your name and email address (if available)
- Your company name, address, and phone number
- The model number and serial number of your systems
- The release number of the operating environment (for example, Solaris 10)
- The release number of Sun Cluster (for example, Sun Cluster 3.2)

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

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

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

Documentation, Support, and Training

The Sun web site provides information about the following additional resources:

- Documentation (http://www.sun.com/documentation/)
- Support (http://www.sun.com/support/)
- Training (http://www.sun.com/training/)

Sun Welcomes Your Comments

Sun is interested in improving its documentation and welcomes your comments and suggestions. To share your comments, go to http://docs.sun.com and click Feedback.

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 .login file. Use ls -a to list all files. machine_name% you have mail.

TABLE P–4 Typographic Conventions   *(Continued)*

Typeface	Meaning	Example
AaBbCc123	What you type, contrasted with onscreen computer output	machine_name% **su** Password:
aabbcc123	Placeholder: replace with a real name or value	The command to remove a file is rm *filename*.
AaBbCc123	Book titles, new terms, and terms to be emphasized	Read Chapter 6 in the *User's Guide*. A *cache* is a copy that is stored locally. Do *not* 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 the C shell, Bourne shell, and Korn shell.

TABLE P–5 Shell Prompts

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

1

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 a Sun™ Cluster environment.

This chapter contains the following main topics:

- "Installing Storage Arrays" on page 11
- "Configuring Storage Arrays" on page 34
- "Maintaining Storage Arrays" on page 40

Installing Storage Arrays

This section contains the procedures for installing single-controller and dual-controller storage array configurations in new and existing Sun 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	"How to Install a Single-Controller Configuration in a New Cluster" on page 17
Installing arrays in a new cluster, using a dual-controller configuration	"How to Install a Dual-Controller Configuration in a New Cluster" on page 20
Adding arrays to an existing cluster, using a single-controller configuration.	"How to Add a Single-Controller Configuration to an Existing Cluster" on page 22

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

Task	Information
Adding arrays to an existing cluster, using a dual-controller configuration.	"How to Add a Dual-Controller Configuration to an Existing Cluster" on page 28

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).
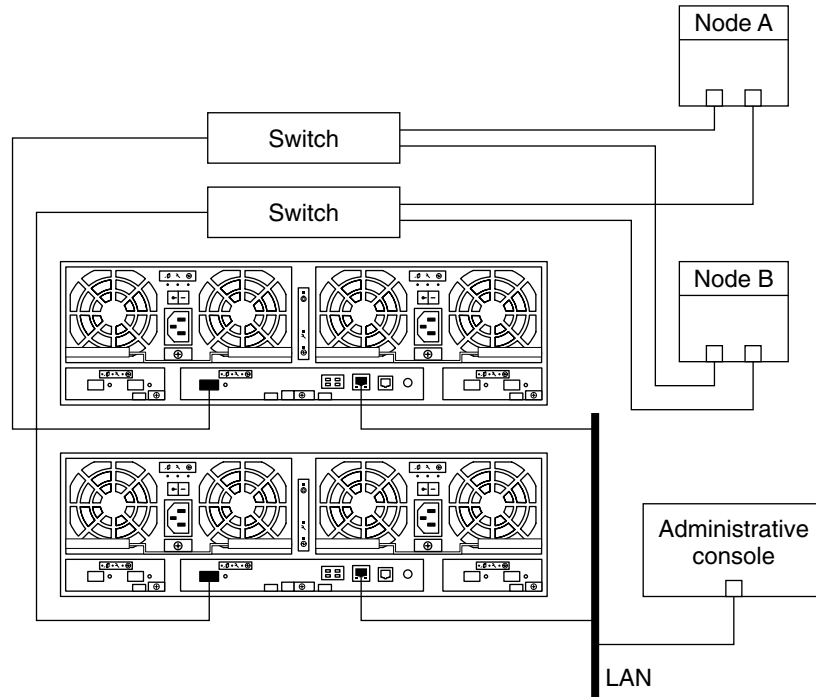
FIGURE 1–1   Installing a 1x1 Configuration With Software RAID-1

