Sun Cluster 3.1 - 3.2 With Sun StorEdge or StorageTek 9900 Series Storage Device Manual for Solaris OS

Sun Microsystems, Inc.
4150 Network Circle
Santa Clara, CA 95054
U.S.A.

Part No: 819–3021–12
January 2009, Revision 12

090130@21808

Contents

Preface

The *Sun Cluster 3.1 - 3.2 With Sun StorEdge or StorageTek 9900 Series Storage Device Manual for Solaris OS* provides procedures specific to Sun StorEdge™ and Sun StorageTek™ 9900 Series storage devices that are placed in a Sun™ Cluster environment.

Use this manual with any version of Sun Cluster 3.1 or 3.2 software. Unless otherwise noted, procedures are the same for all Sun Cluster 3.1 or 3.2 versions. See the "Revision History" on page 6 for a list of changes to this manual.

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.

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

How This Book Is Organized

This book contains the following chapters.

Chapter 1, "Installing and Configuring a Sun StorEdge or StorageTek 9900 Series Storage Array," discusses how to install Sun StorEdge 9900 storage systems and how to configure logical units on Sun StorEdge 9900 storage systems.

Chapter 2, "Installing Multipathing Software in a Sun StorEdge or StorageTek 9900 Series Storage Array," discusses how to add multipathing software to Sun StorEdge 9900 storage systems.

Chapter 3, "Maintaining a Sun StorEdge or StorageTek 9900 Series Storage Array," discusses how to maintain Sun StorEdge 9900 storage systems 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 or StorageTek 9900 Series Storage Device Manual for Solaris OS

Revision Date	New Information
February 2008	Chapter 1: added Sun StorageTek 9985V and 9990V to the list of 9900 Series storage arrays; updated steps in "How to Add a Storage Array to an Existing Cluster"; updated Step #3 in "How to Add a Logical Volume". Chapter 2: removed "Installing Multipathing Software" section; updated Before You Begin section in "How to Add Sun StorEdge Traffic Manager Software". Global: Renamed all chapters and book name to include StorageTek products.
March 2008	Replaced outdated information about SunSolve with information about Sun Connection Update Manager.
January 2009	Updated links in Preface to different versions of Sun Cluster documentation.

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 Installation and Configuration Guide Available on `http://www.sun.com/products-n-solutions/hardware/docs`	817-3674

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	`#`

◆ ◆ ◆ CHAPTER 1

Installing and Configuring a Sun StorEdge or StorageTek 9900 Series Storage Array

This chapter contains a *limited* set of procedures about how to install and configure Sun StorEdge 9900 Series storage arrays in a Sun Cluster environment. Contact your Sun service provider to perform tasks that are not documented in this chapter.

The StorEdge and StorageTek 9900 Series storage arrays includes the following storage arrays:

- Sun StorEdge 9910
- Sun StorEdge 9960
- Sun StorEdge 9970
- Sun StorEdge 9980
- Sun StorEdge 9985
- Sun StorageTek 9985V
- Sun StorEdge 9990
- Sun StorageTek 9990V

You can perform all the procedures in this chapter on all StorEdge 9900 Series storage arrays unless noted otherwise.

This chapter contains the following sections.

For conceptual information on multihost disks, see your Sun Cluster concepts documentation document.

Restrictions

When using storage-based replication, do not configure a replicated volume as a quorum device. Locate any quorum devices on an unreplicated volume. See "Using Storage-Based Data Replication" in *Sun Cluster System Administration Guide for Solaris OS* for more information on storage-based replication.

Installing Storage Arrays

The initial installation of a storage array in a new cluster must be performed by your Sun service provider.

▼ How to Add a Storage Array to an Existing Cluster

Use this procedure to add a new storage array to a running cluster.

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

If you need to add a storage array to more than two nodes, repeat Step 20 through Step 36 for each additional node that connects to the storage array.

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 Power on the storage array.

Note – The storage array requires approximately 10 minutes to boot.

Contact your service provider to power on the storage array.

2 If you plan to use multipathing software, verify that the storage array is configured for multipathing.

Contact your service provider to verify that the storage array is configured for multipathing.

3 Configure the new storage array.

Contact your service provider to create the desired logical volumes.

4 **If you need to install a host adapter in** `Node A`**, and if this host adapter is the first on** `Node A`**, contact your service provider to install the support packages and configure the drivers before you proceed to Step 5.**

Note – If you use multipathing software, each node requires two paths to the same set of LUNs.

