

# Oracle® Solaris Cluster 3.3 With SCSI JBOD 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

---

<b>Preface</b> .....	5
<b>1 Installing a SCSI JBOD Storage Device</b> .....	11
Installing a Storage Array .....	11
Setting SCSI Initiator IDs .....	13
▼ SPARC: How to Install a Storage Array in a New SPARC Based Cluster .....	13
▼ x86: How to Install a Storage Array in a New x86 Based Cluster .....	18
▼ SPARC: How to Add a Storage Array to an Existing SPARC Based Cluster .....	24
▼ x86: How to Add a Storage Array to an Existing x86 Based Cluster .....	30
<b>2 Maintaining a SCSI JBOD Storage Device</b> .....	35
Maintaining Storage Arrays .....	36
FRUs That <i>Do Not</i> Require Oracle Solaris Cluster Maintenance Procedures .....	37
▼ How to Remove a Storage Array .....	39
▼ How to Replace the Chassis .....	41
Disconnecting and Reconnecting a Node from Shared Storage .....	43
▼ How to Replace a SCSI Cable .....	47
Replacing a Host Adapter .....	48
▼ How to Add a Disk Drive .....	62
▼ How to Remove a Disk Drive .....	63
▼ How to Replace a Disk Drive <i>Without</i> Oracle Real Application Clusters .....	65
▼ SPARC: How to Replace a Disk Drive <i>With</i> Oracle Real Application Clusters .....	68
▼ How to Upgrade Disk Drive Firmware .....	71
▼ How to Upgrade Host Adapter Firmware .....	72
<b>3 Cabling Diagrams</b> .....	75
Sun StorEdge 3120 Cabling Diagrams .....	75

Installing a Sun StorEdge 3120 Storage Array .....	75
Sun StorEdge 3310 and 3320 SCSI Cabling Diagrams .....	76
Installing a Sun StorEdge 3310 or 3320 Storage Array .....	76
Adding a Sun StorEdge 3310 or 3320 Storage Array .....	77
SPARC: Sun StorEdge D1000 Cabling Diagrams .....	78
SPARC: Installing a Sun StorEdge D1000 Storage Array .....	78
SPARC: Adding a Sun StorEdge D1000 Storage Array .....	79
SPARC: Replacing a Sun StorEdge D1000 Storage Array .....	81
SPARC: Removing a Sun StorEdge D1000 Storage Array .....	81
SPARC: Sun StorEdge Multipack Cabling Diagrams .....	82
SPARC: Installing a Sun StorEdge Multipack Storage Array .....	82
SPARC: Adding a Sun StorEdge Multipack Storage Array .....	83
SPARC: Replacing a Sun StorEdge Multipack Storage Array .....	84
SPARC: Removing a Sun StorEdge Multipack Storage Array .....	85
SPARC: Sun StorEdge Netra D130/S1 Cabling Diagrams .....	86
SPARC: Installing a Sun StorEdge Netra D130/S1 Storage Array .....	86
SPARC: Adding a Sun StorEdge Netra D130/S1 Storage Array .....	89
SPARC: Replacing a Sun StorEdge Netra D130/S1 Storage Array .....	90
SPARC: Removing a Sun StorEdge Netra D130/S1 Storage Array .....	91
SPARC: Sun StorEdge D2 Cabling Diagrams .....	91
SPARC: Installing a Sun StorEdge D2 Storage Array .....	91
SPARC: Adding a Sun StorEdge D2 Storage Array .....	93
<b>Index .....</b>	<b>97</b>

# Preface

---

The *Oracle Solaris Cluster 3.3 With SCSI JBOD Storage Device Manual* provides procedures specific to SCSI JBOD storage devices in an Oracle Solaris Cluster environment. Use this manual with any version of Oracle Solaris Cluster 3.3 software on SPARC based systems and with any supported versions of Oracle Solaris Cluster 3.3 software on x86 based systems.

---

**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.

---

---

**SPARC only** – In this document, references to Oracle Real Application Clusters also apply to Oracle Parallel Server, unless otherwise stated.

---

## 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 discusses how to install SCSI JBOD storage devices.
- Chapter 2 describes how to maintain SCSI JBOD storage devices in a running cluster.
- Chapter 3 contains cabling diagrams for SCSI JBOD storage devices.

## 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 SCSI JBOD Storage Device Manual

Revision Date	New Information
March 2008	Improved the accuracy of “ <a href="#">How to Replace a Disk Drive Without Oracle Real Application Clusters</a> ” on page 65 by revising selected steps.
January 2009	Updated links in Preface to lead to different versions of Sun Cluster and Oracle Solaris Cluster documentation.
September 2010	Changed product name throughout the book; removed old CLI commands.

## Related Documentation

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

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

The related books are organized by the following products and subjects:

- [Sun StorEdge 3120 storage array](#)
- [Sun StorEdge 3310 SCSI storage array](#)
- [SPARC: Sun StorEdge D1000 storage array](#)
- [SPARC: Sun StorEdge D130/S1 storage array](#)
- [SPARC: Sun StorEdge D2 storage array](#)
- [SPARC: Sun StorEdge Multipack storage array](#)
- [Sun Cluster software and data services](#)

TABLE P-2 Sun StorEdge 3310 SCSI Storage Array

Title	Part Number
<i>OpenBoot 2.x Command Reference Manual</i>	806-2906
<i>Sun StorEdge PCI Dual Ultra3 SCSI Host Adapter Release Notes</i>	816-2157
<i>Sun StorEdge 3310 SCSI Array Release Notes</i>	816-7292
<i>Sun StorEdge 3000 Family FRU Installation Guide</i>	816-7326
<i>Sun StorEdge 3000 Family Installation, Operation, and Service Manual, Sun StorEdge 3120 SCSI Array</i>	816-7956
<i>Sun StorEdge 3000 Family CLI 2.0 User's Guide</i>	817-4951

TABLE P-3 Sun StorEdge 3120 Storage Array

Title	Part Number
<i>Sun StorEdge 3000 Family FRU Installation Guide</i>	816-7326
<i>Sun StorEdge 3000 Family Installation, Operation, and Service Manual, Sun StorEdge 3120 SCSI Array</i>	816-7956
<i>Sun StorEdge 3120 SCSI Array Release Notes</i>	816-7955
<i>Sun StorEdge 3120 SCSI Array Quick Installation Guide, shipped with your storage array.</i>	816-7962
<i>Sun StorEdge 3000 Family CLI 2.0 User's Guide</i>	817-4951

TABLE P-4 SPARC: Sun StorEdge D1000 Storage Array or Sun StorEdge Netra st D1000 Storage Array

Title	Part Number
<i>Netra st A1000/D1000 Installation and Maintenance Manual</i>	805-7147
<i>Netra st A1000/D1000 Product Notes</i>	805-7148
<i>OpenBoot 2.x Command Reference Manual</i>	806-2906
<i>Sun StorEdge A1000 and D1000 Installation, Operations, and Service Manual</i>	805-2624

TABLE P-5 SPARC: Sun StorEdge Netra D130/S1 Storage Array

Title	Part Number
<i>Netra st D130 Installation and Maintenance Manual</i>	806-1489

TABLE P-6 SPARC: Sun StorEdge D2 Storage Array

Title	Part Number
<i>Sun StorEdge D2 Array Installation, Operation, and Service Manual</i>	816-2578

TABLE P-7 SPARC: Sun StorEdge Multipack Storage Array

Title	Part Number
<i>OpenBoot 2.x Command Reference Manual</i>	806-2906
<i>SPARCstorage MultiPack 2 Service Manual</i>	805-1667

TABLE P-8 Oracle Solaris Cluster and Sun Cluster Documentation

Documentation
Oracle Solaris Cluster 3.3
Sun Cluster 3.2
Sun Cluster 3.1

## Using UNIX Commands

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

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

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

## Getting Help

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

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

- The release number of Oracle Solaris Cluster (for example, Oracle Solaris Cluster 3.3)

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

Command	Function
<code>prtconf -v</code>	Displays the size of the system memory and reports information about peripheral devices
<code>psrinfo -v</code>	Displays information about processors
<code>showrev -p</code>	Reports which patches are installed
<code>prtdiag -v</code>	Displays system diagnostic information
<code>/usr/cluster/bin/clnode show-rev</code>	Displays Oracle Solaris Cluster release and package version information for each node

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

## Documentation, Support, and Training

See the following web sites for additional resources:

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

## Oracle Welcomes Your Comments

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

Oracle Technology Network (<http://www.oracle.com/technetwork/index.html>) 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-9 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% you have mail.</code>
<b>AaBbCc123</b>	What you type, contrasted with onscreen computer output	<code>machine_name% su</code> Password:
<i>aabbcc123</i>	Placeholder: replace with a real name or value	The command to remove a file is <i>rm filename</i> .
<i>AaBbCc123</i>	Book titles, new terms, and terms to be emphasized	Read Chapter 6 in the <i>User's Guide</i> . <i>A cache</i> is a copy that is stored locally. Do <i>not</i> save the file. <b>Note:</b> Some emphasized items appear bold online.

## Shell Prompts in Command Examples

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

TABLE P-10 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 a SCSI JBOD Storage Device

---

This chapter describes the procedures about how to install SCSI JBOD storage devices in an Oracle Solaris Cluster environment.

The procedures in this chapter apply to the following SCSI JBOD storage devices:

- Oracle's Sun StorEdge 3120 storage array
- Oracle's Sun StorEdge 3310 SCSI storage array
- Oracle's Sun StorEdge 3320 SCSI storage array
- SPARC: Oracle's Sun StorEdge D1000 storage array
- SPARC: Oracle's Sun StorEdge D130/S1 storage array
- SPARC: Oracle's Sun StorEdge D2 storage array
- SPARC: Oracle's Sun StorEdge Multipack storage array

## Installing a Storage Array

This section contains procedures for installing storage arrays in new clusters and adding them to existing clusters.

If your storage array uses single-ended SCSI specifications, ensure that your bus lengths comply to the following guidelines.

The single-ended SCSI specifications specify that bus lengths are based on speed and number of devices. The bus lengths in the following table outline a typical implementation of the single-ended SCSI specifications for Oracle hardware. For details, see your host adapter and storage documentation.

TABLE 1-1 Typical Single-Ended, Wide SCSI Bus Lengths

Number of Devices <sup>1</sup>	Maximum Length
<i>UltraSCSI (40MB/s: FAST-20, Wide)</i>	

<sup>1</sup> Devices include both targets and initiators.

TABLE 1-1 Typical Single-Ended, Wide SCSI Bus Lengths (Continued)

Number of Devices <sup>1</sup>	Maximum Length
up to 4	3 meters
up to 8	1.5 meters
<i>FastSCSI (20MB/s: FAST-10, Wide)</i>	
up to 16	6 meters

<sup>1</sup> Devices include both targets and initiators.

If you exceed these specifications, you might experience SCSI errors. The host adapter or the driver might recover from these errors by retrying the request. If this action does not succeed, the host adapter or the driver might recover by renegotiating to a less demanding mode of operation. In some cases, the host adapter or the driver might not be able to recover from these errors, and I/O might fail. You experience delays in I/O if the host adapter or the driver needs to perform this recovery.

If your configuration uses UltraSCSI and requires the 6-meter bus length, use the host adapter driver's `scsi-options` property to limit the speed negotiation to FastSCSI operation. Use the following `/kernel/drv/glm.conf` file as an example to set the `scsi-options` property.

```
name="glm" parent="/pci@1f,4000" unit-address="2" scsi-options=0x3f8;
```

This example uses specific hardware. Change this example to accommodate the hardware in your configuration. In this example, the `scsi-options` property sets the following support.

- Disconnect/reconnect
- Synchronous transfer
- Parity
- Fast SCSI
- Wide SCSI

For more information, see your `isp(7D)`, `glm(7D)` or other host adapter driver man page and documentation on <http://sunsolve.sun.com>.

TABLE 1-2 Task Map: Installing Storage Arrays

Task	Information
Install a storage array in a new cluster, before the OS and Oracle Solaris Cluster software are installed.	<p>“SPARC: How to Install a Storage Array in a New SPARC Based Cluster” on page 13</p> <p>“x86: How to Install a Storage Array in a New x86 Based Cluster” on page 18</p>

TABLE 1-2 Task Map: Installing Storage Arrays (Continued)

Task	Information
Add a storage array to an existing cluster.	<p>“SPARC: How to Add a Storage Array to an Existing SPARC Based Cluster” on page 24</p> <p>“x86: How to Add a Storage Array to an Existing x86 Based Cluster” on page 30</p>

## Setting SCSI Initiator IDs

When your cluster configuration contains shared SCSI JBOD arrays, you must ensure the two nodes connected to a shared SCSI JBOD array have unique SCSI initiator IDs.

The installation procedures in this section describe the steps for setting SCSI initiator IDs in a two-node cluster. If your cluster has additional nodes connected to shared SCSI JBOD arrays, apply these steps as appropriate. Some topologies, for example, clustered pairs, use these procedures unchanged. Others, for example the N+1 topology, may require minor changes.

---

**x86 only** – On x86 based systems, setting SCSI initiator IDs is a two-step process. You first set the IDs in the BIOS and then in a configuration file. Until both steps are complete, the IDs are not set and the systems might be unable to boot or the nodes might panic. Set the IDs on one node at a time, as instructed in the procedure.

---

### ▼ SPARC: How to Install a Storage Array in a New SPARC Based Cluster

This procedure assumes that you are installing one or more storage arrays at initial installation of a SPARC based cluster. If you are adding arrays to a running cluster, use the procedure in “SPARC: How to Add a Storage Array to an Existing SPARC Based Cluster” on page 24.

Multihost storage in clusters uses the multi-initiator capability of the small computer system interface (SCSI) specification. When installing arrays in your cluster, you must ensure that each device in each SCSI chain has a unique SCSI address. The procedure that follows has specific instructions for achieving this requirement. For additional information about multi-initiator capability, see “Multi-Initiator SCSI” in *Oracle Solaris Cluster Concepts Guide*.