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.
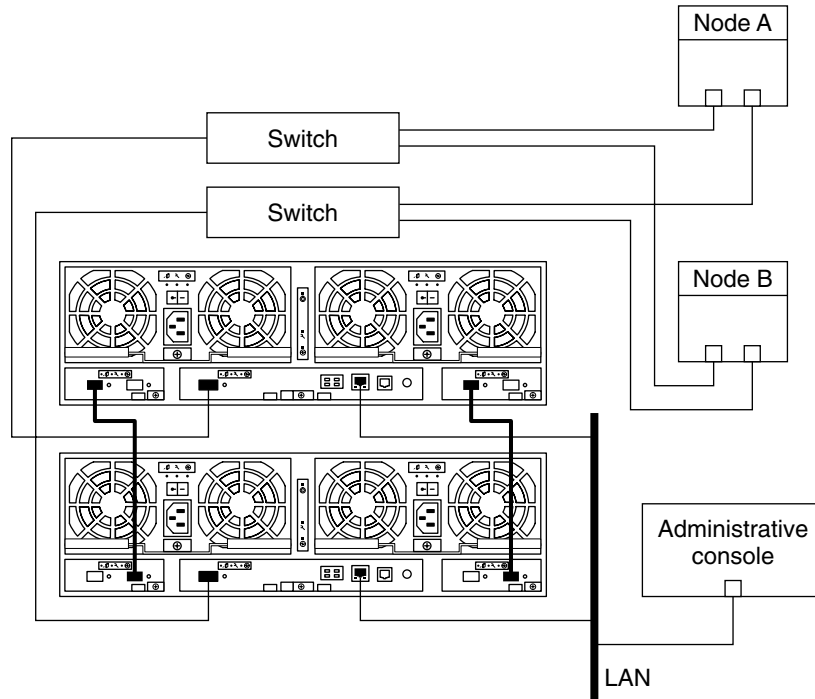
FIGURE 1–3   Installing a 2x2 Configuration

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 Sun Cluster software 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 20
- "How to Add a Single-Controller Configuration to an Existing Cluster" on page 22
- "How to Add a Dual-Controller Configuration to an Existing Cluster" on page 28

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*.

If you use the Solaris 8, Solaris 9, or Solaris 10 Operating System, Sun Connection Update Manager keeps you informed of the latest versions of patches and features. Using notifications and intelligent needs-based updating, Sun Connection helps improve operational efficiency and ensures that you have the latest software patches for your Sun software.

You can download the Sun Connection Update Manager product for free by going to `http://www.sun.com/download/products.xml?id=4457d96d`.

Additional information for using the Sun patch management tools is provided in *Solaris Administration Guide: Basic Administration* at http://docs.sun.com (`http://docs.sun.com`). Refer to the version of this manual for the 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 *Sun Cluster System Administration Guide for Solaris OS* 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 a list of patches that affect Sun Cluster, see the Sun Cluster Wiki Patch Klatch.

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 Solaris operating environment. Apply any required Solaris patches for Sun Cluster software and storage array support.

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

If you use the Solaris 8, Solaris 9, or Solaris 10 Operating System, Sun Connection Update Manager keeps you informed of the latest versions of patches and features. Using notifications and intelligent needs-based updating, Sun Connection helps improve operational efficiency and ensures that you have the latest software patches for your Sun software.

You can download the Sun Connection Update Manager product for free by going to http://www.sun.com/download/products.xml?id=4457d96d.

Additional information for using the Sun patch management tools is provided in *Solaris Administration Guide: Basic Administration* at http://docs.sun.com (http://docs.sun.com). Refer to the version of this manual for the 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 *Sun Cluster System Administration Guide for Solaris OS* 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 a list of patches that affect Sun Cluster, see the Sun Cluster Wiki Patch Klatch.

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 Sun Cluster software installation tasks, see your Sun 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 Solaris operating environment and Sun 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 17
- "How to Add a Single-Controller Configuration to an Existing Cluster" on page 22
- "How to Add a Dual-Controller Configuration to an Existing Cluster" on page 28

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*.