If you do not need to install a host adapter, skip to Step 11.

5 **If your node is enabled with the Solaris dynamic reconfiguration (DR) feature, install the host adapter.**

For the procedure about how to install a host adapter, see the documentation that shipped with your host adapter or updated information on the manufacturer's web site.

If your node is not enabled with DR, you must shut down this node to install the host adapter. Proceed to Step 6.

6 **Determine the resource groups and device groups that are running on** `Node A` **and** `Node B`**.**

Record this information because you use this information in Step 35 and Step 36 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 -n NodeA[ NodeB ...]
# cldevicegroup status -n NodeA[ NodeB ...]
```

`-n` *NodeA*[*NodeB* ...]	The node or nodes for which you are determining resource groups and device groups.

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

```
# scstat
```

For more information, see your Sun Cluster system administration documentation.

7 **Shut down and power off** `Node A`**.**

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

8 **Install the host adapter in** `Node A`**.**

For the procedure about how to install a host adapter, see the documentation that shipped with your host adapter or updated information on the manufacturer's web site.

9 **Power on and boot** `Node A` **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*.

10 If necessary, 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*.

11 Attach the storage array to `Node A`.

Contact your service provider to install a fiber-optic cable between the storage array and your node.

12 Configure the storage array.

Contact your service provider to configure the storage array.

13 If you plan to install the Solaris I/O multipathing multipathing software, use the procedure in "How to Install Sun Multipathing Software" in *Sun Cluster Software Installation Guide for Solaris OS*.

14 SPARC: If you plan to install Sun StorEdge 9900 Dynamic Link Manager (Sun SDLM) software, install it and any required patches for Sun SDLM software support on `Node A`.

For the procedure about how to install the Sun SDLM software, see the documentation that shipped with your storage array.

15 To create the new Solaris device files and links on `Node A`, perform a reconfiguration boot.

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

16 **If you are using Solaris 8 or 9, on** Node A **configure all controllers that are affected by the new physical path.**

```
# cfgadm -c configure cN
```

17 **On** Node A**, use the appropriate multipathing software commands to verify that the same set of LUNs is visible to the expected controllers.**

18 **On** Node A**, update the paths to the device ID 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 commands:**

  ```
  # scgdevs
  ```

19 **(Optional) On** Node A**, verify that the device IDs are assigned to the new storage array.**

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

  ```
  # cldevice list -n NodeA -v
  ```

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

  ```
  # scdidadm -l
  ```

20 **If you need to install a host adapter in** Node B**, and if this host adapter is the first in** Node B**, contact your service provider to install the support packages and configure the drivers before you proceed to Step 21.**

Note – If you use multipathing software, each node requires two paths to the same set of LUNs.

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

If you do not need to install host adapters, skip to Step 26.

21 **If your node is enabled with the Solaris dynamic reconfiguration (DR) feature, install the host adapter.**

For the procedure about how to install a host adapter, see the documentation that shipped with your host adapter or updated information on the manufacturer's web site.

22 **If your node is not enabled with DR, shut down and power off** Node B.

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

23 Install the host adapter in `Node B`.

For the procedure about how to install a host adapter, see the documentation that shipped with your host adapter or updated information on the manufacturer's web site.

24 If necessary, upgrade the host adapter firmware on `Node B`.

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*.

25 Power on and boot `Node B` 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*.

26 Attach the storage array to `Node B`.

Contact your service provider to install a fiber-optic cable between the storage array and your node.

27 If you plan to install Solaris I/O multipathing multipathing software, use the procedure in "How to Install Sun Multipathing Software" in *Sun Cluster Software Installation Guide for Solaris OS*.

28 Install any required patches or software for Solaris I/O multipathing software support on `Node B`.

- For the procedure about how to install the Solaris I/O multipathing software, see the *Sun StorEdge Traffic Manager Installation and Configuration Guide* at `http://www.sun.com/products-n-solutions/hardware/docs/`.

- For a list of required patches for Solaris I/O multipathing software, see the *Sun StorEdge Traffic Manager Software Release Notes* at `http://www.sun.com/storage/san`.

29 SPARC: If you plan to use Sun StorEdge 9900 Dynamic Link Manager (Sun SDLM) multipathing software, install the software and any required patches for Sun SDLM software support on Node B**.**

For the procedure about how to install the Sun SDLM software, see the documentation that shipped with your storage array.

30 To create the new Solaris device files and links on Node B**, perform a reconfiguration boot.**

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

31 If you are using Solaris 8 or 9, on Node B**configure all controllers that are affected by the new physical path.**