---

**Note** – This procedure uses an updated method for setting the `scsi-initiator-id`. The method that was published in earlier documentation is still applicable. However, the method changes if your cluster configuration uses a Sun StorEdge PCI Dual Ultra3 SCSI host adapter to connect to any other shared storage. You then must update your `nvrnrc` script and set the `scsi-initiator-id` by following the steps in this procedure.

---

**Before You Begin** Before performing this procedure, ensure that you have met the following prerequisites. This procedure relies on the following prerequisites and assumptions.

- You have read the entire procedure.
  - You can access necessary patches, drivers, software packages, and hardware.
- 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*.

- You have planned the SCSI address assignments for your host adapters.
  - Your nodes are powered off or are at the OpenBoot PROM.
  - Your arrays are powered off.
  - Your cluster interconnect hardware is connected to the nodes in your cluster.
  - No software is installed.
- 1 Verify that the storage arrays are set up correctly for your planned configuration.**
  - 2 If necessary, install the host adapters in the nodes that you plan to connect to the storage array.** If possible, put each host adapter on a separate I/O board to ensure maximum redundancy.
  - 3 Cable the storage arrays.** For cabling diagrams, see [Chapter 3, “Cabling Diagrams.”](#)

---

**Note** – Ensure that the SCSI bus length does not exceed bus-length specifications. This measurement includes the cables to both nodes, as well as the bus length that is internal to each storage array, node, and the host adapter. For more information about bus-length limitations, see your hardware documentation.

---

**4 Connect the AC or DC power cords for each storage array to a different power source.**

If your storage array has redundant power inputs, connect each power cord from the storage array to a different power source. If the arrays are not mirrors of each other, the arrays can share power sources.

**5 Ensure that each device in the SCSI chain has a unique SCSI address.**

To avoid SCSI-chain conflicts, the following steps instruct you to reserve SCSI address 7 for one host adapter in the SCSI chain and change the other host adapter's global `scsi-initiator-id` to an available SCSI address. Then the steps instruct you to change the `scsi-initiator-id` for local devices back to 7.

---

**Note** – A slot in the storage array might not be in use. However, do not set the `scsi-initiator-id` to a SCSI address for that disk slot. This precaution minimizes future complications if you install additional disk drives.

---

**a. If necessary, power on the storage devices.**

---

**Note** – If necessary, halt the nodes so that you can perform OpenBoot PROM (OBP) Monitor tasks at the ok prompt.

---

For the procedure about powering on a storage device, see the service manual that shipped with your storage device.

**b. If necessary, power on a node, but do not allow it to boot. If necessary, halt the system to continue with OBP Monitor tasks.**

**c. Set the `scsi-initiator-id` for one node to 6.**

```
{1} ok setenv scsi-initiator-id 6
scsi-initiator-id = 6
```

**d. Find the paths to the host adapters that connect to the local disk drives.**

```
{0} ok show-disks
```

Use this information to change the SCSI addresses in the `nvramrc` script. Do not include the `/sd` directories in the device paths.

- e. **Edit the nvramrc script to set the scsi-initiator-id for the local devices on the first node to 7.**

For a full list of commands, see the [OpenBoot 2.x Command Reference Manual](#).



**Caution** – Insert exactly one space after the first double quote and before scsi-initiator-id.

```
{0} ok nvedit
0: probe-all
1: cd /pci@1f,4000/scsi@2
2: 7 encode-int " scsi-initiator-id" property
3: device-end
4: cd /pci@1f,4000/scsi@3
5: 7 encode-int " scsi-initiator-id" property
6: device-end
7: install-console
8: banner[Control C]
{0} ok
```

- f. **Store the changes.**

The changes you make through the nvedit command are recorded on a temporary copy of the nvramrc script. You can continue to edit this copy without risk. After you complete your edits, save the changes. If you are not sure about the changes, discard them.

- **To store the changes, type the following command:**

```
{0} ok nvstore
{1} ok
```

- **To discard the changes, type the following command:**

```
{0} ok nvquit
{1} ok
```

- g. **Verify the contents of the nvramrc script that you created, as shown in the following example.**

If the contents of the nvramrc script are incorrect, use the nvedit command to make corrections.

```
{0} ok printenv nvramrc
nvramrc =
    probe-all
    cd /pci@1f,4000/scsi@2
    7 " scsi-initiator-id" integer-property
    device-end
    cd /pci@1f,4000/scsi@3
    7 " scsi-initiator-id" integer-property
    device-end
    install-console
    banner
{1} ok
```

- h. Instruct the OpenBoot PROM (OBP) Monitor to use the `nvrnrc` script, as shown in the following example.**

```
{0} ok setenv use-nvrnrc? true
use-nvrnrc? = true
{1} ok
```

- 6 Verify that the `scsi-initiator-id` is set correctly on the second node.**

- a. If necessary, power on the second node, but do not allow it to boot. If necessary, halt the system to continue with OBP Monitor tasks.**

- b. Verify that the `scsi-initiator-id` for each host adapter on the second node is set to 7.**

Use the `show-disks` command to find the paths to the host adapters that are connected to these enclosures. Select each host adapter's device tree node, and display the node's properties to confirm that the `scsi-initiator-id` for each host adapter is set to 7.

```
{0} ok cd /pci@6,4000/pci@3/scsi@5
{0} ok .properties
scsi-initiator-id      00000007
...
```

- 7 Install the operating system software.**

- a. Install the Oracle Solaris operating system.**

See your Oracle Solaris Cluster installation documentation for instructions.

- b. Install any unbundled drivers required by your cluster configuration.**

See the host adapter documentation for driver installation procedures .

- c. Apply any required Oracle Solaris patches.**

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

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

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

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

- 8 If you are using Sun StorEdge 3310 JBOD arrays with the Sun StorEdge PCI/PCI-X Dual Ultra320 SCSI host adapter, you must throttle down the speed of the adapter to U160. Add the following entry to your `/kernel/drv/mpt.conf` file:**

```
scsi-options=0x1ff8;
```

- 9 Install the Oracle Solaris Cluster software and volume management software.**

For software installation procedures, see the Oracle Solaris Cluster installation documentation.

- 10 If you are using Solaris Volume Manager as your volume manager, save the disk-partitioning information.**



**Caution** – Do not save disk-partitioning information in `/tmp` because you will lose this file when you reboot. Instead, save this file in `/usr/tmp`.

---

You might need disk-partitioning information if you replace a failed disk drive in the future.

## ▼ **x86: How to Install a Storage Array in a New x86 Based Cluster**

This procedure assumes that you are installing one or more storage arrays at initial installation of an x86 based cluster. If you are adding arrays to a running cluster, use the procedure in “[SPARC: How to Add a Storage Array to an Existing SPARC Based Cluster](#)” on page 24.

Multi-host storage in clusters uses the multi-initiator capability of the small computer system interface (SCSI) specification. When installing arrays in your cluster, you must ensure that each device in each SCSI chain has a unique SCSI address. The procedure that follows has specific instructions for achieving this requirement. For additional information about multi-initiator capability, see “[Multi-Initiator SCSI](#)” in *Oracle Solaris Cluster Concepts Guide*.

---

**Note** – On x86 based systems, setting SCSI initiator IDs is a two-step process. You first set the IDs in the BIOS and then in a configuration file. Until both steps are complete, the IDs are not set and the systems might not boot or the nodes might panic. Set the IDs on one node at a time, as instructed in the procedure.

---

**Before You Begin** Before performing this procedure, ensure that you have met the following prerequisites. This procedure relies on the following prerequisites and assumptions.

- You have read the entire procedure.
- You can access necessary patches, drivers, software packages, and hardware.

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*.

- You have planned the SCSI address assignments for your host adapters.
- Your nodes and arrays are powered off.
- Your cluster interconnect hardware is connected to the nodes in your cluster.
- No software is installed.

- 1 Verify that the storage arrays are set up correctly for your planned configuration.**
- 2 If necessary, install the host adapters in the nodes that you plan to connect to the storage array.** If possible, put each host adapter on a separate BUS to ensure maximum redundancy.
- 3 Power on one node.**
- 4 On the first node, ensure that each device in the SCSI chain has a unique SCSI address by configuring the initiator IDs in the BIOS.**

To avoid SCSI-chain conflicts, perform the following steps.

---

**Note** – Perform these steps on only one cluster node.

---

**a. Access your BIOS settings.**

To access the BIOS on the V40z server with SG-XCPI2SCSI-LM320 Sun StorEdge PCI/PCI-X Dual Ultra320 SCSI Host Adapter cards, press **Ctrl-C** when prompted during reboot.

**b. Verify that the internal controller is set to the default value of 7.**

**c. Set the new host adapter `scsi-initiator-id` to 6.**

**5 Cable the storage arrays to all nodes.**

For cabling diagrams, see [Chapter 3, “Cabling Diagrams.”](#)

---

**Note** – Ensure that the bus length does not exceed SCSI-bus-length specifications. This measurement includes the cables to both nodes, as well as the bus length that is internal to each storage array, node, and the host adapter. For more information about SCSI-bus-length limitations, see your hardware documentation.

---

**6 Connect the AC or DC power cords for each storage array to a different power source.**

If your storage array has redundant power inputs, connect each power cord from the storage array to a different power source. If the arrays are not mirrors of each other, the arrays can share power sources.

**7 Power on the storage array.**

For the procedure about powering on a storage device, see the service manual that shipped with your storage device.

**8 Install the operating system software on the node for which you configured the BIOS in [Step 4](#).**

**a. Install the Solaris operating system.**

See your Oracle Solaris Cluster installation documentation for instructions.

**b. Install any unbundled drivers required by your cluster configuration.**

For driver installation procedures, see the host adapter documentation.

**c. Apply any required Solaris patches.**

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 On the node for which you configured the BIOS in Step 4, finish configuring the SCSI initiator IDs.**

**a. Get the information required for the `mpt.conf` file.**

To create the `mpt.conf` entries, you need the path to your boot disk and the SCSI unit address.

To find this information on X4000 series servers with SG-XCPI2SCSI-LM320 Sun StorEdge PCI/PCI-X Dual Ultra320 SCSI Host Adapter cards, use the following command:

```
# echo | format
Searching for disks...done

AVAILABLE DISK SELECTIONS:
    0. clt0d0 <DEFAULT cyl 8938 alt 2 hd 255 sec 63>
       /pci@0,0/pci1022,7450@a/pci17c2,10@4/sd@0,0
```

**b. Create a `/kernel/drv/mpt.conf` file.**

**c. Include the following entries:**

```
scsi-initiator-id=6;
name="mpt" parent="/pci@0,0/pci1022,7450@a"
  unit-address="4"
  scsi-initiator-id=7;
```

---

**Note** – These entries are based on the foregoing example output of the `format` command. Your entries must include the values output from your `format` command. Also, note that the `parent` and `unit-address` values are strings. The quotation marks are required to form correct values in the `mpt.conf` file.

---

The entries in this example have the following meanings:

`scsi-initiator-id=6;`      Matches your setting in the BIOS for the host adapter ports.

<code>name="mpt";</code>	Indicates that these settings are for the mpt driver.
<code>parent</code>	Specifies the path to your local drive which you discovered in <a href="#">Step a</a> .
<code>unit-address</code>	Specifies the unit address of the local drive. In the example in <a href="#">Step a</a> , this information derives from the <code>pci17c2,10@4</code> portion of the output.
<code>scsi-initiator-id=7;</code>	Sets your node's local drive back to the default SCSI setting of 7.

**d. Reboot the node to activate the `mpt.conf` file changes.**

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

**10 Ensure that each LUN has an associated entry in the `/kernel/drv/sd.conf` file.**

For more information, see the *Sun StorEdge 3000 Family Installation, Operation, and Service Manual*.

**11 To make the changes to the `/kernel/drv/sd.conf` file active, perform the following option.**

Run the `update_drv -f sd` command and then the `devfsadm` command.

**12 Power on the remaining nodes and install the operating system software on them.**

**a. Install the Oracle Solaris operating system.**

See your Oracle Solaris Cluster installation documentation for instructions.

**b. Install any unbundled drivers required by your cluster configuration.**

For driver installation procedures, see the host adapter documentation.

**c. Apply any required Oracle Solaris patches.**

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

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

If you must apply a patch when a node is in noncluster mode, you can apply it in a rolling fashion, one node at a time, unless instructions for a patch require that you shut down the entire cluster. Follow the procedures in [“How to Apply a Rebooting Patch \(Node\)”](#) in *Oracle*

*Solaris Cluster System Administration Guide* to prepare the node and to boot it in noncluster mode. For ease of installation, consider applying all patches at the same time. That is, apply all patches to the node that you place in noncluster mode.

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

- 13 Ensure that each LUN has an associated entry in the `/kernel/drv/sd.conf` file.**  
For more information, see the *Sun StorEdge 3000 Family Installation, Operation, and Service Manual*.
- 14 To make the changes to the `/kernel/drv/sd.conf` file active, perform the following option.**  
Run the `update_drv -f sd` command and then the `devfsadm` command.
- 15 If you are using Sun StorEdge 3310 JBOD arrays with the Sun StorEdge PCI/PCI-X Dual Ultra320 SCSI host adapter, you must throttle down the speed of the adapter to U160. Add the following entry to your `/kernel/drv/mpt.conf` file on each node:**  
`scsi-options=0x1ff8;`
- 16 Ensure that each LUN has an associated entry in the `/kernel/drv/sd.conf` file.**  
For more information, see the *Sun StorEdge 3000 Family Installation, Operation, and Service Manual*.
- 17 To make the changes to the `/kernel/drv/sd.conf` file active, perform the following option.**  
Run the `update_drv -f sd` command and then the `devfsadm` command.
- 18 Install the Oracle Solaris Cluster software and volume management software on each node.**  
For software installation procedures, see the Oracle Solaris Cluster installation documentation.

