

Oracle® Solaris Cluster 3.3 With Sun StorEdge T3 or T3+ Array Manual

SPARC Platform Edition

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

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

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

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

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

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

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

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

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

Contents

| | |
|--|----|
| Preface | 5 |
| 1 Installing and Configuring a Sun StorEdge T3 or T3+ Array | 11 |
| Installing Sun StorEdge T3 and T3+ Arrays | 11 |
| ▼ How to Install a Storage Array in a New Cluster Using a Single-Controller Configuration | 12 |
| ▼ How to Install a Storage Array in a New Cluster, Using a Partner-Group Configuration .. | 16 |
| ▼ How to Add a Storage Array to an Existing Cluster, Using a Single-Controller Configuration | 20 |
| ▼ How to Add a Storage Array to an Existing Cluster, Using a Partner-Group Configuration | 28 |
| Configuring Sun StorEdge T3 and T3+ Arrays | 34 |
| ▼ How to Create a Logical Volume | 34 |
| ▼ How to Remove a Logical Volume | 36 |
| 2 Maintaining and Upgrading a Sun StorEdge T3 or T3+ Array | 41 |
| Maintaining Sun StorEdge T3 or T3+ Storage Array Components | 41 |
| Sun StorEdge T3 and T3+ Array FRUs | 44 |
| ▼ How to Replace a Disk Drive | 44 |
| Replacing a Node-to-Switch Component | 45 |
| ▼ How to Replace a Hub/Switch or Hub/Switch-to-Array Component in a Single-Controller Configuration | 47 |
| ▼ How to Replace an FC Switch or Storage Array-to-Switch Component in a Partner-Group Configuration | 49 |
| ▼ How to Replace a Storage Array Controller | 50 |
| ▼ How to Replace a Chassis | 51 |
| ▼ How to Replace a Host Adapter | 51 |
| ▼ How to Remove a Storage Array in a Single-Controller Configuration | 54 |

- ▼ How to Remove a Partner Group 57
- Upgrading Sun StorEdge T3 or T3+ Storage Arrays 60
 - ▼ How to Upgrade Storage Array Firmware (No Submirrors) 61
 - ▼ How to Upgrade Storage Array Firmware When Using Mirroring 62
 - ▼ How to Upgrade a StorEdge T3 Controller to a StorEdge T3+ Controller 63
 - ▼ How to Migrate From a Single-Controller Configuration to a Partner-Group Configuration 64
- SAN Considerations 67
 - SAN Clustering Considerations 68

- Index**69

Preface

The *Oracle Solaris Cluster 3.3 With Sun StorEdge T3 or T3+ Array Manual* provides procedures specific to Oracle's Sun StorEdge T3 or T3+ arrays in single-controller and partner-group configurations when placed in an Oracle Solaris Cluster environment.

Use this manual with any version of Oracle Solaris Cluster 3.3 software. Unless otherwise noted, procedures are the same for all Oracle Solaris Cluster versions. See the [“Revision History” on page 6](#) for a list of changes to this manual.

Who Should Use This Book

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

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

How This Book Is Organized

This book contains the following chapters.

[Chapter 1, “Installing and Configuring a Sun StorEdge T3 or T3+ Array,”](#) discusses how to install and configure the StorEdge T3 and T3+ storage array.

[Chapter 2, “Maintaining and Upgrading a Sun StorEdge T3 or T3+ Array,”](#) discusses how to maintain the StorEdge T3 and T3+ storage array.

Revision History

The following table lists the information that has been revised or added since the initial release of this documentation. The table also lists the revision date for these changes.

TABLE P-1 Oracle Solaris Cluster 3.3 With Sun StorEdge T3 or T3+ Array Manual

| Revision Date | New Information |
|----------------|---|
| March 2008 | Changed the procedure in “ How to Add a Storage Array to an Existing Cluster, Using a Partner-Group Configuration ” on page 28 so that you do not perform a reconfiguration boot. |
| January 2009 | Updates links to Sun Cluster documentation in Preface. |
| September 2010 | Updated book with new product name and removed old CLI commands for 3.3 release. |