```
# cfgadm -c configure cN
```

32 On Node B**, use the appropriate multipathing software commands to verify that the same set of LUNs is visible to the expected controllers.**

33 On Node B**, update the paths to the device ID 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 commands:**

  ```
  # scgdevs
  ```

34 (Optional) On Node B**, verify that the device IDs are assigned to the new LUNs.**

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

  ```
  # cldevice show
  ```

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

  ```
  # scdidadm -l
  ```

35 (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
  ```

36 (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.
resourcegroup	The resource group that is returned to the node or nodes.
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
  ```

37 Repeat Step 20 through Step 36 for each additional node that connects to the storage array.

38 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 and enable it before installing the Sun Cluster software and enabling 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 Add Solaris I/O multipathing Software" on page 29.

Configuring Storage Arrays

This section contains the procedures about how to configure a storage array in a Sun Cluster environment. The following table lists these procedures. For configuration tasks that are not cluster-specific, see the documentation that shipped with your storage array.

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/c`*X*`t`*Y*`d`*Z*`s`*N* `does not match physical device's id for d`*decimalnumber*`, 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 1–1 Task Map: Configuring a Storage Array

Task	Information
Add a logical volume.	See "How to Add a Logical Volume" on page 19.
Remove a logical volume.	See "How to Remove a Logical Volume" on page 20.

▼ How to Add a Logical Volume

Use this procedure to add a logical volume to a cluster. This procedure assumes that your service provider created your logical volume. This procedure also assumes that all nodes are booted and are attached to the storage array.

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*.

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

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

```
# devfsadm
```

2 **On each node connected to the storage array, use the appropriate multipathing software commands to verify that the same set of LUNs is visible to the expected controllers.**

3 **If you are running Veritas Volume Manager, update the list of devices on all nodes that are attached to the logical volume that you created in Step 2.**

See your Veritas Volume Manager documentation for information about how to use the `vxdctl enable` command. Use this command to update new devices (volumes) in your Veritas Volume Manager list of devices.

Note – You might need to install the Veritas Array Support Library (ASL) package that corresponds to the array. For more information, see your Veritas Volume Manager documentation.

If you are not running Veritas Volume Manager, proceed to Step 4.

4 **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
```

If a volume management daemon such as `vold` is running on your node, and you have a CD-ROM drive that is connected to the node, a `device busy` error might be returned even if no disk is in the drive. This error is expected behavior.

See Also To create a new resource or reconfigure a running resource to use the new logical volume, see your Sun Cluster data services collection.

▼ How to Remove a Logical Volume

Use this procedure to remove a logical volume. This procedure assumes all nodes are booted and are connected to the storage array. This storage array hosts the logical volume that you are removing.

This procedure defines `Node A` as the node with which you begin working. `Node B` is the remaining node.

If you need to remove a storage array from more than two nodes, repeat Step 9 through Step 12 for each additional node. Each node connects to the logical volume.

Caution – During this procedure, you lose access to the data that resides on the logical volume that you are removing.

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*.

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

1 **If necessary, back up all data. Migrate all resource groups and disk device groups to another node.**

2 **If the logical volume that you plan to remove is configured as a quorum device, choose and configure another device to be the new quorum device. Then remove the old quorum device.**

To determine whether this logical volume 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 your Sun Cluster system administration documentation.

3 **Run the appropriate Solaris Volume Manager commands or Veritas Volume Manager commands to remove the reference to the logical volume from any diskset or disk group.**

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

4 **If the cluster is running Veritas Volume Manager, update the list of devices on all nodes. These nodes are attached to the logical volume that you are removing.**

See your Veritas Volume Manager documentation for information about how to use the `vxdisk rm` command to remove devices (volumes) in your Veritas Volume Manager device list.

5 **Remove the logical volume.**

Contact your service provider to remove the logical volume.