#### **Example 1–1** x86: A Completed `mpt.conf` File When Using StorEdge 3320 Arrays

The following `mpt.conf` file shows all entries, assuming the following:

- The output of the `format` command is that shown in [Step a](#).
- You are using Sun StorEdge 3320 JBOD arrays with the Sun StorEdge PCI/PCI-X Dual Ultra320 SCSI host adapter.

```
# more /kernel/drv/mpt.conf
scsi-initiator-id=6;
name="mpt" parent="/pci@0,0/pci1022,7450@a"
    unit-address="4"
    scsi-initiator-id=7;
```

**Example 1-2** x86: A Completed `mpt.conf` File When Using StorEdge 3310 Arrays

The following `mpt.conf` file shows all entries, assuming the following:

- The output of the `format` command is that shown in [Step a](#).
- You are using Sun StorEdge 3310 JBOD arrays with the Sun StorEdge PCI/PCI-X Dual Ultra320 SCSI host adapter.

```
# more /kernel/drv/mpt.conf
scsi-initiator-id=6;
name="mpt" parent="/pci@0,0/pci1022,7450@a"
    unit-address="4"
    scsi-initiator-id=7;
scsi-options=0x1ff8;
```

**Next Steps** If needed, finish setting up your storage arrays, including partitions. If you are using Solaris Volume Manager as your volume manager, save the disk-partitioning information. You might need disk-partitioning information if you replace a failed disk drive in the future.




---

**Caution** – Do not save disk-partitioning information in `/tmp` because you will lose this file when you reboot. Instead, save this file in `/usr/tmp`.

---

## ▼ SPARC: How to Add a Storage Array to an Existing SPARC Based Cluster

This procedure contains instructions for adding storage arrays to an operational cluster. If you need to install storage arrays to a new cluster, use the procedure in “[SPARC: How to Install a Storage Array in a New SPARC Based Cluster](#)” on page 13 or “[x86: How to Install a Storage Array in a New x86 Based Cluster](#)” on page 18.

Adding a storage array enables you to alter your storage pool. You might want to perform this procedure in the following scenarios.

- You need to increase your storage pool.
- You need to upgrade to a higher-quality or to a larger storage array.
  - To upgrade storage arrays, remove the old storage array and then add the new storage array.
  - To replace a storage array with the same type of storage array, see “[How to Replace the Chassis](#)” on page 41.

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

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

- 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, see the Oracle Solaris Cluster system administration documentation, and skip steps that instruct you to shut down the node.

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 Install all software that is specific to the storage array or to any new host adapters.**

Install the software and patches to all nodes that will connect to the new storage array.

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*.

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

```
# clnode evacuate NodeA
```

**3 If you need to install host adapters in Node A, perform the following steps.**

**a. Shut down and power off Node A.**

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

**b. Install host adapters in Node A.**

For the procedure about how to install host adapters, see your host adapters and server documentation.

**4 Connect the storage array to the host adapters on Node A.**

- If necessary, terminate the ports that will connect to Node B.
  - If you have a Netra D130 array, always terminate the ports that connect to Node B.
  - If you have a StorEdge 3310 or 3320 SCSI array, terminate the ports that connect to Node B when using a split-bus configuration.
- If your storage array is a StorEdge 3310 or 3320 SCSI array, do not power on the storage array until the storage array is cabled to Node A.

For cabling diagrams, see [Chapter 3, “Cabling Diagrams.”](#)

---

**Note** – Ensure that the bus length does not exceed SCSI bus-length specifications. This measurement includes the cables to both nodes, as well as the bus length that is internal to each storage array, node, and the host adapter. For more information about SCSI bus-length limitations, see your hardware documentation.

---

**5 Ensure that each device in the SCSI chain has a unique SCSI address.**

To avoid SCSI-chain conflicts, the following steps instruct you to reserve SCSI address 7 for one host adapter in the SCSI chain and change the other host adapter's `scsi-initiator-id` to an available SCSI address. Then the steps instruct you to change the `scsi-initiator-id` for local devices back to 7.

---

**Note** – A slot in the storage array might not be in use. However, do not set the `scsi-initiator-id` to a SCSI address for that disk slot. This precaution minimizes future complications if you install additional disk drives.

---

**a. If necessary, power on the storage devices.**

For the procedure about powering on a storage device, see the service manual that shipped with your storage device.

---

**Note** – If necessary, halt the nodes so that you can perform OpenBoot PROM (OBP) Monitor tasks at the `ok` prompt.

---

**b. If necessary, power on Node A, but do not allow it to boot. If necessary, halt the system to continue with OBP Monitor tasks.**

**c. Set the `scsi-initiator-id` for Node A to 6.**

```
{1} ok setenv scsi-initiator-id 6
scsi-initiator-id = 6
```

**d. Find the paths to the host adapters that connect to the local disk drives.**

```
{0} ok show-disks
```

Use this information to change the SCSI addresses in the `nvrnrc` script. Do not include the `/sd` directories in the device paths.

**e. Edit the `nvrnrc` script to set the `scsi-initiator-id` for the local devices on the first node to 7.**

For a full list of commands, see the [OpenBoot 2.x Command Reference Manual](#).




---

**Caution** – Insert exactly one space after the first double quote and before `scsi-initiator-id`.

---

```
{0} ok nvedit
0: probe-all
1: cd /pci@1f,4000/scsi@2
2: 7 encode-int " scsi-initiator-id" property
3: device-end
4: cd /pci@1f,4000/scsi@3
5: 7 encode-int " scsi-initiator-id" property
6: device-end
7: install-console
8: banner[Control C]
{0} ok
```

**f. Store the changes.**

The changes you make through the `nvedit` command are recorded on a temporary copy of the `nvrnrc` script. You can continue to edit this copy without risk. After you complete your edits, save the changes. If you are not sure about the changes, discard them.

- **To store the changes, type the following command:**

```
{0} ok nvstore
{1} ok
```

- **To discard the changes, type the following command:**

```
{0} ok nvquit
{1} ok
```

**g. Verify the contents of the `nvrarc` script that you created, as shown in the following example.**

If the contents of the `nvrarc` script are incorrect, use the `nvedit` command to make corrections.

```
{0} ok printenv nvrarc
nvrarc =                probe-all
                        cd /pci@1f,4000/scsi@2
                        7 "scsi-initiator-id" integer-property
                        device-end
                        cd /pci@1f,4000/scsi@3
                        7 "scsi-initiator-id" integer-property
                        device-end
                        install-console
                        banner

{1} ok
```

**h. Instruct the OpenBoot PROM (OBP) Monitor to use the `nvrarc` script, as shown in the following example.**

```
{0} ok setenv use-nvrarc? true
use-nvrarc? = true
{1} ok
```

- 6 To create the new Solaris device files and links, perform a reconfiguration boot on Node A by adding `-r` to your boot instruction.**
- 7 If necessary, format and label the disks.**
- 8 On Node A, verify that the device IDs have been assigned to the disk drives in the storage array.**  

```
# cldevice list -n NodeA -v
```
- 9 Move all resource groups and device groups off Node B.**  

```
# clnode evacuate NodeB
```
- 10 If you need to install host adapters in Node B, perform the following steps.**
  - a. Shut down Node B.**  
 For the procedure about how to shut down and power off a node, see the Oracle Solaris Cluster system administration documentation.
  - b. Install the host adapters in Node B.**  
 For the procedure about how to install a host adapter, see your host adapter and server documentation.

**c. Power on and boot Node B.**

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

**11 Connect the storage array to the host adapters on Node B.**

If you added port terminators in [Step 4](#), remove the terminator ports and connect the storage array to Node B.

For cabling diagrams, see [Chapter 3, “Cabling Diagrams.”](#)

---

**Note** – Ensure that the bus length does not exceed bus-length specifications. This measurement includes the cables to both nodes, as well as the bus length that is internal to each storage array, node, and the host adapter. For more information about bus-length limitations, see your hardware documentation.

---

**12 Verify that the `scsi-initiator-id` is set correctly on the second node.**

**a. If necessary, power on the second node, but do not allow it to boot. If necessary, halt the system to continue with OBP Monitor tasks.**

**b. Verify that the `scsi-initiator-id` for each host adapter on the second node is set to 7.**

Use the `show-disks` command to find the paths to the host adapters that are connected to these enclosures. Select each host adapter's device tree node, and display the node's properties to confirm that the `scsi-initiator-id` for each host adapter is set to 7.

```
{0} ok cd /pci@6,4000/pci@3/scsi@5
{0} ok .properties
scsi-initiator-id      00000007
...
```

**13 To create the new Oracle Solaris device files and links, perform a reconfiguration boot by adding `-r` to your boot instruction.**

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

**14 On Node B, verify that the device IDs have been assigned to the disk drives in the storage array.**

```
# cldevice list -n NodeB -v
```

**15 Perform volume management administration to add the disk drives in the storage array to the volume management configuration.**

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

**Next Steps** If needed, finish setting up your storage arrays, including partitions. If you are using Solaris Volume Manager as your volume manager, save the disk-partitioning information. You might need disk-partitioning information if you replace a failed disk drive in the future.



**Caution** – Do not save disk-partitioning information in /tmp because you will lose this file when you reboot. Instead, save this file in /usr/tmp.

## ▼ **x86: How to Add a Storage Array to an Existing x86 Based Cluster**

This procedure contains instructions for adding storage arrays to an operational cluster. If you need to install storage arrays to a new cluster, use the procedure in “[SPARC: How to Install a Storage Array in a New SPARC Based Cluster](#)” on page 13 or “[x86: How to Install a Storage Array in a New x86 Based Cluster](#)” on page 18.

Adding a storage array enables you to alter your storage pool. You might want to perform this procedure in the following scenarios.

- You need to increase your storage pool.
- You need to upgrade to a higher-quality or to a larger storage array.  
To upgrade storage arrays, remove the old storage array and then add the new storage array.  
To replace a storage array with the same type of storage array, see “[How to Replace the Chassis](#)” on page 41.

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

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

- 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, see the Oracle Solaris Cluster system administration documentation, and skip steps that instruct you to shut down the node.

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

- 1 Install all software that is specific to the storage array or to any new host adapters.**  
Install the software and patches to all nodes that will connect to the new storage array.

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*.

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

```
# clnode evacuate NodeA
```

**3 If you need to install host adapters in Node A, perform the following steps.**

**a. Shut down and power off Node A.**

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

**b. Install host adapters in Node A.**

For the procedure about how to install host adapters, see your host adapters and server documentation.

**4 If you installed host adapters in Step 3, or if you intend to use previously unconfigured host adapters, ensure that each device in the SCSI chain has a unique SCSI address by configuring the `scsi-initiator-id` in the BIOS.**

---

**Note** – Perform these steps on one cluster node, the node on which you have configured SCSI initiator IDs for the cluster in the past.

---

**a. Access your host adaptor's BIOS settings.**

To access the BIOS on the V40z server with X4422A Sun Dual Gigabit Ethernet and Dual SCSI Adapter cards, press **Ctrl-C** when prompted during reboot.

**b. Verify that internal controller is set to the default value of 7.**

- c. Select a unique value for each of the new host adapter's ports.
  - d. Set each controller's `scsi-initiator-id` to that value.
- 5 Connect the storage array to the host adapters on Node A.
- If necessary, terminate the ports that will connect to Node B.
    - If you have a Netra D130 array, always terminate the ports that connect to Node B.
    - If you have a StorEdge 3310 or 3320 SCSI array, terminate the ports that connect to Node B when using a split-bus configuration.
  - If your storage array is a StorEdge 3310 or 3320 SCSI array, do not power on the storage array until the storage array is cabled to Node A.

For cabling diagrams, see [Chapter 3, “Cabling Diagrams.”](#)

---

**Note** – Ensure that the bus length does not exceed SCSI bus-length specifications. This measurement includes the cables to both nodes, as well as the bus length that is internal to each storage array, node, and the host adapter. For more information about SCSI bus-length limitations, see your hardware documentation.

---

- 6 If you installed host adapters in [Step 3](#), or if you intend to use previously unconfigured host adapters, finish configuring the SCSI initiator IDs on the same node on which you configured the BIOS in [Step 4](#).

- a. Get the information required for the `mpt.conf` file.

To create the `mpt.conf` entries, you need the path to your boot disk and the SCSI unit address.

To find this information on X4000 series servers with SG-XCPI2SCSI-LM320 Sun StorEdge PCI/PCI-X Dual Ultra320 SCSI Host Adapter cards, use the following command:

```
# echo | format
Searching for disks...done

AVAILABLE DISK SELECTIONS:
  0. clt0d0 <DEFAULT cyl 8938 alt 2 hd 255 sec 63>
     /pci@0,0/pci1022,7450@a/pci17c2,10@4/sd@0,0
```

- b. Create or edit your `/kernel/drv/mpt.conf` file.

- c. Include the following entries:

```
scsi-initiator-id=6;
name="mpt" parent="/pci@0,0/pci1022,7450@a"
  unit-address="4"
scsi-initiator-id=7;
```

---

**Note** – These entries are based on the foregoing example output of the `format` command. Your entries must include the values output from your `format` command. Also, note that the `parent` and `unit-address` values are strings. The quotation marks are required to form correct values in the `mpt.conf` file.

---

The entries in this example have the following meanings:

<code>scsi-initiator-id=6;</code>	Matches your setting in the BIOS for the host adapter ports.
<code>name="mpt"</code>	Indicates that these settings are for the <code>mpt</code> driver.
<code>parent</code>	Is set to the path to your local drive.
<code>unit-address</code>	Specifies the unit address of the local drive. In the example in <a href="#">Step a</a> , this information derives from the <code>pci17c2,10@4</code> portion of the output.
<code>scsi-initiator-id=7;</code>	Sets your node's local drive back to the default SCSI setting of 7.

