

Oracle® Solaris Cluster 3.3 With Sun StorEdge 9900 Series Storage Device Manual

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

This software and related documentation are provided under a license agreement containing restrictions on use and disclosure and are protected by intellectual property laws. Except as expressly permitted in your license agreement or allowed by law, you may not use, copy, reproduce, translate, broadcast, modify, license, transmit, distribute, exhibit, perform, publish, or display any part, in any form, or by any means. Reverse engineering, disassembly, or decompilation of this software, unless required by law for interoperability, is prohibited.

The information contained herein is subject to change without notice and is not warranted to be error-free. If you find any errors, please report them to us in writing.

If this is software or related software documentation that is delivered to the U.S. Government or anyone licensing it on behalf of the U.S. Government, the following notice is applicable:

U.S. GOVERNMENT RIGHTS Programs, software, databases, and related documentation and technical data delivered to U.S. Government customers are "commercial computer software" or "commercial technical data" pursuant to the applicable Federal Acquisition Regulation and agency-specific supplemental regulations. As such, the use, duplication, disclosure, modification, and adaptation shall be subject to the restrictions and license terms set forth in the applicable Government contract, and, to the extent applicable by the terms of the Government contract, the additional rights set forth in FAR 52.227-19, Commercial Computer Software License (December 2007). Oracle America, Inc., 500 Oracle Parkway, Redwood City, CA 94065.

This software or hardware is developed for general use in a variety of information management applications. It is not developed or intended for use in any inherently dangerous applications, including applications which may create a risk of personal injury. If you use this software or hardware in dangerous applications, then you shall be responsible to take all appropriate fail-safe, backup, redundancy, and other measures to ensure its safe use. Oracle Corporation and its affiliates disclaim any liability for any damages caused by use of this software or hardware in dangerous applications.

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

AMD, Opteron, the AMD logo, and the AMD Opteron logo are trademarks or registered trademarks of Advanced Micro Devices. Intel and Intel Xeon are trademarks or registered trademarks of Intel Corporation. All SPARC trademarks are used under license and are trademarks or registered trademarks of SPARC International, Inc. UNIX is a registered trademark licensed through X/Open Company, Ltd.

This software or hardware and documentation may provide access to or information on content, products, and services from third parties. Oracle Corporation and its affiliates are not responsible for and expressly disclaim all warranties of any kind with respect to third-party content, products, and services. Oracle Corporation and its affiliates will not be responsible for any loss, costs, or damages incurred due to your access to or use of third-party content, products, or services.

Contents

Preface	5
1 Installing and Configuring a Sun StorEdge or StorageTek 9900 Series Storage Array	11
Restrictions	12
Installing Storage Arrays	12
▼ How to Add a Storage Array to an Existing Cluster	12
Configuring Storage Arrays	17
▼ How to Add a Logical Volume	18
▼ How to Remove a Logical Volume	19
2 Enabling Multipathing Software in a Sun StorEdge or StorageTek 9900 Series Storage Array	23
Enabling Solaris I/O Multipathing Software	23
3 Maintaining a Sun StorEdge or StorageTek 9900 Series Storage Array	25
Maintaining Storage Arrays	25
▼ How to Remove a Storage Array	26
▼ How to Replace a Host Adapter	29
▼ How to Replace an FC Switch or Storage Array-to-Switch Component	32
Replacing a Node-to-Switch Component	32
Index	35

Preface

The *Oracle Solaris Cluster 3.3 With Sun StorEdge 9900 Series Storage Device Manual* provides procedures specific to Oracle's Sun StorEdge and Sun StorageTek 9900 Series storage devices that are placed in an Oracle Solaris Cluster environment.

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

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

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

Who Should Use This Book

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

This document is intended for experienced system administrators with extensive knowledge of Oracle software and hardware. Do not use this document as a planning or 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, “Enabling 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 Oracle Solaris Cluster 3.3 With Sun StorEdge 9900 Series Storage Device Manual

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.
September 2010	Updated product name; removed old CLI commands for 3.3 release; removed section with steps on how to add multipathing for OS versions prior to Oracle Solaris 10.

Related Documentation

The following books provide conceptual information or procedures to administer hardware and applications. If you plan to use this documentation in a hardcopy format, ensure that you have these books available for your reference.