6 **Determine the resource groups and device groups that are running on `Node A` and `Node B`.**

Record this information because you use this information in Step 11 and Step 12 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 -n NodeA[ NodeB ...]
# cldevicegroup status -n NodeA[ NodeB ...]
```

`-n` *NodeA*[*NodeB*…] The node or nodes for which you are determining resource groups and device groups.

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

```
# scstat
```

7 **Shut down and reboot `Node A` by using the `shutdown` command with the `-i6` option.**

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

8 **On `Node A`, update the `/devices` and `/dev` entries.**

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

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

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

```
# devfsadm -C
# scdidadm -C
```

9 **Shut down and reboot `Node B` by using the `shutdown` command with the `-i6` option.**

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*.

10 **On `Node B`, update the `/devices` and `/dev` entries.**

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

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

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

```
# devfsadm -C
# scdidadm -C
```

11 (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
  ```

12 (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 restored.
resourcegroup1[*resourcegroup2* …]	The resource group or groups that you are restoring to the node or nodes.

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

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

13 Repeat Step 9 through Step 12 for each additional node that connects to the logical volume.

See Also To create a logical volume, see "How to Add a Logical Volume" on page 19.

◆ ◆ ◆ CHAPTER 2

Installing Multipathing Software in a Sun StorEdge or StorageTek 9900 Series Storage Array

This chapter contains procedures on how to install and add multipathing software. Multipathing software enables you to define and control redundant physical paths to I/O devices such as storage arrays and networks. If the active path to a device becomes unavailable, the multipathing software can automatically switch to an alternate path to maintain availability. This capability is known as automatic failover. To maximize multipathing capabilities, your servers must be configured with redundant hardware. Redundant hardware is two or more host bus adapters from each node, that are connected to the same dual-ported storage array.

This chapter contains the following procedures.

Adding Multipathing Software

This section contains a procedure on how to add multipathing software in a running cluster. Choose one of the following multipathing solutions.

- SPARC: Sun StorEdge 9900 Dynamic Link Manager software
- Solaris I/O multipathing software

▼ SPARC: How to Add Sun StorEdge 9900 Dynamic Link Manager Software

Use this procedure to add Sun StorEdge 9900 Dynamic Link Manager (Sun SDLM) software in a running cluster.

Do not use this procedure to convert from host-based mirroring to a multipathing solution. For the procedure on how to convert from host-based mirroring to a multipathing solution, contact your Sun service provider.

Perform this procedure on one node at a time. This procedure defines `Node N` as the node on which you are installing the multipathing software.

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.

Before You Begin This procedure assumes that you installed and configured your storage array.

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

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 these nodes.

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

  ```
  # clresourcegroup status -n NodeN
  # cldevicegroup status -n NodeN
  ```

  `-n` *NodeN* 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
  ```

For more information, see the Sun Cluster system administration documentation.

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
  ```

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

  ```
  # scswitch -S -h from-node
  ```

For more information, see the Sun Cluster system administration documentation.

3 If you need to install additional physical paths between `Node N` and the storage, shut down and power off `Node N`.

Note – If you use multipathing software, each node requires two paths to the same set of LUNs.

If you do not need additional physical paths, skip to Step 7.

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

4 **Install the host adapters and the cables between `Node N` and the storage.**

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

5 **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*.

6 **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*.

7 **Install and configure Sun SDLM software on `Node N`, and apply required patches for Sun SDLM software support on `Node N`.**

For instructions on how to install and configure the Sun SDLM software, see the documentation that shipped with your storage array.

8 **If you did not perform a reconfiguration reboot when you configured Sun SDLM software, perform a reconfiguration reboot to create the new Solaris device files and links on** Node N**.**

```
# boot -r
```

9 **On** Node N**, update the paths to the device ID 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 commands:**

```
# scgdevs
```

10 **(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
```

11 **(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
```

12 **On all remaining nodes and one node at a time, repeat Step 2 through Step 11.**

▼ How to Add Solaris I/O multipathing Software

The best way to enable multipathing for a cluster is to install the multipathing software (or enable multipathing in Solaris 10) 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, use this procedure and be careful to perform the steps that clean up the device IDs.

Do not use this procedure to convert from host-based mirroring to a multipathing solution. For the procedure on how to convert from host-based mirroring to a multipathing solution, contact your Sun service provider.

Perform this procedure on one node at a time. This procedure defines `Node N` as the node on which you are installing the multipathing software.

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*.