**7 To create the new Solaris device files and links, perform a reconfiguration boot on Node A by adding `-r` to your boot instruction.**

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

**8 If necessary, format and label the disks.**

**9 On Node A, verify that the device IDs have been assigned to the disk drives in the storage array.**

```
# cldevice show -v
```

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

```
# clnode evacuate NodeB
```

**11 If you need to install host adapters in Node B, perform the following steps.**

**a. Shut down Node B.**

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

**b. Install the host adapters in Node B.**

For the procedure about how to install a host adapter, see your host adapter and server documentation.

**c. Power on and boot Node B.**

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

**12 Connect the storage array to the host adapters on Node B.**

If you added port terminators in [Step 5](#), remove the terminator ports and connect the storage array to Node B.

For cabling diagrams, see [Chapter 3, “Cabling Diagrams.”](#)

---

**Note** – Ensure that the bus length does not exceed bus-length specifications. This measurement includes the cables to both nodes, as well as the bus length that is internal to each storage array, node, and the host adapter. For more information about bus-length limitations, see your hardware documentation.

---

**13 Verify that the `scsi-initiator-id` is set correctly on the second node.**

**a. Access your BIOS settings.**

To access the BIOS on X4000 series servers with SG-XCPI2SCSI-LM320 Sun StorEdge PCI/PCI-X Dual Ultra320 SCSI Host Adapter cards, press **Ctrl-C** when prompted during reboot.

**b. Verify that the `scsi-initiator-id` for each host adapter on the second node is set to 7.**

**14 To create the new Solaris device files and links, perform a reconfiguration boot by adding `-r` to your boot instruction.**

**15 On Node B, verify that the device IDs have been assigned to the disk drives in the storage array.**

```
# cldevice show -v
```

**16 Perform volume management administration to add the disk drives in the storage array to the volume management configuration.**

For more information, see your Solaris Volume Manager documentation.

**Next Steps** If needed, finish setting up your storage arrays, including partitions. If you are using Solaris Volume Manager as your volume manager, save the disk-partitioning information. You might need disk-partitioning information if you replace a failed disk drive in the future.




---

**Caution** – Do not save disk-partitioning information in `/tmp` because you will lose this file when you reboot. Instead, save this file in `/usr/tmp`.

---

## Maintaining a SCSI JBOD Storage Device

---

This chapter describes the procedures about how to maintain SCSI JBOD storage devices in an Oracle Solaris Cluster environment.

The procedures in this chapter apply to the following SCSI JBOD storage devices:

- Sun StorEdge 3120 storage array
- Sun StorEdge 3310 SCSI storage array
- Sun StorEdge 3320 SCSI storage array
- SPARC: Sun StorEdge D1000 storage array
- SPARC: Sun StorEdge Multipack storage array
- SPARC: Sun StorEdge D130/S1 storage array
- SPARC: Sun StorEdge D2 storage array

This chapter contains the following procedures:

- [“How to Remove a Storage Array” on page 39](#)
- [“How to Replace the Chassis” on page 41](#)
- [“How to Replace a SCSI Cable” on page 47](#)
- [“Disconnecting and Reconnecting a Node from Shared Storage” on page 43](#)
- [“How to Replace a Host Adapter When Using Failover and Scalable Data Services Only” on page 49](#)
- [“How to Replace a Host Adapter When Using Oracle Real Application Clusters Only” on page 52](#)
- [“How to Replace a Host Adapter When Using Both Failover and Scalable Data Services and Oracle Real Application Clusters” on page 55](#)
- [“How to Add a Disk Drive” on page 62](#)
- [“How to Remove a Disk Drive” on page 63](#)
- [“How to Replace a Disk Drive \*Without\* Oracle Real Application Clusters” on page 65](#)
- [“SPARC: How to Replace a Disk Drive \*With\* Oracle Real Application Clusters” on page 68](#)
- [“How to Upgrade Disk Drive Firmware” on page 71](#)
- [“How to Upgrade Host Adapter Firmware” on page 72](#)

## Maintaining Storage Arrays

The maintenance procedures in “[FRUs That Do Not Require Oracle Solaris Cluster Maintenance Procedures](#)” on page 37 are performed the same as in a noncluster environment. [Table 2–1](#) lists the procedures that require cluster-specific steps.

TABLE 2–1 Task Map: Maintaining a Storage Array

Task	Information
Remove a storage array	<a href="#">“How to Remove a Storage Array” on page 39</a>
Replace a storage array	<a href="#">“SPARC: How to Add a Storage Array to an Existing SPARC Based Cluster” on page 24</a>
To replace a storage array, remove the storage array. Add a new storage array to the configuration.	<a href="#">“How to Remove a Storage Array” on page 39</a>
Add a JBOD as an expansion unit.	<i>Sun StorEdge 3000 Family Installation, Operation, and Service Manual</i>
Follow the same procedure that you use in a noncluster environment.	
Remove a JBOD expansion unit.	<i>Sun StorEdge 3000 Family Installation, Operation, and Service Manual</i>
Follow the same procedure that you use in a noncluster environment.	
Replace a chassis	<a href="#">“How to Replace the Chassis” on page 41</a>
Replace a SCSI cable	<a href="#">“How to Replace a SCSI Cable” on page 47</a>
Replace a host adapter	<a href="#">“Replacing a Host Adapter” on page 48</a>
Add a disk drive	<a href="#">“How to Add a Disk Drive” on page 62</a>
Remove a disk drive	<a href="#">“How to Remove a Disk Drive” on page 63</a>
Replace a disk drive	<a href="#">“How to Replace a Disk Drive <i>Without</i> Oracle Real Application Clusters” on page 65</a> <a href="#">“SPARC: How to Replace a Disk Drive <i>With</i> Oracle Real Application Clusters” on page 68</a>
Upgrade disk drive firmware	<a href="#">“How to Upgrade Disk Drive Firmware” on page 71</a>
Upgrade host adapter firmware	<a href="#">“How to Upgrade Host Adapter Firmware” on page 72</a>

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

Each storage device has a different set of FRUs that do not require cluster-specific procedures. Choose among the following storage devices:

- Sun StorEdge 3120
- Sun StorEdge 3310 SCSI
- (SPARC) Sun StorEdge D1000
- (SPARC) Sun StorEdge D130/S1
- (SPARC) Sun StorEdge D2
- (SPARC) Sun StorEdge Multipack

### Sun StorEdge 3120 Storage Array FRUs

The following is a list of administrative tasks that require no cluster-specific procedures. See the *Sun StorEdge 3000 Family FRU Installation Guide* for the procedures for the following FRUs. For a URL to this storage documentation, see [“Related Documentation” on page 6](#).

- Air management blank sled
- Fan module
- Power supply

### Sun StorEdge 3310 and 3320 SCSI Storage Array FRUs

The following is a list of administrative tasks that require no cluster-specific procedures. See the *Sun StorEdge 3000 Family FRU Installation Guide* for instructions on replacing the following FRUs. <http://www.tokensbeads.com/education.htm>

- Replace a disk drive
- Replace an AC or DC power supply and fan module
- Replace a battery
- Replace an LED and reset switch module
- Replace an EMU module
- Replace an air management blank sled
- Replace a fan module
- Replace a power supply

## SPARC: Sun StorEdge D1000 Storage Array FRUs

The following is a list of administrative tasks that require no cluster-specific procedures. See the *Sun StorEdge A1000 and D1000 Installation, Operations, and Service Manual* and the *Sun StorEdge A1000 and D1000 Installation, Operations, and Service Manual* for the procedures for the following FRUs. For a URL to this storage documentation, see [“Related Documentation” on page 6](#).

- Controller interface board
- Cooling canister
- Input power cable
- Power supply

## SPARC: Sun StorEdge Netra D130/S1 Storage Array FRUs

The following is a list of administrative tasks that require no cluster-specific procedures. See the *Netra st D130 Installation and Maintenance Manual* for the procedures for the following FRUs. For a URL to this storage documentation, see [“Related Documentation” on page 6](#).

- Input power cable
- Internal fan (cooling system)
- Power supply

## SPARC: Sun StorEdge D2 Storage Array FRUs

The following is a list of administrative tasks that require no cluster-specific procedures. See the *Sun StorEdge D2 Array Installation, Operation, and Service Manual* for the procedures for the following FRUs. For a URL to this storage documentation, see [“Related Documentation” on page 6](#).

- Environmental Services Module (ESM)
- Fan
- Fan canister
- Power supply

## SPARC: Sun StorEdge Multipack Storage Array FRUs

The following is a list of administrative tasks that require no cluster-specific procedures. See the *SPARCstorage MultiPack Service Manual* for the procedures for the following FRUs. For a URL to this storage documentation, see [“Related Documentation” on page 6](#).

- Backplane
- Power supply

## ▼ How to Remove a Storage Array

Removing a storage array enables you to downsize or reallocate your existing storage pool.

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

- You plan to remove the references to the disk drives in the array.
- You do not need to *replace* the storage array's chassis.
  - If you need to replace your storage array's chassis, see [“How to Replace the Chassis” on page 41](#).
- Your nodes are not configured with dynamic reconfiguration functionality.
  - If your nodes are configured for dynamic reconfiguration, see Oracle Solaris Cluster system administration documentation, and skip steps that instruct you to shut down the node.

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 the storage array that you want to remove contains a quorum device, add a new quorum device that will not be affected by this procedure. Then remove the old quorum device.**

To determine whether the affected array contains a quorum device, use the following command.

```
# clquorum show
```

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

- 2 If necessary, back up the metadvice or volume.**

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

- 3 On each node that is connected to the storage array, perform volume management administration to remove the storage array from the configuration.**

If a volume manager does manage the disk drives, run the appropriate volume manager commands to remove the disk drives from any diskset or disk group. For more information, see your Solaris Volume Manager or Veritas Volume Manager documentation. See the following paragraph for additional Veritas Volume Manager commands that are required.

---

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

---

```
# vxdisk offline cNtXdY
# vxdisk rm cNtXdY
```

**4 Identify the disk drives that you plan to remove.**

```
# cfgadm -al
```

**5 On all nodes, remove references to the disk drives in the storage array that you plan to remove.**

```
# cfgadm -c unconfigure cN::dsk/cNtXdY
```

**6 Disconnect the SCSI cables from the storage array.**

**7 On all nodes, update device namespaces.**

```
# devfsadm -C
```

**8 On all nodes, remove all obsolete device IDs.**

```
# cldevice clear
```

**9 Power off the storage array. Disconnect the storage array from the power source.**

For the procedure about how to power off a storage array, see your storage documentation. For a list of storage documentation, see [“Related Documentation” on page 6](#).

**10 Remove the storage array.**

For the procedure about how to remove a storage array, see your storage documentation. For a list of storage documentation, see [“Related Documentation” on page 6](#).

**11 If you plan to remove a host adapter that has an entry in the `nvrarc` script, delete the references to the host adapters in the `nvrarc` script.**

---

**Note** – If other parallel SCSI devices are connected to the nodes, you can delete the contents of the `nvrarc` script. Then, at the OpenBoot PROM, set `setenv use-nvrarc? false`. Afterward, reset the `scsi-initiator-id` to 7 as outlined in [“Installing a Storage Array” on page 11](#).

---

**12 If necessary, remove any unused host adapters from the nodes.**

For the procedure about how to remove a host adapter, see your host adapter and server documentation.

**13 From any node, verify that the configuration is correct.**

```
# cldevice list -v
```

## ▼ How to Replace the Chassis

You might need to replace a chassis if the chassis fails. With this procedure, you are able to replace the chassis and retain the storage array's disk drives and the references to those disk drives. By retaining the storage array's disk drives, you save time because you no longer need to resynchronize your mirrors or restore your data.

If you need to replace the entire storage array, see [“How to Remove a Storage Array” on page 39](#).

**Before You Begin** This procedure relies on the following assumptions.

- You plan to retain the disk drives in the storage array that you need to replace.
- You plan to retain the references to the disk drives.
- Your replacement chassis is unpacked and ready to be placed into the cluster.
- Your nodes are *not* configured with dynamic reconfiguration functionality.

If your nodes are configured for dynamic reconfiguration, see the Oracle Solaris Cluster system administration documentation, and skip steps that instruct you to shut down the node.

- You are *not* using a StorEdge D2, a StorEdge 3310 SCSI, or a StorEdge 3320 SCSI storage array in split-bus mode.

If you are running your storage array in split-bus mode, you must shut down the entire cluster to replace the chassis of the storage array.

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

**1 If the chassis you are replacing contains a drive that is configured as quorum device, add a new quorum device that will not be affected by this procedure. Then remove the old quorum device.**

To determine whether a quorum device will be affected by this procedure, use one of the following commands.

```
# clquorum show
```

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

**2 If possible, back up the metadevices or volumes that reside in the storage array.**

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

- 3 On each node that is connected to the storage array, perform volume management administration to remove the storage array from the configuration.**

---

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

---

If a volume manager manages the disk drives, run the appropriate volume manager commands to remove the disk drives from any diskset or disk group. For more information, see your Solaris Volume Manager or Veritas Volume Manager documentation.

```
# vxdisk offline cNtXdY# vxdisk rm cNtXdY
```

- 4 Disconnect the SCSI cables from the storage array.**

You can disconnect the cables in any order.

- 5 Power off and disconnect the storage array from the power source.**

For more information, see your storage documentation. For a list of storage documentation, see [“Related Documentation” on page 6](#).

- 6 Connect the new storage array to the power sources.**

- 7 Connect the SCSI cables to the new storage array.**

You can connect the cables in any order.

Ensure the cable does not exceed bus-length limitations. For more information on bus-length limitations, see your hardware documentation.

- 8 One disk drive at a time, remove the disk drives from the old storage array. Insert the disk drives into the same slots in the new storage array.**

Move your other components to the new chassis as well. For the procedures about how to replace your storage array's components, see your storage documentation. For a list of storage documentation, see [“Related Documentation” on page 6](#).