If you use the Solaris 8, Solaris 9, or Solaris 10 Operating System, Sun Connection Update Manager keeps you informed of the latest versions of patches and features. Using notifications and intelligent needs-based updating, Sun Connection helps improve operational efficiency and ensures that you have the latest software patches for your Sun software.

You can download the Sun Connection Update Manager product for free by going to `http://www.sun.com/download/products.xml?id=4457d96d`.

Additional information for using the Sun patch management tools is provided in *Solaris Administration Guide: Basic Administration* at `http://docs.sun.com` (`http://docs.sun.com`). Refer to the version of this manual for the 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 *Sun Cluster System Administration Guide for Solaris OS* 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 a list of patches that affect Sun Cluster, see the Sun Cluster Wiki Patch Klatch.

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** `mpxio`**.**

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 Solaris operating system and apply the required Solaris patches for Sun Cluster software and storage array support.

For the procedure about how to install the Solaris operating environment, see "How to Install Solaris Software" in *Sun Cluster Software Installation Guide for Solaris OS*.

14   Install any required patches or software for Solaris I/O multipathing software support to nodes and enable multipathing.

For the procedure about how to install the Solaris I/O multipathing software, see "How to Install Sun Multipathing Software" in *Sun Cluster Software Installation Guide for Solaris OS*.

15   Confirm that all storage arrays that you installed are visible to all nodes.

```
# luxadm probe
```

See Also   To continue with Sun Cluster software installation tasks, see your Sun 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 17
- "How to Install a Dual-Controller Configuration in a New Cluster" on page 20
- "How to Add a Dual-Controller Configuration to an Existing Cluster" on page 28

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*.

If you use the Solaris 8, Solaris 9, or Solaris 10 Operating System, Sun Connection Update Manager keeps you informed of the latest versions of patches and features. Using notifications and intelligent needs-based updating, Sun Connection helps improve operational efficiency and ensures that you have the latest software patches for your Sun software.

You can download the Sun Connection Update Manager product for free by going to http://www.sun.com/download/products.xml?id=4457d96d.

Additional information for using the Sun patch management tools is provided in *Solaris Administration Guide: Basic Administration* at http://docs.sun.com (http://docs.sun.com). Refer to the version of this manual for the 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 *Sun Cluster System Administration Guide for Solaris OS* 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 a list of patches that affect Sun Cluster, see the Sun Cluster Wiki Patch Klatch.

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 Sun Cluster commands. Most commands also have short forms. Except for the forms of the command names, the commands are identical. For a list of the commands and their short forms, see Appendix A, "Sun Cluster Object-Oriented Commands," in *Sun Cluster 3.1 - 3.2 Hardware Administration Manual for Solaris OS*.

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.

    - **If you are using Sun Cluster 3.2, use the following commands:**
      ```
      # clresourcegroup status +
      # cldevicegroup status +
      ```

    - **If you are using Sun Cluster 3.1, use the following command:**
      ```
      # scstat
      ```

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

    - **If you are using Sun Cluster 3.2, use the following command:**
      ```
      # clnode evacuate nodename
      ```

    - **If you are using Sun Cluster 3.1, use the following command:**
      ```
      # scswitch -S -h 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 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 Sun 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 *Sun Cluster System Administration Guide for Solaris OS*.

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

If you use the Solaris 8, Solaris 9, or Solaris 10 Operating System, Sun Connection Update Manager keeps you informed of the latest versions of patches and features. Using notifications and intelligent needs-based updating, Sun Connection helps improve operational efficiency and ensures that you have the latest software patches for your Sun software.

You can download the Sun Connection Update Manager product for free by going to `http://www.sun.com/download/products.xml?id=4457d96d`.

Additional information for using the Sun patch management tools is provided in *Solaris Administration Guide: Basic Administration* at http://docs.sun.com (`http://docs.sun.com`). Refer to the version of this manual for the 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 *Sun Cluster System Administration Guide for Solaris OS* 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 a list of patches that affect Sun Cluster, see the Sun Cluster Wiki Patch Klatch.

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 Solaris patches for storage array support on** Node N**.**

For a list of required 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 *Sun Cluster System Administration Guide for Solaris OS*.