Before You Begin This procedure assumes that you have already installed and configured your storage array and added multipathing software. If you use Solaris 8 or Solaris 9, you must install Sun StorEdge TrafficManager Software to enable multipathing. (In Solaris 10, you just need to enable the multipathing software.) For instructions, see "How to Install Sun Multipathing Software" in *Sun Cluster Software Installation Guide for Solaris OS*.

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

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

Record this information because you use this information 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 -n NodeN
# cldevicegroup status -n NodeN
```

`-n` *NodeN* 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
```

For more information, see the Sun Cluster system administration documentation.

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
  ```

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

  ```
  # scswitch -S -h from-node
  ```

For more information, see the Sun Cluster system administration documentation.

3 **If you need to install additional physical paths between** `Node N` **and the storage, shut down and power off** `Node N`**.**

Note – If you use multipathing software, each node requires two paths to the same set of LUNs. For the full procedure on how to shut down and power off a node, see Sun Cluster system administration documentation.

---

4 **Install the host adapters and the cables between** `Node N` **and the storage.**

Note – If you use multipathing software, each node requires two paths to the same set of LUNs.

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

5 **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*.

6 **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*.

7 Install any required patches or software for Solaris I/O multipathing software support on `Node N`.

For instructions on installing the software, see the *Sun StorEdge Traffic Manager Installation and Configuration Guide* at `http://www.sun.com/products-n-solutions/hardware/docs/`.

8 If you are using Solaris OS version 9, edit the `/kernel/drv/scsi_vhci.conf` file to activate the Solaris I/O multipathing software. Set the `mpxio-disable` parameter to `no`.

```
mpxio-disable="no"
```

9 (Optional) Activate multipathing functionality.

- **If you are using the Solaris 10 operating system, issue the following commands to enable multipathing:**

  ```
  # /usr/sbin/stmsboot -e
  ```

- **If you are using the Solaris 8 or 9 operating system, edit the `/kernel/drv/scsi_vhci.conf` file.**

  Set the `mpxio-disable` parameter to `no`.

  For more information, see the STMS software documentation.

10 If you are using Solaris 8 or Solaris 9 OS software, boot `Node N` into cluster mode.

For more information on booting nodes, see Chapter 3, "Shutting Down and Booting a Cluster," in *Sun Cluster System Administration Guide for Solaris OS*.

11 On each node connected to the storage array, use the appropriate multipathing software commands to verify that the same set of LUNs is visible to the expected controllers.

12 If you are using Solaris 8 or Solaris 9 OS software, on `Node N` configure all controllers that are affected by the new physical path.

```
# cfgadm -c configure cN
```

13 On Node N, update the paths to the device ID instances.

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

```
 # cldevice clear
# cldevice refresh
```

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

```
# scdidadm -C
# scdidadm -r
```

`scdidadm -C` and `-r`	Update DID mappings with new device names while preserving DID instance numbers for disks that are connected to multiple cluster nodes. DID instance numbers of the local disks might not be preserved. For this reason, the DID disk names for local disks might change.

14 (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
```

15 (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
```

16 **On each node connected to the storage array, use the `format(1M)` command. Use this command to verify that you see half the number of disks you saw in Step 11.**

```
# format
```

See the `format` command man page for more information about by using the command.

17 **On all remaining nodes and one node at a time, repeat Step 2 through Step 16.**

◆ ◆ ◆ CHAPTER 3

Maintaining a Sun StorEdge or StorageTek 9900 Series Storage Array

This chapter contains a limited set of procedures about how to maintain a storage array. Contact your service provider to add, remove, or replace any storage array components.

This chapter contains the following 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.

Maintaining Storage Arrays

This section contains the procedures for maintaining a storage system in a running cluster. Table 3–1 lists these procedures.

TABLE 3–1 Task Map: Maintaining a Storage Array

Task	Information
Remove a storage array.	"How to Remove a Storage Array" on page 36
Add a node to the storage array.	Sun Cluster system administration documentation
Remove a node from the storage array.	Sun Cluster system administration documentation
Replace a node's host adapter.	"How to Replace a Host Adapter" on page 40
Replace an FC switch or storage array-to-switch component.	"How to Replace an FC Switch or Storage Array-to-Switch Component" on page 43
Replace a node-to-switch/storage component.	"Replacing a Node-to-Switch Component" on page 43

▼ How to Remove a Storage Array

Use this procedure to permanently remove a storage array. This procedure provides the flexibility to remove the host adapters from the nodes that are attached to 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.

If you need to remove a storage array from more than two nodes, repeat Step 15 through Step 23 for each additional node that connects to the storage array.

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 **If necessary, back up all data and migrate all resource groups and disk device groups to another node.**

2 **If the storage array that you plan to remove contains a quorum device, choose and configure another device to be the new quorum device. Then remove the old quorum device.**

To determine whether this logical volume 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
```

To add or remove a quorum device in your configuration, see your Sun Cluster system administration documentation.