- 9 Power on the storage array.**

For the procedure about how to shut down and power off the storage array, see the documentation that came with your array.

- 10 On each node that is attached to the storage array, run the `devfsadm (1M)` command.**

```
# devfsadm
```

- 11 **From one node, attach the new storage array to the global device namespace.**

```
# cldevice populate
```

- 12 **One at a time, shut down and reboot the nodes that are connected to the storage array.**

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

- 13 **Perform volume management administration to add the storage array back into the configuration.**

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

## Disconnecting and Reconnecting a Node from Shared Storage

Use this procedure to *temporarily* disconnect a node from shared storage. You need to temporarily disconnect a node from shared storage if you intend to replace an HBA.

This procedure relies on the following assumptions.

- You intend to reconnect the shared storage to the same host adapter that the storage array was connected before you disconnected the SCSI cable.
- This procedure defines Node A as the node you want to disconnect from the shared storage. Node B is the remaining node.
- This procedure assumes that you might have more than one storage array connected to Node A.

These procedures provide 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.

### ▼ How to Disconnect the Node from Shared Storage

You must maintain proper SCSI-bus termination during this procedure. The process by which you disconnect the node from shared storage depends on whether you have host adapters available on Node B (see [Figure 2–1](#)). If you do not have host adapters available on Node B and your storage device does not have auto-termination, you must use terminators (see [Figure 2–2](#)).

---

**Note** – To determine the specific terminator that your storage array supports, see your storage documentation. For a list of storage documentation, see [“Related Documentation”](#) on page 6.

---

- 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 “[How to Reconnect the Node to Shared Storage](#)” on page 46 to return resource groups and device groups to this node.

Use the following command:

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

**3 Identify the submirrors on the storage array that is connected to Node A.****4 Move all resource groups and device groups off Node A.**

```
# clnode evacuate NodeA
```

**5 (Optional) If necessary, detach the submirrors on the storage array that is connected to Node A.**

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


**6 Shut down Node A.**

To shut down and power off a node, see your Oracle Solaris Cluster system administration documentation.

**7 To maintain proper SCSI-bus termination during this procedure, perform one of the following steps. The approach you choose depends on whether you have an available host adapter on Node B.**


- **Disconnect the SCSI cable between Node A and the storage array. Attach this SCSI cable to Node B on the storage array. For an illustration, see [Figure 2–1](#).**

FIGURE 2-1 Disconnecting the Node from Shared Storage by Using Host Adapters on Node B



- Disconnect the SCSI cable between Node A and the storage array. Install an appropriate SCSI terminator to this SCSI connector on the storage array. For an illustration, see [Figure 2-2](#).

FIGURE 2-2 Disconnecting the Node from Shared Storage by Using Terminators



Ⓣ Terminator

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

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

**9 If you detached the submirrors in Step 5, reattach the submirrors. Wait for the submirrors to resynchronize.**

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

**10 Repeat Step 5 through Step 9 for each remaining storage array that is connected to Node A.**

**11 Reconnect the node to the shared storage as outlined in “How to Reconnect the Node to Shared Storage” on page 46.**

▼ **How to Reconnect the Node to Shared Storage**

**1 Become superuser or assume a role that provides `soLaris.cLuster.mOdiFY` RBAC authorization.**

**2 (Optional) If necessary, detach the submirrors on the storage array.**

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

**3 Disconnect the SCSI terminator. Reattach the SCSI cable between the storage array and Node A.**




---

**Caution** – Connect this storage array to the same host adapter that the storage array was connected before you disconnected the SCSI cable.

---

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

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

**5 If you detached submirrors in Step 1, reattach the submirrors on the storage array. Wait for the submirrors to resynchronize.**

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

**6 Repeat Step 1 through Step 5 for the remaining storage array that you want to reconnect to Node A.**



**3 Detach the submirror or submirrors on the storage array.**

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

**4 Replace the SCSI cable.**

**5 Reattach the submirror or submirrors on the storage array.**

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

**6 If you relocated a quorum device in [Step 1](#), relocate the quorum device function to this storage array.**

To add and remove quorum devices, see the Oracle Solaris Cluster system administration documentation.

## Replacing a Host Adapter

You need to replace a host adapter if your host adapter fails, if it becomes unstable, or if you want to upgrade to a newer version. These procedures define Node A as the node with the host adapter that you plan to replace.

Choose the procedure that corresponds to your cluster configuration.

---

**x86 only** – If your cluster is x86 based, Oracle RAC services are not supported. Follow the instructions in [“How to Replace a Host Adapter When Using Failover and Scalable Data Services Only”](#) on page 49.

---

Cluster Configuration	Instructions
Oracle Solaris Cluster failover and scalable data services only, using the recommended HBA configuration	<a href="#">“How to Replace a Host Adapter When Using Failover and Scalable Data Services Only”</a> on page 49
Oracle Parallel Server/Real Application Clusters (OPS/RAC) only, using the recommended HBA configuration	<a href="#">“How to Replace a Host Adapter When Using Oracle Real Application Clusters Only”</a> on page 52
Both failover and scalable data services and OPS/RAC, using the recommended HBA configuration	<a href="#">“How to Replace a Host Adapter When Using Both Failover and Scalable Data Services and Oracle Real Application Clusters”</a> on page 55

Cluster Configuration	Instructions
All clusters using a single, dual-port HBA to provide both paths to shared data	<a href="#">“How to Replace a Host Adapter When Using a Single, Dual-Port HBA to Provide Both Paths to Shared Data”</a> on page 60

**Note** – The first three procedures in this section assume that you are using the recommended HBA configuration: two redundant hardware paths to shared data. If you choose to use a single HBA configuration, see [“Configuring Cluster Nodes With a Single, Dual-Port HBA”](#) in *Oracle Solaris Cluster 3.3 Hardware Administration Manual* for the risks and restrictions of that configuration and use [“How to Replace a Host Adapter When Using a Single, Dual-Port HBA to Provide Both Paths to Shared Data”](#) on page 60.

## ▼ How to Replace a Host Adapter When Using Failover and Scalable Data Services Only

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

- You are *not* using Oracle Real Application Clusters. If you are using Oracle Real Application Clusters, follow the instructions in [“How to Replace a Host Adapter When Using Oracle Real Application Clusters Only”](#) on page 52 or [“How to Replace a Host Adapter When Using Both Failover and Scalable Data Services and Oracle Real Application Clusters”](#) on page 55.
- Except for the failed host adapter, your cluster is fully operational and all nodes are powered on.

If your nodes are configured for dynamic reconfiguration, see the Oracle Solaris Cluster system administration documentation, and skip steps that instruct you to shut down the node.

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 later in this procedure to return resource groups and device groups to Node A.

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

**3 (Optional) If necessary, move all resource groups and device groups off Node A.**

```
# clnode evacuate NodeA
```

- 4 If the storage device that is attached to the failed host adapter is configured as a quorum device, add a new quorum device on a storage device that is not affected by this procedure. Then remove the old quorum device.**

To determine whether the affected device contains a quorum device, use the following command.

```
# clquorum show
```

To add and remove quorum devices, see the Oracle Solaris Cluster system administration documentation.

- 5 Detach the Solaris Volume Manager submirrors or Veritas Volume Manager plexes on the storage array.**

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

- 6 Record the details of disk groups and volumes affected by the failed host adapter.**

Record this information because you use it in [Step 16](#) of this procedure to reattach submirrors on the storage array. To determine which submirrors or plexes are affected, see your Solaris Volume Manager or Veritas Volume Manager documentation.

- 7 If Node A is enabled with the Oracle Solaris DR feature, perform any DR-specific steps and skip to [Step 10](#).**

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

- 8 Shut down Node A.**

To shut down and power off a node, see your Oracle Solaris Cluster system administration documentation.

- 9 Power off Node A.**

- 10 Replace the failed host adapter.**

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

- 11 If Node A is enabled with the Oracle Solaris DR feature, perform any DR-specific steps and skip to [Step 15](#).**

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

- 12 Power on Node A.**

**13 x86: Set the HBA ports to ensure that each array has a unique SCSI address.**

For instructions on setting SCSI initiator IDs in x86 based systems, see “x86: How to Install a Storage Array in a New x86 Based Cluster” on page 18.

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

For the procedure about booting cluster nodes, see Chapter 3, “Shutting Down and Booting a Cluster,” in *Oracle Solaris Cluster System Administration Guide*.

**15 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*.

**16 Reattach the Solaris Volume Manager submirrors or Veritas Volume Manager plexes on the storage array.**

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

**17 (Optional) If you moved device groups off the node in Step 4 of the disconnect procedure, move all device groups back to the node.**

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

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

-n NodeA

The node to which you are restoring device groups.

devicegroup1[ devicegroup2 ...]

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

- 18 (Optional) If you moved resource groups off the node in Step 4 of the disconnect procedure, move all resource groups back to the node.**

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

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

*-n NodeA*

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.

- 19 If you relocated a quorum device in Step 4, and if you want the cluster configured as it was before replacing the host adapter, relocate the quorum device function back to this storage array.**

To add and remove quorum devices, see your Oracle Solaris Cluster system administration documentation.

## ▼ How to Replace a Host Adapter When Using Oracle Real Application Clusters Only

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

- You are using Oracle Real Application Clusters *only*. If you are using other failover or scalable data services, follow the instructions in “[How to Replace a Host Adapter When Using Failover and Scalable Data Services Only](#)” on page 49 or “[How to Replace a Host Adapter When Using Both Failover and Scalable Data Services and Oracle Real Application Clusters](#)” on page 55.
- Except for the failed host adapter, your cluster is fully operational and all nodes are powered on.

If your nodes are configured for dynamic reconfiguration, see the Oracle Solaris Cluster system administration documentation, and skip steps that instruct you to shut down the node.

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 Oracle instance that is running on Node A.**

```
# ps -ef | grep oracle
```

- 3 Shut down the Oracle Real Application Clusters instance and any other process on Node A that should be stopped before shutting down the node.**

To shut down and restart an Oracle instance in the RAC environment, refer to your Oracle documentation.

- 4 If the storage devices that are attached to the failed host adapter contain a quorum device, add a new quorum device on a different storage device. Then remove the old quorum device.**

To determine whether the affected device contains a quorum device, use the following command.

```
# clquorum show
```

To add and remove quorum devices, see the Oracle Solaris Cluster system administration documentation.

- 5 Detach the Veritas Volume Manager plexes on the storage array attached to the failed host adapter.**

For more information, see your Veritas Volume Manager documentation.

- 6 Record the details of disk groups and volumes affected by the failed host adapter**

Record this information because you use it in [Step 15](#) of this procedure to reattach plexes on the storage array. To determine which plexes are affected, see your Veritas Volume Manager documentation.

- 7 If Node A is enabled with the Oracle Solaris dynamic reconfiguration (DR) feature, perform any DR-specific steps and skip to [Step 10](#).**

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

- 8 Shut down Node A.**

To shut down and power off a node, see your Oracle Solaris Cluster system administration documentation.

- 9 Power off Node A.**

- 10 Replace the failed host adapter.**

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

- 11 If Node A enabled with the Oracle Solaris dynamic reconfiguration feature, perform any DR-specific steps and skip to [Step 14](#).**

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

**12 Power on Node A.****13 Boot Node A into cluster mode.**

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

**14 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*.

**15 Reattach the Veritas Volume Manager plexes on the storage array to their respective volumes.**

For more information, see your Veritas Volume Manager documentation.

**16 (Optional) Bring your Oracle Real Application Clusters instance online. This is the instance you identified in [Step 1](#).**

To shut down and restart an Oracle instance in the RAC environment, refer to the *Oracle 9iRAC Administration Guide*.

**17 (Optional) If you relocated a quorum device in [Step 4](#), and if you want your configuration to use the same quorum structure after the host adapter replacement, relocate the quorum device function back to this storage array.**

To add and remove quorum devices, see the Oracle Solaris Cluster system administration documentation.

## ▼ How to Replace a Host Adapter When Using Both Failover and Scalable Data Services and Oracle Real Application Clusters

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

- You are using both Oracle Real Application Clusters and failover or scalable data services.  
If you are using only Oracle Real Application Clusters, follow the instructions in “[How to Replace a Host Adapter When Using Oracle Real Application Clusters Only](#)” on page 52. If you are using only failover and scalable data services, follow the instructions in “[How to Replace a Host Adapter When Using Failover and Scalable Data Services Only](#)” on page 49.
- Except for the failed host adapter, your cluster is operational and all nodes are powered on.

If your nodes are configured for dynamic reconfiguration, see the Oracle Solaris Cluster system administration documentation, and skip steps that instruct you to shut down the node.

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 Oracle instance that is running on Node A.**

```
# ps -ef | grep oracle
```

**3 Shut down the Oracle Real Application Clusters instance identified in [Step 1](#).**

To shut down and restart an Oracle instance in the RAC environment, refer to the *Oracle 9iRAC Administration Guide*.

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

Record this information because you use it later in this procedure to return resource groups and device groups to Node A.

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

**5 (Optional) If necessary, move all resource groups and device groups off Node A.**

```
# clnode evacuate NodeA
```

**6 If the storage device that is connected to the failed host adapter contains a quorum device, add a new quorum device on a different storage device. Then remove the old quorum device.**

To determine whether the affected device contains a quorum device, use the following command.

```
# clquorum show
```

To add and remove quorum devices, see the Oracle Solaris Cluster system administration documentation.

**7 Detach the Veritas Volume Manager plexes on the storage array that is attached to the failed host adapter.**

For more information, see your Veritas Volume Manager documentation.

**8 Record the details of plexes that are affected by the failed host adapter.**

Record this information because you use it in [Step 17](#) of this procedure to reattach plexes on the storage array. To determine which plexes are affected, see your Veritas Volume Manager documentation.