15 **On the node, update the paths to the DID instances.**

- **If you are using Sun Cluster 3.2, use the following command:**
  ```
  # cldevice populate
  ```

- **If you are using Sun Cluster 3.1, use the following command:**
  ```
  # scgdevs
  ```

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.**

- **If you are using Sun Cluster 3.2, use the following commands:**
  ```
  # cldevice clear
  # cldevice list -v
  ```

- **If you are using Sun Cluster 3.1, use the following commands:**
  ```
  # scdidadm -C
  # scdidadm -l
  ```

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.

- **If you are using Sun Cluster 3.2, use the following command:**
  ```
  # 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.

- **If you are using Sun Cluster 3.1, use the following command:**
  ```
  # scswitch -z -D devicegroup -h nodename
  ```

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.

- **If you are using Sun Cluster 3.2, use the following command:**

  ```
  # 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.

- **If you are using Sun Cluster 3.1, use the following command:**

  ```
  # scswitch -z -g resourcegroup -h nodename
  ```

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:

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** mpxio.

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 Sun Cluster commands. Most commands also have short forms. Except for the forms of the command names, the commands are identical. For a list of the commands and their short forms, see Appendix A, "Sun Cluster Object-Oriented Commands," in *Sun Cluster 3.1 - 3.2 Hardware Administration Manual for Solaris OS*.

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.

- **If you are using Sun Cluster 3.2, use the following commands:**

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

- **If you are using Sun Cluster 3.1, use the following command:**

  ```
  # scstat
  ```

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

- **If you are using Sun Cluster 3.2, use the following command:**

  ```
  # clnode evacuate nodename
  ```

- If you are using Sun Cluster 3.1, use the following command:

      # **scswitch -S -h** *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 *Sun Cluster System Administration Guide for Solaris OS*.

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 *Sun Cluster System Administration Guide for Solaris OS*.

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

   If you use the Solaris 8, Solaris 9, or Solaris 10 Operating System, Sun Connection Update Manager keeps you informed of the latest versions of patches and features. Using notifications and intelligent needs-based updating, Sun Connection helps improve operational efficiency and ensures that you have the latest software patches for your Sun software.

   You can download the Sun Connection Update Manager product for free by going to http://www.sun.com/download/products.xml?id=4457d96d.

   Additional information for using the Sun patch management tools is provided in *Solaris Administration Guide: Basic Administration* at http://docs.sun.com (http://docs.sun.com). Refer to the version of this manual for the 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 *Sun Cluster System Administration Guide for Solaris OS* 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 a list of patches that affect Sun Cluster, see the Sun Cluster Wiki Patch Klatch.

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.

If you use the Solaris 8, Solaris 9, or Solaris 10 Operating System, Sun Connection Update Manager keeps you informed of the latest versions of patches and features. Using notifications and intelligent needs-based updating, Sun Connection helps improve operational efficiency and ensures that you have the latest software patches for your Sun software.

You can download the Sun Connection Update Manager product for free by going to `http://www.sun.com/download/products.xml?id=4457d96d`.

Additional information for using the Sun patch management tools is provided in *Solaris Administration Guide: Basic Administration* at http://docs.sun.com (`http://docs.sun.com`). Refer to the version of this manual for the 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 *Sun Cluster System Administration Guide for Solaris OS* 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 a list of patches that affect Sun Cluster, see the Sun Cluster Wiki Patch Klatch.

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 *Sun Cluster System Administration Guide for Solaris OS*.

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

- **If you are using Sun Cluster 3.2, use the following command:**
  ```
  # cldevice populate
  ```

- **If you are using Sun Cluster 3.1, use the following command:**
  ```
  # scgdevs
  ```

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.**

- **If you are using Sun Cluster 3.2, use the following commands:**
  ```
  # cldevice clear
  # cldevice list -v
  ```

- **If you are using Sun Cluster 3.1, use the following commands:**
  ```
  # scdidadm -C
  # scdidadm -l
  ```

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.

- **If you are using Sun Cluster 3.2, use the following command:**
  ```
  # 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.

- **If you are using Sun Cluster 3.1, use the following command:**
  ```
  # scswitch -z -D devicegroup -h nodename
  ```

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.

- **If you are using Sun Cluster 3.2, use the following command:**

  ```
  # 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.