Related Documentation

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

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

TABLE P-2 Hardware Documentation

| Title | Part Number |
|--|-------------|
| <i>Sun StorEdge T3 and T3+ Array Configuration Guide</i> | 816-4771 |
| Available on http://www.sun.com/products-n-solutions/hardware/docs | |
| <i>Sun StorEdge T3 and T3+ Array Installation, Operation, and Service Manual</i> | 816-0773 |
| Available on http://www.sun.com/products-n-solutions/hardware/docs | |
| <i>Sun StorEdge T3 and T3+ Array Release Notes</i> | 816-3652 |
| Available on http://www.sun.com/products-n-solutions/hardware/docs | |
| <i>Sun StorEdge T3 Array Controller Upgrade Manual</i> | 816-0780 |
| Available on http://www.sun.com/products-n-solutions/hardware/docs | |

TABLE P-3 Oracle Solaris Cluster and Sun Cluster Documentation

Documentation

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

Using UNIX Commands

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

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

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

Getting Help

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

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

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

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

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

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

Documentation, Support, and Training

See the following web sites for additional resources:

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

Oracle Welcomes Your Comments

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

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

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

Typographic Conventions

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

TABLE P-4 Typographic Conventions

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

Shell Prompts in Command Examples

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

TABLE P-5 Shell Prompts

| Shell | Prompt |
|--|----------------------------|
| Bash shell, Korn shell, and Bourne shell | \$ |
| Bash shell, Korn shell, and Bourne shell for superuser | # |
| C shell | <code>machine_name%</code> |
| C shell for superuser | <code>machine_name#</code> |

Installing and Configuring a Sun StorEdge T3 or T3+ Array

This chapter contains the procedures about how to install and configure Sun StorEdge T3 and Sun StorEdge T3+ arrays in a single-controller (noninterconnected) or partner-group (interconnected) configuration in an Oracle Solaris Cluster environment.

This chapter contains the following topics.

- “Installing Sun StorEdge T3 and T3+ Arrays” on page 11
- “Configuring Sun StorEdge T3 and T3+ Arrays” on page 34

For information about how to use a storage array in a storage area network (SAN), see “SAN Solutions in an Oracle Solaris Cluster Environment” in *Oracle Solaris Cluster 3.3 Hardware Administration Manual*.

Installing Sun StorEdge T3 and T3+ Arrays

This section contains the procedures about how to install storage arrays in a cluster. The following table lists these procedures.

TABLE 1-1 Task Map: Installing a Storage Array

| Task | Information |
|--|---|
| Installing arrays in a new cluster, using a single-controller configuration | “How to Install a Storage Array in a New Cluster Using a Single-Controller Configuration” on page 12 |
| Installing arrays in a new cluster, using a partner-group configuration | “How to Install a Storage Array in a New Cluster, Using a Partner-Group Configuration” on page 16 |
| Adding arrays to an existing cluster, using a single-controller configuration. | “How to Add a Storage Array to an Existing Cluster, Using a Single-Controller Configuration” on page 20 |

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

| Task | Information |
|---|---|
| Adding arrays to an existing cluster, using a partner-group configuration. | “How to Add a Storage Array to an Existing Cluster, Using a Partner-Group Configuration” on page 28 |
| Upgrading a T3 storage array to a T3+ array. | “How to Upgrade a StorEdge T3 Controller to a StorEdge T3+ Controller” on page 63 |
| Migrating a single-controller array configuration to a partner-group configuration. | “How to Migrate From a Single-Controller Configuration to a Partner-Group Configuration” on page 64 |

▼ How to Install a Storage Array in a New Cluster Using a Single-Controller Configuration

Use this procedure to install and configure the first storage array in a new cluster, using a single-controller configuration. Perform the steps in this procedure in conjunction with the procedures in the Oracle Solaris Cluster software installation documentation and your server hardware manual.

The following procedures contain instructions for other array-installation situations:

- [“How to Install a Storage Array in a New Cluster, Using a Partner-Group Configuration” on page 16](#)
- [“How to Add a Storage Array to an Existing Cluster, Using a Single-Controller Configuration” on page 20](#)
- [“How to Add a Storage Array to an Existing Cluster, Using a Partner-Group Configuration” on page 28](#)

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

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

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

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

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

2 Install the FC hubs/switches.

To install FC hubs/switches, see the documentation that shipped with your FC hub/switch hardware.

Note – If you are using two FC switches and Sun SAN software to create a storage area network (SAN), see “SAN Solutions in an Oracle Solaris Cluster Environment” in *Oracle Solaris Cluster 3.3 Hardware Administration Manual* for more information.