The following books support the Oracle Solaris Cluster 3.3 release. You can also access the documentation for the Sun Cluster 3.1 and 3.2 releases. All Sun Cluster and Oracle Solaris Cluster documentation is available at <http://docs.sun.com>. Documentation that is not available at <http://docs.sun.com> is listed with the appropriate URL.

TABLE P-2 Oracle Solaris Cluster and Sun Cluster Documentation

Documentation

[Oracle Solaris Cluster 3.3](#)
[Sun Cluster 3.2](#)
[Sun Cluster 3.1](#)

Using UNIX Commands

This document contains information about commands that are used to install, configure, or upgrade an Oracle Solaris Cluster configuration. This document might not contain complete information about basic UNIX commands and procedures such as shutting down the system, booting the system, and configuring devices.

See one or more of the following sources for this information:

- Online documentation for the Oracle Solaris Operating System (Oracle Solaris OS)
- Other software documentation that you received with your system
- Oracle Solaris Operating System man pages

Getting Help

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

- Your name and email address (if available)
- Your company name, address, and phone number
- The model number and serial number of your systems
- The release number of the operating environment (for example, Oracle Solaris 10)
- The release number of Oracle Solaris Cluster (for example, Oracle Solaris Cluster 3.3)

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

Command	Function
<code>prtconf -v</code>	Displays the size of the system memory and reports information about peripheral devices
<code>psrinfo -v</code>	Displays information about processors

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

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

Documentation, Support, and Training

See the following web sites for additional resources:

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

Oracle Welcomes Your Comments

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

[Oracle Technology Network \(http://www.oracle.com/technetwork/index.html\)](http://www.oracle.com/technetwork/index.html) offers a range of resources related to Oracle software:

- Discuss technical problems and solutions on the [Discussion Forums \(http://forums.oracle.com\)](http://forums.oracle.com).
- Get hands-on step-by-step tutorials with [Oracle By Example \(http://www.oracle.com/technology/obe/start/index.html\)](http://www.oracle.com/technology/obe/start/index.html).
- Download [Sample Code \(http://www.oracle.com/technology/sample_code/index.html\)](http://www.oracle.com/technology/sample_code/index.html).

Typographic Conventions

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

TABLE P-3 Typographic Conventions

Typeface	Meaning	Example
AaBbCc123	The names of commands, files, and directories, and onscreen computer output	Edit your <code>.login</code> file. Use <code>ls -a</code> to list all files. <code>machine_name%</code> you have mail.
AaBbCc123	What you type, contrasted with onscreen computer output	<code>machine_name%</code> su Password:
<i>aabbcc123</i>	Placeholder: replace with a real name or value	The command to remove a file is <i>rm filename</i> .
<i>AaBbCc123</i>	Book titles, new terms, and terms to be emphasized	Read Chapter 6 in the <i>User's Guide</i> . <i>A cache</i> is a copy that is stored locally. Do <i>not</i> save the file. Note: Some emphasized items appear bold online.

Shell Prompts in Command Examples

The following table shows the default UNIX system prompt and superuser prompt for shells that are included in the Oracle Solaris OS. Note that the default system prompt that is displayed in command examples varies, depending on the Oracle Solaris release.

TABLE P-4 Shell Prompts

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

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 an Oracle Solaris Cluster environment. Contact your Oracle 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:

- Oracle's Sun StorEdge 9910
- Oracle's Sun StorEdge 9960
- Oracle's Sun StorEdge 9970
- Oracle's Sun StorEdge 9980
- Oracle's Sun StorEdge 9985
- Oracle's Sun StorageTek 9985V
- Oracle's Sun StorEdge 9990
- Oracle's 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.

- [“Restrictions” on page 12](#)
- [“Installing Storage Arrays” on page 12](#)
- [“Configuring Storage Arrays” on page 17](#)

For conceptual information on multihost disks, see your Oracle Solaris 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 Within a Cluster](#)” in *Oracle Solaris Cluster System Administration Guide* 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 Oracle 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 18](#) through [Step 31](#) for each additional node that connects to the storage array.

This procedure provides the long forms of the Oracle Solaris Cluster commands. Most commands also have short forms. Except for the forms of the command names, the commands are identical.

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

1 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 Oracle 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 30](#) and [Step 31](#) of this procedure to return resource groups and device groups to these nodes.