**9 If Node A is enabled with the Oracle Solaris dynamic reconfiguration (DR) feature, perform any DR-specific steps and skip to [Step 12](#).**

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

**10 Shut down Node A.**

To shut down and power off a node, see your Oracle Solaris Cluster system administration documentation.

**11 Power off Node A.**

**12 Replace the failed host adapter.**

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

**13 If Node A enabled with the Oracle Solaris DR feature, perform any DR-specific steps and skip to [Step 16](#).**

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

**14 Power on Node A.**

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

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

**16 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*.

**17 Reattach the Veritas Volume Manager plexes on the storage array to their respective volumes.**

For more information, see your Veritas Volume Manager documentation.

**18 (Optional) If you moved device groups off the node in Step 4 of the disconnect procedure, move all device groups back to the node.**

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

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

-n NodeA

The node to which you are restoring device groups.

devicegroup1[ devicegroup2 ...]

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

**19 (Optional) If you moved resource groups off the node in Step 4 of the disconnect procedure, move all resource groups back to the node.**

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

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

-n NodeA

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

resourcegroup1[ resourcegroup2 ...]

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

**20 (Optional) Bring your Oracle Real Application Clusters instance online. This is the instance you identified in Step 1.**

To shut down and restart an Oracle instance in the RAC environment, refer to the *Oracle 9iRAC Administration Guide*.

- 21 (Optional) If you relocated a quorum device in [Step 6](#) and you want your configuration to use the same quorum structure after the host adapter replacement, relocate the quorum device function back to this storage array.

To add and remove quorum devices, see the Oracle Solaris Cluster system administration documentation.

### Example 2–1 SPARC: Replacing a Host Adapter in a Running Cluster

In the following example, a two-node cluster is running Oracle Real Application Clusters and Veritas Volume Manager. In this situation, you begin the host adapter replacement by determining the Oracle instance name.

```
# ps -ef | grep oracle
oracle 14716 14414 0 14:05:47 console 0:00 grep oracle
oracle 14438 1 0 13:05:44 ? 0:02 ora_lmon_tpcc1
.
.
oracle 14434 1 0 13:05:43 ? 0:00 ora_pmon_tpcc1
oracle 14458 1 0 13:05:50 ? 0:00 ora_d000_tpcc1
```

This output identifies the Oracle Real Application Clusters instance as **tpcc1**.

Shutting down the Oracle Real Application Clusters instance on Node A involves several steps, as shown in the following example.

```
# su - oracle
Sun Microsystems Inc. SunOS 5.9 Generic May 2002
$ ksh
$ ORACLE_SID=tpcc1
$ ORACLE_HOME=/export/home/oracle/OraHome1
$ export ORACLE_SID ORACLE_HOME
$ sqlplus "/as sysdba "
```

```
SQL*Plus: Release 9.2.0.4.0 - Production on Mon Jan 5 14:12:28 2004
Copyright (c) 1982, 2002, Oracle Corporation. All rights reserved.
```

```
Connected to:
Oracle9i Enterprise Edition Release 9.2.0.4.0 - 64bit Production
With the Partitioning, Real Application Clusters, OLAP and Oracle Data
Mining options
JServer Release 9.2.0.4.0 - Production
```

```
SQL> shutdown immediate ;
Database closed.
Database dismounted.
ORACLE instance shut down.
```

```
SQL> exit
Disconnected from Oracle9i Enterprise Edition Release 9.2.0.4.0 - 64bit
```

```

Production With the Partitioning, Real Application Clusters, OLAP and
Oracle Data Mining options
JServer Release 9.2.0.4.0 - Production
$ lsnrctl

```

```
LSNRCTL for Solaris: Version 9.2.0.4.0 - Production on 05-JAN-2004 14:15:09
```

```
Copyright (c) 1991, 2002, Oracle Corporation. All rights reserved.
```

```
Welcome to LSNRCTL, type "help" for information.
```

```
LSNRCTL> stop
```

```
Connecting to (DESCRIPTION=(ADDRESS=(PROTOCOL=IPC)(KEY=EXTPROC)))
```

```
The command completed successfully
```

```
LSNRCTL>
```

1. After you have stopped the Oracle Real Application Clusters instance, check for a quorum device and, if necessary, reconfigure the quorum device.
2. When you are certain that the node with the failed adapter does not contain a quorum device, proceed to determine the affected plexes.

Record this information for use in reestablishing the original storage configuration. In the following example, c2 is the controller with the failed host adapter.

```

# vxprint -ht -g tpcc | grep c2
dm tpcc01      c2t0d0s2      sliced  4063    8374320  -
dm tpcc02      c2t1d0s2      sliced  4063    8374320  -
dm tpcc03      c2t2d0s2      sliced  4063    8374320  -
dm tpcc04      c2t3d0s2      sliced  4063    8374320  -
dm tpcc09      c2t8d0s2      sliced  4063    8374320  -
dm tpcc10      c2t9d0s2      sliced  4063    8374320  -
sd tpcc02-01  control_001-01 tpcc02  0        41040    0        c2t1d0  ENA
.
.
sd tpcc03-06  temp_0_0-02   tpcc03  2967840 276480  0        c2t2d0  ENA
sd tpcc03-04  ware_0_0-02   tpcc03  2751840 95040   0        c2t2d0  ENA

```

From this output, you can easily determine which plexes and subdisks are affected by the failed adapter. These are the plexes you detach from the storage array.

```

/usr/sbin/vxplex -g tpcc det control_001-02
/usr/sbin/vxplex -g tpcc det temp_0_0-02

```

1. After the plexes are detached, you can safely shut down the node, if necessary.
2. Proceed with replacing the failed host adapter, following instructions that accompanied that device.
3. After you replace the failed host adapter, and Node A is in cluster mode, reattach the plexes and replace any quorum device to reestablish your original cluster configuration.

## ▼ How to Replace a Host Adapter When Using a Single, Dual-Port HBA to Provide Both Paths to Shared Data

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

- You are using a single, dual-port HBA to provide the redundant paths to your shared data.  
For an explanation of the limitations and risks of using this configuration, see “[Configuring Cluster Nodes With a Single, Dual-Port HBA](#)” in *Oracle Solaris Cluster 3.3 Hardware Administration Manual*.
- Except for the failed host adapter, your cluster is operational and all nodes are powered on.
- Your nodes are not configured for dynamic reconfiguration.  
You cannot use dynamic reconfiguration for this procedure when using a single HBA.
- Node A is the node with the failed host adapter.

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 you are using scalable and failover services, determine the resource groups and device groups that are running on Node A.**  
Record this information because you use it in [Step 17](#) of this procedure to return resource groups and device groups to Node A.  

```
# clresourcegroup status -n NodeA
# cldevicegroup status -n NodeA
```
- 3 Record the details of metadevices that are affected by the failed host adapter.**
- 4 SPARC: If you are using Oracle Real Application Clusters, shut down all RAC instances running in your cluster.**  
To shut down and restart an Oracle instance in the RAC environment, refer to the *Oracle 9iRAC Administration Guide*.
- 5 Shut down the cluster.**  
To shut down a cluster, see your Oracle Solaris Cluster system administration documentation.
- 6 Power off Node A.**
- 7 Replace the failed host adapter.**  
To remove and add host adapters, see the documentation that shipped with your nodes.

**8 Power on Node A.****9 Boot all nodes into cluster mode.**

For the procedure about booting cluster nodes, see [Chapter 3, “Shutting Down and Booting a Cluster,”](#) in *Oracle 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 Perform any volume management maintenance procedures that are necessary to fix any metadevices affected by this procedure.**

For more information, see your volume manager software documentation.

**12 (Optional) If necessary, move the device groups back to the original node.**

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

```
# cldevicegroup switch -n NodeA devicegroup
```

**13 (Optional) If necessary, move the resource groups back to the node.**

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

```
# clresourcegroup switch -n NodeA resourcegroup
```

*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.

**14 (Optional) Bring all Oracle Real Application Clusters instance online.**

To shut down and restart an Oracle instance in the RAC environment, refer to the *Oracle 9iRAC Administration Guide*.

**▼ How to Add a Disk Drive**

Adding a disk drive enables you to increase your storage space after a storage array has been added to your cluster.



**Caution** – (*Sun StorEdge Multipack Enclosures Only*) SCSI-reservations failures have been observed when clustering storage arrays that contain a particular model of Quantum disk drive: SUN4.2G VK4550J. Do not use this particular model of Quantum disk drive for clustering with storage arrays.

If you must use this model of disk drive, you must set the `scsi-initiator-id` of Node A to 6. If you use a six-slot storage array, you must set the storage array for the 9-through-14 SCSI target address range. For more information, see the *Sun StorEdge MultiPack Storage Guide*.

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

- Your storage array contains an empty disk slot.
- Your nodes are not configured with dynamic reconfiguration functionality.  
If your nodes are configured for dynamic reconfiguration, see the Oracle Solaris Cluster system administration documentation, and skip steps that instruct you to shut down the node.

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

**1 Locate an empty disk slot in the storage array for the disk drive you are adding.**

Identify the empty disk slot and note the target number. Refer to your storage array documentation.

**2 Install the disk drive.**

For the procedure about how to install a disk drive, see your storage array documentation.

**3 On all nodes that are attached to the storage array, configure the disk drive.**

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

- 4 **On all nodes that are attached to the storage device, probe all devices and write the new disk drive to the `/dev/rdisk` directory.**

```
# devfsadm
```

Depending on the number of devices that are connected to the node, the `devfsadm` command can require at least five minutes to complete.

- 5 **On all nodes, verify that the entries for the disk drive have been added to the `/dev/rdisk` directory.**

```
# ls -l /dev/rdisk
```

- 6 **If necessary, use the `format(1M)` command or the `fmthard(1M)` command to partition and label the disk drive.**

- 7 **From any node, update the global device namespace.**

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 an expected behavior.

```
# cldevice populate
```

- 8 **On one node, verify that a device ID has been assigned to the disk drive.**

The new device ID that is assigned to the new disk drive might not be in sequential order in the storage array.

```
# cldevice list -v
```

- 9 **Perform volume management administration to add the new disk drive to the configuration.**

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

- 10 **If you want this new disk drive to be a quorum device, add the quorum device.**

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

## ▼ How to Remove a Disk Drive

Removing a disk drive can allow you to downsize or reallocate your existing storage pool. You might want to perform this procedure in the following scenarios.

- You no longer need to make data accessible to a particular node.
- You want to migrate a portion of your storage to another storage array.

For conceptual information about quorum, quorum devices, global devices, and device IDs, see the Oracle Solaris Cluster concepts documentation.

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

- You do not need to remove the entire storage array.  
If you need to remove the storage array, see [“How to Remove a Storage Array” on page 39](#).
- You do not need to replace the storage array's chassis.  
If you need to replace your storage array's chassis, see [“How to Replace the Chassis” on page 41](#).
- Your nodes are not configured with dynamic reconfiguration functionality.  
If your nodes are configured for dynamic reconfiguration, see the Oracle Solaris Cluster system administration documentation, and skip steps that instruct you to shut down the node.

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

**1 Identify the disk drive that you are removing. Also identify the slot that the disk drive needs to be removed from.**

If the disk error message reports the drive problem by device ID, use the `cldevice list` to determine the Oracle Solaris device name. To list all configurable hardware information, use the `cfgadm -al` command.

```
# cldevice list -v
# cfgadm -al
```

**2 If the disk you are removing is configured as quorum device, add a new quorum device that will not be affected by this procedure. Then remove the old quorum device.**

To determine whether a quorum device will be affected by this procedure, use one of the following commands.

```
# clquorum show +
```

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

**3 If possible, back up the metadvice or volume.**

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

**4 Perform volume management administration to remove the disk drive from the configuration.**

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

- 5 **On all nodes, unconfigure the disk drive. Remove references to the disk drive from the operating system and clustering environment.**

```
# cfgadm -c unconfigure cN::disk/cNtXdY
```

- 6 **Remove the disk drive.**

For the procedure about how to remove a disk drive, see your storage documentation. For a list of storage documentation, see [“Related Documentation” on page 6](#).

- 7 **On both nodes, delete the paths to the disk drive that you plan to remove.**

```
# devfsadm -C
```

- 8 **On all nodes, remove all obsolete device IDs.**

```
# cldevice clear
```

## ▼ **How to Replace a Disk Drive *Without* Oracle Real Application Clusters**

You need to replace a disk drive if the disk drive fails or when you want to upgrade to a higher-quality or to a larger disk.

For conceptual information about quorum, quorum devices, global devices, and device IDs, see the Oracle Solaris Cluster concepts documentation.

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

- Your cluster is operational.
- Your system is *not* running Oracle Real Application Clusters.
  - If your system is running Oracle Real Application Clusters, see [“SPARC: How to Replace a Disk Drive \*With\* Oracle Real Application Clusters” on page 68](#).
- (*Solaris Volume Manager Only*) If the disk drive failure prevents Solaris Volume Manager from reading the disk label, you have a backup of the disk-partitioning information.
- Your nodes are not configured with dynamic reconfiguration functionality.
  - If your nodes are configured for dynamic reconfiguration, see the Oracle Solaris Cluster system administration documentation, and skip steps that instruct you to shut down the node.

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

- 1 **Become superuser or assume a role that provides `solaris.cluster.read` and `solaris.cluster.modify` RBAC authorization.**
- 2 **Identify the failed disk drive.**  

```
# cldevice show -v cNtNdN
```
- 3 **Record the device identifier (DID) of the failed disk drive because you assign it to the replaced disk drive later in this procedure.**

- 4 **If the disk drive you are removing is configured as a quorum device, add a new quorum device that will not be affected by this procedure. Then remove the old quorum device.**

To determine whether a quorum device will be affected by this procedure, use one of the following commands.

```
# clquorum show +
```

To add and remove quorum devices, see the Oracle Solaris Cluster system administration documentation.