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

This RARP server enables you to assign an IP address to the new storage array by using each storage array’s unique MAC address.

To set up a RARP server, see the *Sun StorEdge T3 and T3+ Array Installation, Operation, and Service Manual*.

4 If you are not adding a StorEdge T3+ array, install the media interface adapters (MIAs) in the storage array that you are installing, as shown in Figure 1–1.

To install a media interface adapter (MIA), see the *Sun StorEdge T3 and T3+ Array Configuration Guide*.

5 If you are adding a StorEdge T3+ array, and if it is necessary, install gigabit interface converters (GBICs) or Small Form-Factor Pluggables (SFPs) in the FC hubs/switches, as shown in Figure 1–1.

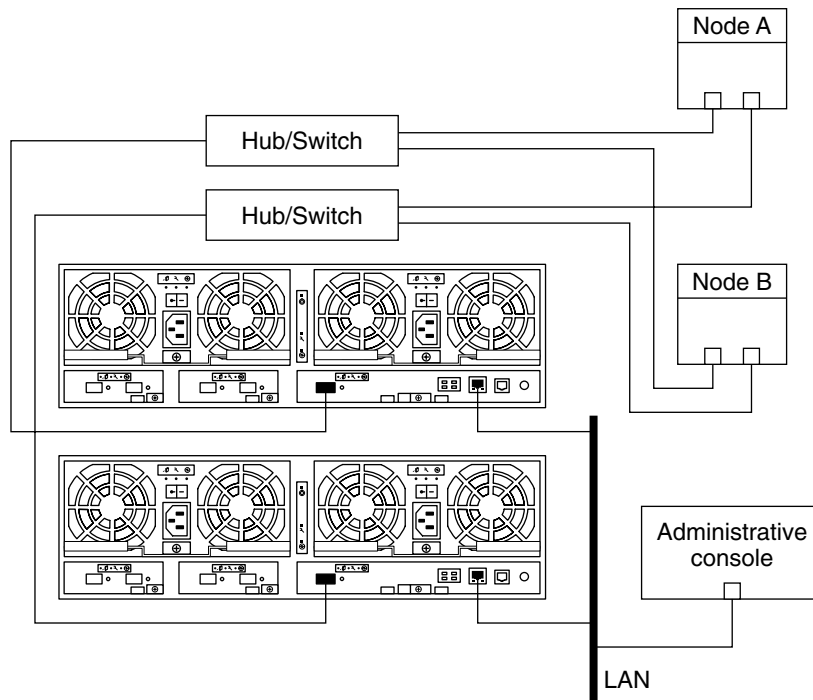
The GBICs or SFPs let you connect the FC hubs/switches to the storage array that you are installing. To install an FC hub/switch GBIC or an SFP, see the documentation that shipped with your FC hub/switch hardware.

6 Install fiber-optic cables between the FC hubs/switches and the storage array, as shown in Figure 1–1.

To install a fiber-optic cable, see the *Sun StorEdge T3 and T3+ Array Configuration Guide*.

7 Install fiber-optic cables between the FC hubs/switches and the nodes, as shown in Figure 1–1.**8 Install the Ethernet cables between the storage array and the Local Area Network (LAN), as shown in Figure 1–1.**

FIGURE 1-1 Installing a Single-Controller Configuration



Note – Figure 1-1 shows how to cable two storage arrays to enable data sharing and host-based mirroring. This configuration prevents a single-point of failure.

- 9 Install power cords to each storage array that you are installing.
- 10 Power on the storage array and confirm that all components are powered on and functional.

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

To power on a storage array, see the *Sun StorEdge T3 and T3+ Array Installation, Operation, and Service Manual*.

- 11 (Optional) Configure the storage array with logical volumes.

To configure the storage array with logical volumes, see the *Sun StorEdge T3 and T3+ Array Administrator's Guide*.

12 Access each storage array that you are adding. Install the required controller firmware for the 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*.

13 Ensure that this new storage array has a unique target address.

To verify and assign a target address, see the *Sun StorEdge T3 and T3+ Array Configuration Guide*.

14 Reset the storage array.

To reboot or reset a storage array, see the *Sun StorEdge T3 and T3+ Array Installation, Operation, and Service Manual*.