3 **If necessary, detach the submirrors from the storage array that you are removing in order to stop all I/O activity to the storage array.**

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

4 **Run the appropriate Solaris Volume Manager or Veritas Volume Manager commands to remove the references to the logical volumes from any diskset or disk group.**

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

5 **Are your nodes enabled with the Solaris dynamic reconfiguration (DR) feature?**

- **If yes, disconnect the fiber-optic cables and, if desired, remove the host adapters from both nodes. Then perform Step 23 on each node that was connected to the storage array**

  For the procedure about how to remove a host adapter, see the documentation that shipped with your host adapter or updated information on the manufacturer's web site.

- **If no, proceed to Step 6.**

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

Record this information because you use this information in Step 21 and Step 22 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 -n NodeA[ NodeB ...]
# cldevicegroup status -n NodeA[ NodeB ...]
```

-n *NodeA*[*NodeB* …] The node or nodes for which you are determining resource groups and device groups.

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

```
# scstat
```

For more information, see your Sun Cluster system administration documentation.

7 **If you want to remove any multipathing software, remove the multipathing software packages.**

8 **Shut down** Node A**.**

For the procedure about how to shut down a node, see your Sun Cluster system administration documentation.

9 **Disconnect the fiber-optic cable between the storage array and** Node A**.**

10 **If you do not want to remove host adapters from** Node A**, skip to Step 13.**

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

12 **Remove the host adapter from** Node A**.**

For the procedure about how to remove host adapters, see the documentation that shipped with your host adapter or updated information on the manufacturer's web site.

13 **Power on** Node A **and allow it to boot into cluster mode.**

For more information, see the documentation that shipped with your server. 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*.

14 **On** Node A**, update the device namespace.**

```
# devfsadm -C
```

15 **Shut down** Node B**.**

For the procedure about how to shut down a node, see your Sun Cluster system administration documentation.

16 **Disconnect the fiber-optic cable between the storage array and** `Node B`**.**

17 **If you do not want to remove host adapters from** `Node B`**, skip to Step 20.**

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

19 **Remove the host adapter from** `Node B`**.**

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

20 **Power on** `Node B` **and allow it to boot into cluster mode.**

For more information, see the documentation that shipped with your server. 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*.

21 **(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
```

22 **(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
```

23 **On** Node B**, update the device namespace.**

```
# devfsadm -C
```

24 **Repeat Step 15 through Step 22 for each additional node that connects to the storage array.**

25 **From one node, remove device ID references to the storage array that was removed.**

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

  ```
  # cldevice clear
  ```

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

  ```
  # scdidadm -C
  ```

▼ 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
  ```

How to Replace an FC Switch or Storage Array-to-Switch Component

Use this procedure to replace an FC switch, or the following storage array-to-switch components in a running cluster.

- Fiber-optic cable that connects an FC switch to a storage array
- GBIC on an FC switch, connecting to a storage array
- FC switch

- **Replace the component by using the following references.**
  - For the procedure about how to replace a fiber-optic cable between a storage array and an FC switch, see the documentation that shipped with your switch hardware.
  - For the procedure about how to replace a GBIC on an FC switch, see the documentation that shipped with your switch hardware.
  - For the procedure about how to replace an SFP on the storage array, contract your service provider.
  - For the procedure about how to replace an FC switch, see the documentation that shipped with your switch hardware.

Note – If you are replacing an FC switch and you intend to save the switch configuration for restoration to the replacement switch, do not connect the cables to the replacement switch until *after* you recall the Fabric configuration to the replacement switch. For more information about how to save and recall switch configurations see the documentation that shipped with your switch hardware.

Before you replace an FC switch, be sure that the `probe_timeout` parameter of your data service software is set to more than 90 seconds. Increasing the value of the `probe_timeout` parameter to more than 90 seconds avoids unnecessary resource group restarts when one of the FC switches is powered off.

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 40.

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 44.

▼ 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
  ```

Index

R

S

T