- 5 **If possible, back up the metadvice or volume.**

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

- 6 *(Solaris Volume Manager Only)* **If your disk drive failure prevents Solaris Volume Manager from reading the disk label, use the disk partitioning information that you saved.**

You saved this information when you performed one of the following tasks:

- Installed your storage array in an initial cluster as outlined in [Step 10](#) of “[SPARC: How to Install a Storage Array in a New SPARC Based Cluster](#)” on page 13
- Added the storage array to an operational cluster.

- 7 *(Solaris Volume Manager Only)* **If your disk drive failure does not prevent Solaris Volume Manager from reading the disk label, save the disk partitioning information now if you have not already done so.**



**Caution** – Do not save disk partitioning information under `/tmp` because you will lose this file when you reboot. Instead, save this file under `/usr/tmp`.

---

```
# prtvtoc /dev/rdisk/cNtNdNsN > filename
```

Use this information when you partition the new disk drive.

## 8 Replace the failed disk.

- a. Determine which node owns the device group.

```
# cldevicegroup status devgroup1
```

- b. If you are using Veritas Volume Manager, remove the disk drives from the Veritas Volume Manager control on a node that does not have ownership of the device group.

```
# vxdisk offline cNtNdN
# vxdisk rm cNtNdN
```

- c. On a node that does not have ownership of the device group, suspend activity on the SCSI bus.

```
# cfgadm -x replace_device cN::disk/cNtNdN
```

When prompted, type **y** to suspend activity on the SCSI bus.

- d. If the message `cfgadm: Component system is busy, try again: failed to offline` is displayed, follow these steps:

- i. Become superuser.

- ii. Temporarily rename the file named `es_rcm.pl`.

- If you are using Version 4.1 of Veritas Volume Manager or a version of Veritas Volume Manager that was released *after* 4.1, type:

```
# mv /usr/lib/rcm/scripts/es_rcm.pl /usr/lib/rcm/scripts/DONTUSE
```

- If you are using a version of Veritas Volume Manager that was released *before* Version 4.1, type:

```
# mv /etc/rcm/scripts/es_rcm.pl /etc/rcm/scripts/DONTUSE
```

- iii. Reissue the `cfgadm` command that you tried to issue previously.

```
# cfgadm -x replace_device cN::disk/cNtNdN
```

- iv. Rename the `DONTUSE` file to its original name.

- If you are using Version 4.1 of Veritas Volume Manager or a version of Veritas Volume Manager that was released *after* 4.1, type:

```
# mv /usr/lib/rcm/scripts/DONTUSE /usr/lib/rcm/scripts/es_rcm.pl
```

- If you are using a version of Veritas Volume Manager that was released *before* Version 4.1, type:

```
# mv /etc/rcm/scripts/DONTUSE /etc/rcm/scripts/es_rcm.pl
```

**e. After SCSI bus activity stops, replace the disk and type y at the prompt.**

After replacing the disk, warning messages might be displayed. Ignore these messages.

**9 On all nodes that are attached to the device, run the `devfsadm` command to probe all devices and to update the device tree.**

```
# devfsadm
```

Depending on the number of devices that are connected to the node, the `devfsadm(1M)` command can require at least five minutes to complete.

**10 Label the new disk drive by using the `format` command.**

**11 (Solaris Volume Manager Only) If you successfully saved the disk partitioning information in [Step 7](#), from any node that is connected to the device, partition the new disk drive by using the partitioning you saved when you installed or added the storage array.**

```
# fmthard -s filename /dev/rdisk/cNtNdNsN
```

**12 On all nodes, repair the device instance for the replaced disk drive.**

```
# cldevice repair DID_number
```

**13 Perform volume management administration to add the disk drive back to its diskset or disk group.**

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

**14 If you want this new disk drive to be a quorum device, add the quorum device.**

To add and remove quorum devices, see the Oracle Solaris Cluster system administration documentation.

## ▼ **SPARC: How to Replace a Disk Drive *With Oracle Real Application Clusters***

You need to replace a disk drive if the disk drive fails or when you want to upgrade to a higher-quality or to a larger disk.

For conceptual information about quorum, quorum devices, global devices, and device IDs, see the Oracle Solaris Cluster concepts documentation.

While performing this procedure, ensure that you use the correct controller number. Controller numbers can be different on each node.

---

**Note** – If the following warning message is displayed, ignore the message. Continue with the next step.

```
vxvm:vxconfigd: WARNING: no transactions on slave
vxvm:vxassist: ERROR: Operation must be executed on master
```

---

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

- Your cluster is operational.
- Your system is running Oracle Real Application Clusters.
 

If your system is *not* running Oracle Real Application Clusters, see [“How to Replace a Disk Drive \*Without\* Oracle Real Application Clusters”](#) on page 65.
- Your nodes are not configured with dynamic reconfiguration functionality.
 

If your nodes are configured for dynamic reconfiguration, see the Oracle Solaris Cluster system administration documentation, and skip steps that instruct you to shut down the node.

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

**1 On each node, check the status of the failed disk drive.**

```
# vxdisk list
```

**2 On each node, identify the failed disk's instance number.**

You use this instance number in [Step 12](#).

```
# ls -l /dev/dsk/cWtXdYsZ
# cat /etc/path_to_inst | grep "device_path"
```

---

**Note** – Ensure that you do not use the `ls` command output as the device path. The following example demonstrates how to find the device path and how to identify the `sd` instance number by using the device path.

```
# ls -l /dev/dsk/c4t0d0s2
lrwxrwxrwx 1 root root 40 Jul 31 12:02 /dev/dsk/c4t0d0s2 ->
../../../../devices/pci@4,2000/scsi@1/sd@0,0:cb41d
# cat /etc/path_to_inst | grep "/pci@4,2000/scsi@1/sd@0,0"
"/node@2/pci@4,2000/scsi@1/sd@0,0" 60 "sd"
```

The `node@2` portion of this output is not present. Oracle Solaris 10 does not add this prefix for cluster nodes.

---

**3 On one node, identify the failed disk's disk ID number.**

You will use the disk ID number in [Step 15](#) and [Step 16](#).

```
# cldevice show -v cNtXdY
```

**4 If the disk drive you are removing is configured as quorum device, add a new quorum device that will not be affected by this procedure. Then remove the old quorum device.**

To determine whether a quorum device will be affected by this procedure, use one of the following commands.

```
# clquorum show +
```

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

**5 On each node, remove the failed disk from Veritas Volume Manager control.**

```
# vxdisk offline cXtYdZ  
# vxdisk rm cXtYdZ
```

**6 On each node, verify that you removed the disk entry.**

```
# vxdisk list
```

**7 Remove the failed disk from the storage array.**

For the procedure about how to remove a disk drive, see your storage documentation. For a list of storage documentation, see [“Related Documentation”](#) on page 6.

**8 On each node, unconfigure the failed disk.**

```
# cfgadm -c unconfigure cX::dsk/cXtYdZ
```

**9 On each node, remove the paths to the disk drive that you are removing.**

```
# devfsadm -C
```

**10 On each node, verify that you removed the disk.**

```
# cfgadm -al | grep cXtYdZ  
# ls /dev/dsk/cXtYdZ
```

**11 Add the new disk to the storage array.**

For the procedure about how to add a disk drive, see your storage documentation. For a list of storage documentation, see [“Related Documentation”](#) on page 6.

**12 On each node, configure the new disk.**

Use the instance number that you identified in [Step 2](#).

```
# cfgadm -c configure cX::sd_instance_Number  
# devfsadm
```

- 13 **Verify that you added the disk.**  

```
# ls /dev/dsk/cXtYdZ
```
- 14 **On one node, update the device ID numbers.**  
 Use the device ID number that you identified in [Step 3](#).  

```
# cldevice repair
```
- 15 **Verify that you updated the disk's device ID number.**  
 Use the device ID number that you identified in [Step 3](#).  

```
# cldevice show | grep DID_number
```
- 16 **Verify that the disk ID number is different than the disk ID number that you identified in [Step 3](#).**
- 17 **On each node, add the new disk to the Veritas Volume Manager database.**  

```
# vxctl enable
```
- 18 **Verify that you added the new disk.**  

```
# vxdisk list |grep cXtYdZ
```
- 19 **Determine the master node.**  

```
# vxctl -c mode
```
- 20 **Perform disk recovery tasks on the master node.**  
 Depending on your configuration and volume layout, select the appropriate Veritas Volume Manager menu item to recover the failed disk.  

```
# vxdiskadm  
# vxtask list
```

## ▼ How to Upgrade Disk Drive Firmware

Upgrade your disk drive firmware if you want to apply bug fixes or enable new functionality.

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

- Your cluster is operational.
- Your nodes are not configured with dynamic reconfiguration functionality.  
 If your nodes are configured for dynamic reconfiguration, see the Oracle Solaris Cluster system administration documentation, and skip steps that instruct you to shut down the node.

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

- 1 If the disk drive on which you want to upgrade the firmware is configured as quorum device, add a new quorum device that will not be affected by this procedure. Then remove the old quorum device.**

To determine whether a quorum device will be affected by this procedure, use the following command.

```
# clquorum show
```

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

- 2 If possible, back up the metadevices or volumes that reside in the disk drive.**

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

- 3 Perform volume management administration to remove the disk drive from the configuration.**

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

- 4 Install the firmware.**

For more information about how to install firmware, see the patch installation instructions.

- 5 Perform volume management administration to add the disk drive back into the configuration.**

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

## ▼ How to Upgrade Host Adapter Firmware

Upgrade your host adapter firmware if you want to apply bug fixes or enable new functionality.

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

- Your cluster is operational.
- This procedure defines Node A as the node on which you are upgrading the host adapter firmware.

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 the node on which you plan to upgrade the host adapter firmware.**

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

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

```
# clnode evacuate NodeA
```

- 3 Perform the firmware upgrade.**

For the procedure about how to upgrade your host adapter firmware, see the patch documentation.

- 4 (Optional) Return the device groups back to Node A.**

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

```
# cldevicegroup switch -n NodeA devicegroup
```

- 5 (Optional) Return the resource groups back to Node A.**

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

```
# clresourcegroup switch -n NodeA resourcegroup
```

*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.



## Cabling Diagrams

---

This chapter contains illustrations to supplement the procedures in this book that involve cabling your storage array.

- [Sun StorEdge 3120 Storage Array](#)
- [Sun StorEdge 3310 SCSI Storage Array](#)
- [SPARC: Sun StorEdge D1000 Storage Array](#)
- [SPARC: Sun StorEdge Multipack Storage Array](#)
- [SPARC: Sun StorEdge Netra D130/S1 Storage Array](#)
- [SPARC: Sun StorEdge D2 Storage Array](#)

### Sun StorEdge 3120 Cabling Diagrams

#### Installing a Sun StorEdge 3120 Storage Array

For the procedure about how to install and cable the storage arrays, see the *Sun StorEdge 3120 SCSI Array Quick Installation Guide*. For a URL to this storage documentation, see [“Related Documentation”](#) on page 6.

---

**Note** – In the configuration that is outlined in [Figure 3–1](#), you must mirror the Sun StorEdge 3120 storage array against another storage array. This storage array need not be a Sun StorEdge 3120 storage array.

---

FIGURE 3-1 Cabling a Single-Bus, Multi-Initiator Sun StorEdge 3120 Storage Array





FIGURE 3-2 Cabling a Dual-Bus, Multi-Initiator Sun StorEdge 3120 Storage Array




## Sun StorEdge 3310 and 3320 SCSI Cabling Diagrams

### Installing a Sun StorEdge 3310 or 3320 Storage Array

Ensure that the cable does not exceed bus-length limitations. For more information about bus-length limitations, see the documentation that shipped with your hardware.

For the procedure about how to cable the storage arrays, see the *Sun StorEdge 3000 Family Installation, Operation, and Service Manual*. For a URL to this storage documentation, see [“Related Documentation”](#) on page 6.

FIGURE 3-3 Cabling a Dual-Bus, Multi-Initiator Sun StorEdge 3310 or 3320 Storage Array



## Adding a Sun StorEdge 3310 or 3320 Storage Array

If the storage array is in dual-bus, multi-initiator mode, see [Figure 3-4](#) for a diagram about how to cable and terminate the storage arrays.




**Caution** – Do not power on the storage array until the storage array is cabled to the first node.

Ensure that the cable does not exceed bus-length limitations. For more information about bus-length limitations, see the documentation that shipped with your hardware.

For the procedure about how to cable the storage arrays, see the *Sun StorEdge 3000 Family Installation, Operation, and Service Manual*. For a URL to this storage documentation, see “[Related Documentation](#)” on page 6.

FIGURE 3-4 Cabling a Dual-Bus, Multi-Initiator Sun StorEdge 3310 or 3320 Storage Array: Part 1




Ⓣ Terminator

If the storage array is in dual-bus, multi-initiator mode, remove the terminators and connect the storage array to the second node. See [Figure 3-5](#) for a diagram about how to cable the storage array to both nodes.

For the procedure about how to cable the storage arrays, see the *Sun StorEdge 3000 Family Installation, Operation, and Service Manual*. For a URL to this storage documentation, see [“Related Documentation”](#) on page 6.

Ensure that the cable does not exceed bus-length limitations. For more information about bus-length limitations, see the documentation that shipped with your hardware.

FIGURE 3-5 Cabling a Dual-Bus, Multi-Initiator Sun StorEdge 3310 or 3320 Storage Array: Part 2




## SPARC: Sun StorEdge D1000 Cabling Diagrams

### SPARC: Installing a Sun StorEdge D1000 Storage Array

Connect the cables to the storage arrays, as shown in [Figure 3-6](#).

Ensure that the *entire* bus length that is connected to each storage array is less than 25 m. This measurement includes the cables to both nodes, as well as the bus length that is internal to each storage array, node, and the host adapter.