Use the following command:

```
# 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.
```

For more information, see your Oracle Solaris 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 Oracle Solaris 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 *Oracle Solaris Cluster System Administration Guide*.

- 10 **If necessary, upgrade the host adapter firmware on Node A.**

The Oracle Enterprise Manager Ops Center 2.5 software helps you patch and monitor your data center assets. Oracle Enterprise Manager Ops Center 2.5 helps improve operational efficiency

and ensures that you have the latest software patches for your software. Contact your Oracle representative to purchase Oracle Enterprise Manager Ops Center 2.5.

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

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

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

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 Oracle Solaris 10 automatically installs Solaris I/O multipathing (MPxIO). If you plan to use it, verify that the paths to the storage device are functioning. See “How to Install Solaris Software” in *Oracle Solaris Cluster Software Installation Guide* and *Solaris Fibre Channel Storage Configuration and Multipathing Support Guide*.

14 To create the new Oracle 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 *Oracle Solaris Cluster System Administration Guide*.

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

16 On Node A, update the paths to the device ID instances.

```
# cldevice populate
```

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

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

- 18 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 19.**

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 24](#).

- 19 If your node is enabled with the Oracle 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.

- 20 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 Oracle Solaris Cluster system administration documentation.

- 21 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.

- 22 If necessary, upgrade the host adapter firmware on Node B.**

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

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

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

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

23 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 *Oracle Solaris Cluster System Administration Guide*.

24 Attach the storage array to Node B.

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

25 Verify that the multipathing paths to the storage device are functioning.

Oracle Solaris 10 automatically installs Solaris I/O multipathing (MPxIO) for the Oracle Solaris 10 OS, which was previously called Sun StorEdge Traffic Manager in the Solaris 9 OS. For more information, see the [Solaris Fibre Channel Storage Configuration and Multipathing Support Guide](#).

26 To create the new Oracle 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 *Oracle Solaris Cluster System Administration Guide*.

27 On Node B, use the appropriate multipathing software commands to verify that the same set of LUNs is visible to the expected controllers.**28 On Node B, update the paths to the device ID instances.**

```
# cldevice populate
```

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

```
# cldevice show
```

30 (Optional) Restore the device groups to the original node.

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

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

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

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

31 (Optional) Restore the resource groups to the original node.

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

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

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

- 32 Repeat Step 18 through Step 31 for each additional node that connects to the storage array.**
- 33 Perform volume management administration to incorporate the new logical volumes into the cluster.**

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

Configuring Storage Arrays

This section contains the procedures about how to configure a storage array in an Oracle Solaris 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` command, the following error message appears on your console if the device ID changed unexpectedly.

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

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

TABLE 1-1 Task Map: Configuring a Storage Array

Task	Information
Add a logical volume.	See “ How to Add a Logical Volume ” on page 18.
Remove a logical volume.	See “ How to Remove a Logical Volume ” on page 19.

▼ 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 Oracle Solaris Cluster commands. Most commands also have short forms. Except for the forms of the command names, the commands are identical.

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

1 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 `vxctl 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.

```
# cldevice populate
```

If a volume management daemon such as `vol` 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 Oracle Solaris 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 Oracle Solaris Cluster commands. Most commands also have short forms. Except for the forms of the command names, the commands are identical.

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

- 1 **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 the following command.

```
#c1quorum show
```


For procedures about how to add and remove quorum devices, see your Oracle Solaris 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 `vxdiskrm` 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.

Use the following command:

```
# 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.
```

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 Oracle Solaris Cluster system administration documentation.

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

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

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 *Oracle Solaris Cluster System Administration Guide*.

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

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

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.

```
# 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.
```

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.

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

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

- 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 18.

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

This chapter contains general instructions on installing and enabling 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 section.

- [“Enabling Solaris I/O Multipathing Software” on page 23](#)

Enabling Solaris I/O Multipathing Software

Solaris I/O multipathing (MPxIO) for the Oracle Solaris 10 OS was previously called Sun StorEdge Traffic Manager in the Solaris 9 OS.

To enable multipathing for a cluster, install the Oracle Solaris 10 OS before installing the Oracle Solaris Cluster software and establishing the cluster. For this procedure, see [“How to Install Solaris Software”](#) in *Oracle Solaris Cluster Software Installation Guide*.