15 Install to the nodes the Oracle Solaris operating environment. Apply any required Oracle Solaris patches for Oracle Solaris Cluster software and storage array support.

To install the Oracle Solaris operating environment, see your Oracle Solaris Cluster software installation documentation. For the location of required Oracle Solaris patches and installation instructions for Oracle Solaris Cluster software support, see your Oracle Solaris Cluster release notes documentation. For a list of required Oracle Solaris patches for storage array support, see the *Sun StorEdge T3 and T3+ Array Release Notes*.

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

▼ How to Install a Storage Array in a New Cluster, Using a Partner-Group Configuration

Use this procedure to install and configure the first storage array partner groups in a new cluster. Perform the steps in this procedure in conjunction with the procedures in the Oracle Solaris Cluster software installation documentation and your server hardware manual.

Make certain that you are using the correct procedure. This procedure contains instructions about how to install a partner group into a new cluster, before the cluster is operational. The following procedures contain instructions for other array-installation situations:

- [“How to Add a Storage Array to an Existing Cluster, Using a Partner-Group Configuration” on page 28](#)
- [“How to Install a Storage Array in a New Cluster Using a Single-Controller Configuration” on page 12](#)
- [“How to Add a Storage Array to an Existing Cluster, Using a Single-Controller Configuration” on page 20](#)

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

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

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

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

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

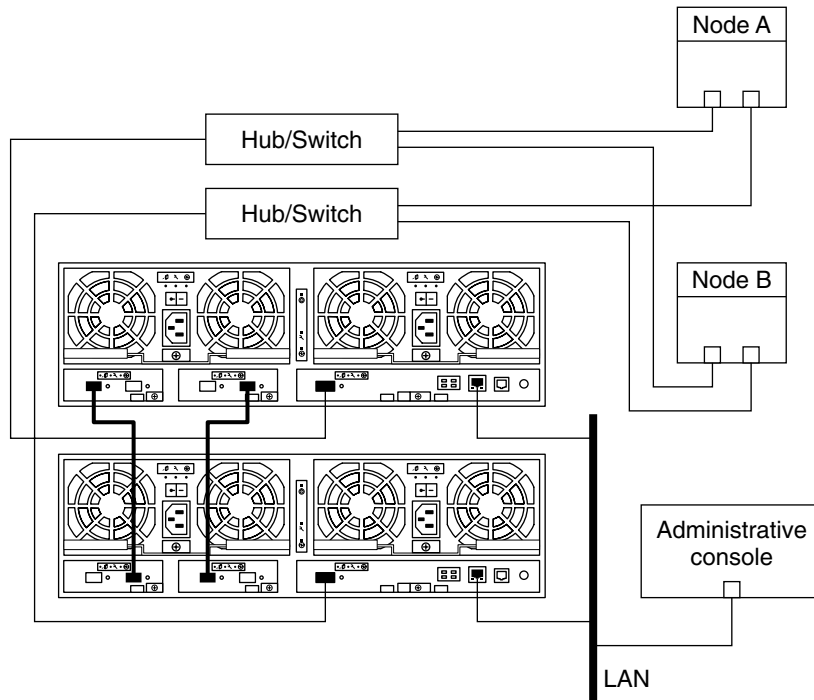
2 Install the Fibre Channel (FC) switches.

To install an FC switch, see the documentation that shipped with your switch hardware.

Note – If you are using two FC switches and Sun SAN software to create a storage area network (SAN), see [“SAN Considerations” on page 67](#) for more information.

3 If you are installing Sun StorEdge T3+ arrays, install the media interface adapters (MIAs) in the Sun StorEdge T3 arrays that you are installing, as shown in [Figure 1–2](#).

FIGURE 1-2 Installing a Partner-Group Configuration



To install a media interface adapter (MIA), see the *Sun StorEdge T3 and T3+ Array Configuration Guide*.

- 4 **If necessary, install GBICs or SFPs in the FC switches, as shown in Figure 1-2.**
To install a GBIC or an SFP to an FC switch, see the documentation that shipped with your FC switch hardware.
- 5 **Set up a Reverse Address Resolution Protocol (RARP) server on the network on which the new storage arrays are to reside.**
This RARP server enables you to assign an IP address to the new storage arrays. Assign an IP address by using the storage array's unique MAC address. For the procedure about how to set up a RARP server, see the *Sun StorEdge T3 and T3+ Array Installation, Operation, and Service Manual*.
- 6 **Cable the storage arrays, as shown in Figure 1-2.**
 - a. **Connect the storage arrays to the FC switches by using fiber-optic cables.**
 - b. **Connect the Ethernet cables from each storage array to the LAN.**

c. Connect interconnect cables between the two storage arrays of each partner group.

d. Connect power cords to each storage array.