- **If you are using Sun Cluster 3.1, use the following command:**

  ```
  # scswitch -z -g resourcegroup -h nodename
  ```

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.

Next Steps   The best way to enable multipathing for a cluster is to install the multipathing software and enable multipathing before installing the Sun Cluster software and establishing the cluster. For this procedure, see "How to Install Sun Multipathing Software" in *Sun Cluster Software Installation Guide for Solaris OS*. If you need to add multipathing software to an established cluster, see "How to Install Sun Multipathing Software" in *Sun Cluster Software Installation Guide for Solaris OS* and follow the troubleshooting steps to clean up the device IDs.

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	"How to Create a Logical Volume" on page 35
Remove a LUN	"How to Remove a Logical Volume" on page 37

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 – 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 Solaris 10 OS, previously called Sun StorEdge Traffic Manager in the Solaris 9 OS, verify that it is installed and configured and the path to the storage device is functioning. To configure the Traffic Manager for the Solaris 9 OS, see the *Sun StorEdge Traffic Manager Installation and Configuration Guide*. To configure multipathing for the Solaris 10 OS, see the *Solaris Fibre Channel Storage Configuration and Multipathing Support Guide*.

This procedure provides the long forms of the Sun Cluster commands. Most commands also have short forms. Except for the forms of the command names, the commands are identical. For a list of the commands and their short forms, see Appendix A, "Sun Cluster Object-Oriented Commands," in *Sun Cluster 3.1 - 3.2 Hardware Administration Manual for Solaris OS*.

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 .**

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 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 the Traffic Manager for the Solaris 9 OS, see the *Sun StorEdge Traffic Manager Installation and Configuration Guide*. To configure multipathing for the Solaris 10 OS, 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.**

- **If you are using Sun Cluster 3.2, use the following command:**

    ```
    # cldevice populate
    ```

- **If you are using Sun Cluster 3.1, use the following command:**

    ```
    # scgdevs
    ```

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 *Sun Cluster System Administration Guide for Solaris OS*.

- To create a new resource or configure a running resource to use the new logical volume, see Chapter 2, "Administering Data Service Resources," in *Sun Cluster Data Services Planning and Administration Guide for Solaris OS*.

▼ 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 Sun Cluster commands. Most commands also have short forms. Except for the forms of the command names, the commands are identical. For a list of the commands and their short forms, see Appendix A, "Sun Cluster Object-Oriented Commands," in *Sun Cluster 3.1 - 3.2 Hardware Administration Manual for Solaris OS*.

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 one of the following commands.

- **If you are using Sun Cluster 3.2, use the following command:**

  ```
  # clquorum show
  ```

- **If you are using Sun Cluster 3.1, use the following command:**

  ```
   # scstat -q
  ```

For procedures about how to add and remove quorum devices, see Chapter 6, "Administering Quorum," in *Sun Cluster System Administration Guide for Solaris OS*.

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 Sun 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.

- **If you are using Sun Cluster 3.2, use the following commands:**
  ```
  # clresourcegroup status +
  # cldevicegroup status +
  ```

- **If you are using Sun Cluster 3.1, use the following command:**
  ```
  # scstat
  ```

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

- **If you are using Sun Cluster 3.2, use the following command:**
  ```
  # clnode evacuate nodename
  ```

- **If you are using Sun Cluster 3.1, use the following command:**
  ```
  # scswitch -S -h 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 *Sun Cluster System Administration Guide for Solaris OS*.

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

- **If you are using Sun Cluster 3.2, use the following command:**
  ```
  # devfsadm -C
  # cldevice clear
  ```

- **If you are using Sun Cluster 3.1, use the following command:**
  ```
  # devfsadm -C
  # scdidadm -C
  ```

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.

- **If you are using Sun Cluster 3.2, use the following command:**
  ```
  # 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.

- **If you are using Sun Cluster 3.1, use the following command:**

   # **scswitch -z -D** *devicegroup* **-h** *nodename*

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.