After you install the OS, you must enable MPxIO and verify that the paths to the storage device are functioning. To enable multipathing, see the *Solaris Fibre Channel Storage Configuration and Multipathing Support Guide*.

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 Oracle service provider.

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:

- “How to Remove a Storage Array” on page 26
- “How to Replace a Host Adapter” on page 29
- “How to Replace an FC Switch or Storage Array-to-Switch Component” on page 32
- “Replacing a Node-to-Switch Component” on page 32

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

```
device id for nodename:/dev/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` 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 26
Add a node to the storage array.	Oracle Solaris Cluster system administration documentation
Remove a node from the storage array.	Oracle Solaris Cluster system administration documentation
Replace a node's host adapter.	“How to Replace a Host Adapter” on page 29
Replace an FC switch or storage array-to-switch component.	“How to Replace an FC Switch or Storage Array-to-Switch Component” on page 32
Replace a node-to-switch/storage component.	“Replacing a Node-to-Switch Component” on page 32

▼ 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 Oracle Solaris Cluster commands. Most commands also have short forms. Except for the forms of the command names, the commands are identical.

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

- 1 **If necessary, back up all 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 the following command.

```
#clquorum show
```

To add or remove a quorum device in your configuration, see your Oracle Solaris 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.

Use the following command:

```
# 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.

For more information, see your Oracle Solaris 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 Oracle Solaris 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 *Oracle Solaris Cluster System Administration Guide*.

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 Oracle Solaris 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 *Oracle Solaris Cluster System Administration Guide*.

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.

```
# 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.

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.

```
# 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.

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.**

```
# cldevice clear
```

▼ How to Replace a Host Adapter

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

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

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

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

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

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

This procedure provides the long forms of the Oracle Solaris Cluster commands. Most commands also have short forms. Except for the forms of the command names, the commands are identical.

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

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

Record this information because you use this information in [Step 10](#) and [Step 11](#) of this procedure to return resource groups and device groups to Node A.

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

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

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

```
# clnode evacuate nodename
```

4 Shut down Node A.

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

5 Power off Node A.

6 Replace the failed host adapter.

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

- 7 If you need to upgrade the node's host adapter firmware, boot Node A into noncluster mode by adding -x to your boot instruction. Proceed to Step 8.**

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

- 8 Upgrade the host adapter firmware on Node A.**

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

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

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

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

- 9 Boot Node A into cluster mode.**

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

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

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

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

-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.

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

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

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

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

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

▼ How to 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 29](#).

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 33](#).
- If your cluster does *not* use multipathing, see [“How to Replace a Node-to-Switch Component in a Cluster Without Multipathing” on page 33](#).

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

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

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

Before You Begin

This procedure provides the long forms of the Oracle Solaris Cluster commands. Most commands also have short forms. Except for the forms of the command names, the commands are identical.

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

You have completed this procedure.

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

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

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

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

```
# clnode evacuate nodename
```

- 5 **Replace the node-to-switch component.**

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

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

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

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

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

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

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

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

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

`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.

Index

A

adding
 nodes, 26
 storage array s, 12–17
arrays, *See* storage array s

C

configuring, storage arrays, 17–21
creating, logical volumes, 18

D

deleting, *See* removing

H

HBAs, replacing, 29–32
help, 7–8
host adapters
 replacing, 26, 29–32

I

installing, storage arrays, 12–17

L

logical unit numbers, *See* logical volumes
logical units, *See* logical volumes
logical volumes
 creating, 18
 removing, 19–21
LUNs, *See* logical volumes

M

maintaining, storage arrays, 25–34
MPxIO, enabling, 23
multipathing, enabling, 23

N

node-to-switch components, replacing, 32–34
node-to-switch/storage components, replacing, 26
nodes
 adding, 26
 removing, 26

R

removing
 logical volumes, 19–21
 nodes, 26
 storage arrays, 26–29
replacing
 host adapters, 26, 29–32

replacing (*Continued*)

- node-to-switch components, 32–34
- node-to-switch/storage components, 26
- storage array -to-switch components, 26
- storage array-to-switch components, 32

S

storage array-to-switch components

- replacing, 26, 32

storage arrays

- adding, 12–17
- configuring, 17–21
- installing, 12–17
- maintaining, 25–34
- removing, 26–29

storage-based replication, 12

T

technical support, 7–8

TrueCopy, 12