FIGURE 3-6 SPARC: Installing a Sun StorEdge D1000 Mirrored-Pair Configuration: An Example



## SPARC: Adding a Sun StorEdge D1000 Storage Array


Ensure that the *entire* SCSI bus length to each enclosure is less than 25 m. This measurement includes the cables to both nodes, as well as the bus length that is internal to each enclosure, node, and host adapter. Refer to the documentation that shipped with the enclosure for other restrictions about SCSI operation.

FIGURE 3-7 SPARC: Adding a Sun StorEdge D1000 Mirrored-Pair Configuration: Part 1



Connect the storage array to the host adapters by using differential SCSI cables as shown in see [Figure 3-8](#).


FIGURE 3-8 SPARC: Adding a Sun StorEdge D1000 Mirrored-Pair Configuration: Part 2



## SPARC: Replacing a Sun StorEdge D1000 Storage Array

Disconnect the SCSI cables from the storage array. Remove the SCSI cables in the order that is shown in [Figure 3–9](#).


FIGURE 3–9 SPARC: Disconnecting the SCSI Cables From a Sun StorEdge D1000 Storage Array



## SPARC: Removing a Sun StorEdge D1000 Storage Array

Disconnect the SCSI cables from the storage array. Remove the SCSI cables in the order that is shown in [Figure 3–10](#).

FIGURE 3–10 SPARC: Disconnecting the SCSI Cables From a Sun StorEdge D1000 Storage Array




# SPARC: Sun StorEdge Multipack Cabling Diagrams

## SPARC: Installing a Sun StorEdge Multipack Storage Array

Connect the cables to the storage array, as shown in [Figure 3–11](#).

Ensure that the *entire* SCSI bus length to each storage array conforms to the distances in [Table 1–1](#). This measurement includes the cables to both nodes, as well as the bus length that is internal to each storage array, node, and host adapter. Refer to the documentation that shipped with the storage array for other restrictions about SCSI operation.

FIGURE 3–11 SPARC: Installing a Sun StorEdge Multipack Mirrored-Pair Configuration: An Example




## SPARC: Adding a Sun StorEdge Multipack Storage Array

Connect the single-ended SCSI cable between the node and the storage arrays, as shown in [Figure 3-12](#).


Ensure that the *entire* SCSI bus length to each storage array conforms to the distances in [Table 1-1](#). This measurement includes the cables to both nodes, as well as the bus length that is internal to each storage array, node, and host adapter. Refer to the documentation that shipped with the storage array for other restrictions about SCSI operation.

FIGURE 3-12 SPARC: Adding a Sun StorEdge Multipack Mirrored-Pair Configuration: Part 1



Connect the storage arrays to the host adapters by using single-ended SCSI cables, as shown in [Figure 3-13](#).


FIGURE 3-13 SPARC: Adding a Sun StorEdge Multipack Mirrored-Pair Configuration: Part 2



## SPARC: Replacing a Sun StorEdge Multipack Storage Array

Disconnect the SCSI cables from the storage array. Disconnect the SCSI cable on the SCSI OUT connector, and then disconnect the cable on the SCSI IN connector, as shown in [Figure 3-14](#).


FIGURE 3-14 SPARC: Replacing a Sun StorEdge Multipack Storage Array



## SPARC: Removing a Sun StorEdge Multipack Storage Array

Disconnect the SCSI cables from the storage array. Disconnect the SCSI cables in the order that is shown in [Figure 3-15](#).

FIGURE 3-15 SPARC: Removing a Sun StorEdge Multipack Storage Array



# SPARC: Sun StorEdge Netra D130/S1 Cabling Diagrams


## SPARC: Installing a Sun StorEdge Netra D130/S1 Storage Array

If your host has only one SCSI port, see “SPARC: Single-SCSI-Port Hosts” on page 86. If your host has two SCSI ports, see “SPARC: Dual-SCSI-Port Hosts” on page 88.

### SPARC: Single-SCSI-Port Hosts

When you install the storage array on single-SCSI-port hosts use the Netra E1 PCI Expander for the second host SCSI port. Figure 3–16 shows an overview of the installation. The storage arrays are cabled so that no single point of failure is in the cluster. Netra E1 PCI Expanders provide the second SCSI port for the 1RU form-factor hosts such as the Netra t1, x1, or t1 200.

FIGURE 3–16 SPARC: Sun StorEdge Netra D130/S1 Mirrored-Pair Configuration With Single SCSI Port Hosts: An Example



To cable the storage arrays correctly, adhere to the bus-length restrictions that are listed in the documentation that shipped with your host adapter. In an Oracle Solaris Cluster environment,

this cable-length measurement includes the cables to both nodes. This cable-length measurement also includes the bus length that is internal to each storage array, node, and host adapter.

FIGURE 3-17 SPARC: Cabling a Sun StorEdge Netra D130/S1 Mirrored-Pair Configuration With Single-SCSI-Port Hosts






FIGURE 3-18 SPARC: Cabling a Sun StorEdge Netra D130/S1 Mirrored-Pair Configuration With Single-SCSI-Port Hosts and Netra E1 PCI Expanders



## SPARC: Dual-SCSI-Port Hosts

To cable the storage arrays correctly, adhere to the bus-length restrictions that are listed in the documentation that shipped with your host adapter. In an Oracle Solaris Cluster environment, this cable-length measurement includes the cables to both nodes. This cable-length measurement also includes the bus length that is internal to each storage array, node, and host adapter.

FIGURE 3-19 SPARC: Cabling a Sun StorEdge Netra D130/S1 Mirrored-Pair Configuration With Dual-SCSI-Port Hosts



## SPARC: Adding a Sun StorEdge Netra D130/S1 Storage Array

To cable the storage arrays correctly, adhere to the bus-length restrictions that are listed in the documentation that shipped with your host adapter. In an Oracle Solaris Cluster environment, this cable-length measurement includes the cables to both nodes. This cable-length measurement also includes the bus length that is internal to each storage array, node, and host adapter.

FIGURE 3–20 SPARC: Adding a Sun StorEdge Netra D130/S1 Storage Array Mirrored-Pair Configuration: Part 1






FIGURE 3–21 SPARC: Adding a Sun StorEdge Netra D130/S1 Mirrored-Pair Configuration: Part 2



## SPARC: Replacing a Sun StorEdge Netra D130/S1 Storage Array

Disconnect the SCSI cables from the storage arrays. Disconnect the cable on the SCSI OUT connector, and then disconnect the cable on the SCSI IN connector, as shown in [Figure 3–22](#).


FIGURE 3–22 SPARC: Replacing a Sun StorEdge Netra D130/S1 Storage Array



## SPARC: Removing a Sun StorEdge Netra D130/S1 Storage Array

Disconnect the SCSI cables from the storage arrays. Disconnect the SCSI cable in the order that is shown in [Figure 3–23](#).

FIGURE 3–23 SPARC: Removing a Sun StorEdge Netra D130/S1 Storage Array




## SPARC: Sun StorEdge D2 Cabling Diagrams

### SPARC: Installing a Sun StorEdge D2 Storage Array

You can install a single storage array by using a split-bus configuration. You can install two storage arrays by using a single-bus configuration.


FIGURE 3–24 SPARC: Installing a Sun StorEdge D2 Split-Bus Configuration: An Example



To use a single storage array in a Oracle Solaris Cluster environment, you must use the split-bus configuration.

As [Figure 3–24](#) shows, the split-bus configuration requires that the storage array contain two Environmental Service Modules (ESM). Each ESM is connected to two nodes. Each ESM controls half the disks in the storage array and the data is mirrored by using a volume manager.

FIGURE 3-25 SPARC: Installing a Sun StorEdge D2 Single-Bus Configuration: An Example



As [Figure 3-25](#) shows, the single-bus configuration requires two storage arrays that each contain a single Environmental Service Module (ESM). The ESM controls all disks on the storage array in which it is installed and the volume manager implements mirroring between the two storage arrays. Each ESM must be connected to two nodes.

## SPARC: Adding a Sun StorEdge D2 Storage Array

Ensure that the *entire* bus length that is connected to the storage array is less than 12 m. The bus length includes the cables from the ESM to the host adapter. See the documentation that shipped with your hardware for internal bus length.

- For split-bus configurations, see [Figure 3-26](#) and [Figure 3-28](#).
- For single-bus configurations, see [Figure 3-27](#) and [Figure 3-29](#).

FIGURE 3-26 SPARC: Adding a Sun StorEdge D2 Split-Bus Configuration: Part 1




FIGURE 3-27 SPARC: Adding a Sun StorEdge D2 Single-Bus Configuration: Part 1




FIGURE 3–28 SPARC: Adding a Sun StorEdge D2 Split-Bus Configuration: Part 2





FIGURE 3–29 SPARC: Adding a Sun StorEdge D2 Single-Bus Configuration: Part 2





# Index

---

## Numbers and Symbols

- 3120 storage array
  - cabling, 75–76
  - FRUs, 37
  - product documentation, 7
- 3310 SCSI storage array
  - cabling, 76–78
  - FRUs, 37
  - product documentation, 7
- 3320 SCSI storage array
  - cabling, 76–78
  - FRUs, 37

## A

- adding
  - See also* installing
  - disk drives, 62–63
  - storage arrays to SPARC clusters, 24–30
  - storage arrays to x86 clusters, 30–34
- air management blank sleds
  - replacing, 37

## B

- backplanes, replacing, 38
- batteries, replacing, 37

## C

- cables, replacing, 47–48
- cabling
  - Sun StorEdge 3120 storage array, 75–76
  - Sun StorEdge 3310 SCSI storage array, 76–78
  - Sun StorEdge 3320 SCSI storage array, 76–78
  - Sun StorEdge D1000 storage array, 78–81
  - Sun StorEdge D2 storage array, 91–95
  - Sun StorEdge Multipack storage array, 82–85
  - Sun StorEdge Netra D130/S1 storage array, 86–91
- chassis, replacing, 41–43
- controller interface boards, replacing, 38
- cooling canisters, replacing, 38

## D

- D1000 storage array
  - cabling, 78–81
  - FRUs, 38
  - product documentation, 7
- D130/S1 storage array
  - cabling, 86–91
  - FRUs, 38
  - product documentation, 7–8
- D2 storage array
  - cabling, 91–95
  - FRUs, 38
  - product documentation, 8
- disk drive, replacing, 37
- disk drives
  - adding, 62–63

disk drives (*Continued*)  
removing, 63–65  
replacing, 65–68, 68–71  
upgrading firmware, 71–72

## E

EMU modules, replacing, 37  
ESMs, replacing, 38

## F

fan canisters, replacing, 38  
fan modules  
replacing, 37  
fans, replacing, 38  
firmware  
upgrading disk drive firmware, 71–72  
upgrading host adapter firmware, 72–73  
FRUs  
Sun StorEdge 3120 storage array, 37  
Sun StorEdge 3310 SCSI storage array, 37  
Sun StorEdge 3320 SCSI storage array, 37  
Sun StorEdge D1000 storage array, 38  
Sun StorEdge D2 storage array, 38  
Sun StorEdge Multipack storage array, 38  
Sun StorEdge Netra D130/S1 storage array, 38

## H

HBAs, replacing, 48–62  
help, 8–9  
host adapters  
replacing, 48–62  
upgrading firmware, 72–73

## I

input power cables  
replacing, 38

installing  
*See also* adding  
disk drives, 62–63  
storage arrays in SPARC clusters, 13–18  
storage arrays in x86 clusters, 18–24  
internal fans, replacing, 38

## L

LEDs, replacing, 37

## M

maintaining, storage arrays, 36–73  
Multipack storage array  
cabling, 82–85  
FRUs, 38  
product documentation, 8

## N

Netra D130/S1 storage array  
cabling, 86–91  
FRUs, 38  
product documentation, 7–8  
Netra st D1000 storage array  
cabling, 78–81  
FRUs, 38  
product documentation, 7  
node  
disconnecting from shared storage, 43–47  
reconnecting to shared storage, 46–47

## O

Oracle Parallel Server, 5  
Oracle Real Application Clusters, 68–71

**P**

power supplies  
replacing, 37, 38

**R**

removing  
disk drives, 63–65  
storage arrays, 39–41

replacing  
air management blank sleds, 37  
backplanes, 38  
batteries, 37  
chassis, 41–43  
controller interface boards, 38  
cooling canisters, 38  
disk drives, 37, 65–68, 68–71  
EMU modules, 37  
ESMs, 38  
fan canisters, 38  
fan modules, 37  
fans, 38  
host adapters, 48–62  
input power cables, 38  
internal fans, 38  
LEDs, 37  
power modules, 37  
power supplies, 37, 38  
reset switch modules, 37  
SCSI cables, 47–48

reset switch modules, replacing, 37

**S**

SCSI cables, replacing, 47–48  
SCSI initiator IDs, 13  
SCSI termination, 43  
setting SCSI initiator IDs, 13  
storage arrays  
adding to SPARC clusters, 24–30  
adding to x86 clusters, 30–34  
installing in SPARC clusters, 13–18  
installing in x86 clusters, 18–24

storage arrays (*Continued*)

maintaining, 36–73  
removing, 39–41

Sun StorEdge 3120 storage array  
cabling, 75–76  
FRUs, 37  
product documentation, 7

Sun StorEdge 3310 SCSI storage array  
cabling, 76–78  
FRUs, 37  
product documentation, 7

Sun StorEdge 3320 SCSI storage array  
cabling, 76–78  
FRUs, 37

Sun StorEdge D1000 storage array  
cabling, 78–81  
FRUs, 38  
product documentation, 7

Sun StorEdge D2 storage array  
cabling, 91–95  
FRUs, 38  
product documentation, 8

Sun StorEdge Multipack storage array  
cabling, 82–85  
FRUs, 38  
product documentation, 8

Sun StorEdge Netra D130/S1 storage array  
cabling, 86–91  
FRUs, 38  
product documentation, 7–8

**T**

technical support, 8–9

**U**

upgrading  
disk drive firmware, 71–72  
host adapter firmware, 72–73

**X**

x86 based systems, 13