- **If you are using Sun Cluster 3.2, use the following command:**

   # **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.

- **If you are using Sun Cluster 3.1, use the following command:**

   # **scswitch -z -g** *resourcegroup* **-h** *nodename*

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 or
scdidadm -c command, the following error message appears on your console if the device ID
changed unexpectedly.

```
device id for nodename:/dev/rdsk/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 or scdidadm -R command for
each affected device.

TABLE 1–3 Task Map: Maintaining a Storage Array

Task	Information
Upgrade storage array firmware.	"How to Upgrade Storage Array Firmware" on page 42
Remove a storage array or partner group.	"How to Remove a Single-Controller Configuration" on page 43
	"How to Remove a Dual-Controller Configuration" on page 46
Replace a node-to-switch component. ■ Node-to-switch fiber-optic cable ■ FC host adapter ■ FC switch ■ GBIC or SFP	"Replacing a Node-to-Switch Component" on page 49
Replace a node's host adapter.	"How to Replace a Host Adapter" on page 52
Add a node to the storage array.	Sun Cluster system administration documentation
Remove a node from the storage device.	Sun 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` or `scdidadm -c` command, the following error message appears on your console if the device ID changed unexpectedly.

```
device id for nodename:/dev/rdsk/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` or `scdidadm -R` 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 Sun Cluster commands. Most commands also have short forms. Except for the forms of the command names, the commands are identical. For a list of the commands and their short forms, see Appendix A, "Sun Cluster Object-Oriented Commands," in *Sun Cluster 3.1 - 3.2 Hardware Administration Manual for Solaris OS*.

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.

- **If you are using Sun Cluster 3.2, use the following commands:**
  ```
  # clresourcegroup status -n NodeA
  # cldevicegroup status -n NodeA
  ```

- **If you are using Sun Cluster 3.1, use the following command:**
  ```
  # scstat
  ```

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 *Sun Cluster System Administration Guide for Solaris OS*.

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 Sun Cluster system administration documentation.

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

For the procedure about how to boot nodes, seeChapter 3, "Shutting Down and Booting a Cluster," in *Sun Cluster System Administration Guide for Solaris OS*.

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 *Sun Cluster System Administration Guide for Solaris OS*.

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 Sun Cluster system administration documentation.

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

For the procedure about how to boot nodes, seeChapter 3, "Shutting Down and Booting a Cluster," in *Sun Cluster System Administration Guide for Solaris OS*

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

- **If you are using Sun Cluster 3.2, use the following command:**
    ```
    # devfsadm -C
    # cldevice clear
    ```

- **If you are using Sun Cluster 3.1, use the following command:**
    ```
    # devfsadm -C
    # scdidadm -C
    ```

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.

- **If you are using Sun Cluster 3.2, use the following command:**
    ```
    # 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.

- **If you are using Sun Cluster 3.1, use the following command:**
    ```
    # scswitch -z -D devicegroup -h nodename
    ```

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.

- **If you are using Sun Cluster 3.2, use the following command:**
    ```
    # 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.

- **If you are using Sun Cluster 3.1, use the following command:**

  # **scswitch -z -g** *resourcegroup* **-h** *nodename*

▼ 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 Sun Cluster commands. Most commands also have short forms. Except for the forms of the command names, the commands are identical. For a list of the commands and their short forms, see Appendix A, "Sun Cluster Object-Oriented Commands," in *Sun Cluster 3.1 - 3.2 Hardware Administration Manual for Solaris OS*.

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.

- **If you are using Sun Cluster 3.2, use the following commands:**

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

- **If you are using Sun Cluster 3.1, use the following command:**

    ```
    # scstat
    ```

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 *Sun Cluster System Administration Guide for Solaris OS*.

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 Sun 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 *Sun Cluster System Administration Guide for Solaris OS*.

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 *Sun Cluster System Administration Guide for Solaris OS*.

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 Sun Cluster system administration documentation.

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

For more information about how to boot nodes, seeChapter 3, "Shutting Down and Booting a Cluster," in *Sun Cluster System Administration Guide for Solaris OS*.

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

- **If you are using Sun Cluster 3.2, use the following command:**
  ```
  # devfsadm -C
  # cldevice clear
  ```

- **If you are using Sun Cluster 3.1, use the following command:**
  ```
  # devfsadm -C
  # scdidadm -C
  ```

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.

- **If you are using Sun Cluster 3.2, use the following command:**
  ```
  # 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.

- **If you are using Sun Cluster 3.1, use the following command:**

  `# scswitch -z -D` *devicegroup* `-h` *nodename*

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.

- **If you are using Sun Cluster 3.2, use the following command:**

  `# 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.

- **If you are using Sun Cluster 3.1, use the following command:**

  `# scswitch -z -g` *resourcegroup* `-h` *nodename*

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 52.

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 50.
- If your cluster does *not* use multipathing, see "How to Replace a Node-to-Switch Component in a Cluster Without Multipathing" on page 50.

▼ 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 Sun Cluster commands. Most commands also have short forms. Except for the forms of the command names, the commands are identical. For a list of the commands and their short forms, see Appendix A, "Sun Cluster Object-Oriented Commands," in *Sun Cluster 3.1 - 3.2 Hardware Administration Manual for Solaris OS*.

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**.**

- **If you are using Sun Cluster 3.2, use the following commands:**

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

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

- **If you are using Sun Cluster 3.1, use the following command:**

  ```
  # scstat
  ```

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

- **If you are using Sun Cluster 3.2, use the following command:**

  ```
  # clnode evacuate nodename
  ```

- **If you are using Sun Cluster 3.1, use the following command:**

  ```
  # scswitch -S -h 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.

- **If you are using Sun Cluster 3.2, use the following command:**

  ```
  # 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.

- **If you are using Sun Cluster 3.1, use the following command:**

  ```
  # scswitch -z -D devicegroup -h nodename
  ```

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.

- **If you are using Sun Cluster 3.2, use the following command:**

  ```
  # 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.

- **If you are using Sun Cluster 3.1, use the following command:**

  `# `**`scswitch -z -g`** *resourcegroup* **`-h`** *nodename*

▼ 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 *Sun Cluster Hardware Administration Manual for Solaris OS* 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 *Sun Cluster 3.1 - 3.2 Hardware Administration Manual for Solaris OS*.

  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 Sun Cluster commands. Most commands also have short forms. Except for the forms of the command names, the commands are identical. For a list of the commands and their short forms, see Appendix A, "Sun Cluster Object-Oriented Commands," in *Sun Cluster 3.1 - 3.2 Hardware Administration Manual for Solaris OS*.

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.

- **If you are using Sun Cluster 3.2, use the following commands:**

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

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

- **If you are using Sun Cluster 3.1, use the following command:**

  ```
  # scstat
  ```

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

- **If you are using Sun Cluster 3.2, use the following command:**

  ```
  # clnode evacuate nodename
  ```

- **If you are using Sun Cluster 3.1, use the following command:**

  ```
  # scswitch -S -h 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 *Sun Cluster System Administration Guide for Solaris OS*.

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**.**

If you use the Solaris 8, Solaris 9, or Solaris 10 Operating System, Sun Connection Update Manager keeps you informed of the latest versions of patches and features. Using notifications and intelligent needs-based updating, Sun Connection helps improve operational efficiency and ensures that you have the latest software patches for your Sun software.

You can download the Sun Connection Update Manager product for free by going to http://www.sun.com/download/products.xml?id=4457d96d.

Additional information for using the Sun patch management tools is provided in *Solaris Administration Guide: Basic Administration* at http://docs.sun.com (`http://docs.sun.com`). Refer to the version of this manual for the 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 *Sun Cluster System Administration Guide for Solaris OS* 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 a list of patches that affect Sun Cluster, see the Sun Cluster Wiki Patch Klatch.

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 *Sun Cluster System Administration Guide for Solaris OS*.

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.

- **If you are using Sun Cluster 3.2, use the following command:**

  `# 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.

- **If you are using Sun Cluster 3.1, use the following command:**

  `# scswitch -z -D` *devicegroup* `-h` *nodename*

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.

- **If you are using Sun Cluster 3.2, use the following command:**

  `# 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.

- **If you are using Sun Cluster 3.1, use the following command:**

  `# `**`scswitch -z -g`** *resourcegroup* **`-h`** *nodename*

Index