To install fiber-optic, Ethernet, and interconnect cables, see the *Sun StorEdge T3 and T3+ Array Installation, Operation, and Service Manual*.

7 Power on the storage arrays. Verify that all components are powered on and functional.

To power on the storage arrays and verify the hardware configuration, see the *Sun StorEdge T3 and T3+ Array Installation, Operation, and Service Manual*.

8 Administer the storage arrays' network settings.

Use the `telnet` command to access the *master* controller unit and administer the storage arrays. To administer the storage array network addresses and settings, see the *Sun StorEdge T3 and T3+ Array Installation, Operation, and Service Manual*.

The master controller unit is the storage array that has the interconnect cables attached to the second port of each interconnect card, when viewed from the rear of the storage arrays. For example, [Figure 1–2](#) shows the master controller unit of the partner group as the lower storage array. In this diagram, the interconnect cables are connected to the second port of each interconnect card on the master controller unit.

9 Install any required storage array controller firmware.

For partner-group configurations, use the `telnet` command to access the *master* controller unit. Install the required controller firmware.

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

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

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

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

- 10 Ensure that each storage array has a unique target address.**
To verify and assign a target address to a storage array, see the *Sun StorEdge T3 and T3+ Array Configuration Guide*.
 - 11 Ensure that the cache and mirror settings for each storage array are set to auto.**
For more information, see the *Sun StorEdge T3 and T3+ Array Administrator's Guide*.
 - 12 Ensure that the mp_support parameter for each storage array is set to mpzio.**
For more information, see the *Sun StorEdge T3 and T3+ Array Administrator's Guide*.
 - 13 Ensure that both storage array controllers are online.**
For more information about how to correct the situation if both controllers are not online, see the *Sun StorEdge T3 and T3+ Array Administrator's Guide*.
 - 14 (Optional) Configure the storage arrays with the desired logical volumes.**
To create and initialize a logical volume, see the *Sun StorEdge T3 and T3+ Array Administrator's Guide*. To mount a logical volume, see the *Sun StorEdge T3 and T3+ Array Installation, Operation, and Service Manual*.
 - 15 Reset the storage arrays.**
To reboot or reset a storage array, see the *Sun StorEdge T3 and T3+ Array Installation, Operation, and Service Manual*.
 - 16 On all nodes, install the Oracle Solaris operating environment. Apply the required Oracle Solaris patches for Oracle Solaris Cluster software and storage array support.**
To install the Oracle Solaris operating environment, see “[How to Install Solaris Software](#)” in *Oracle Solaris Cluster Software Installation Guide*
 - 17 On all nodes, update the /devices and /dev entries.**

```
# devfsadm -C
```
 - 18 On all nodes, confirm that all storage arrays that you installed are visible.**

```
# luxadm display
```
- See Also** To continue with Oracle Solaris Cluster software installation tasks, see your Oracle Solaris Cluster software installation documentation.

▼ How to Add a Storage Array to an Existing Cluster, Using a Single-Controller Configuration

This procedure contains instructions about how to add a new storage array to a running cluster in a single-controller configuration. The following procedures contain instructions for other array-installation situations:

- “How to Add a Storage Array to an Existing Cluster, Using a Partner-Group Configuration” on page 28
- “How to Install a Storage Array in a New Cluster Using a Single-Controller Configuration” on page 12
- “How to Install a Storage Array in a New Cluster, Using a Partner-Group Configuration” on page 16

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

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` role-based access control (RBAC) authorization.**
- 2 Set up a Reverse Address Resolution Protocol (RARP) server on the network on which you want the new storage arrays to be located.**
- 3 Assign an IP address to the new storage arrays.**

This RARP server enables you to assign an IP address to the new storage array by using the storage array’s unique MAC address.

To set up a RARP server, see the *Sun StorEdge T3 and T3+ Array Installation, Operation, and Service Manual*.

- 4 If you are adding a StorEdge T3+ array, install the media interface adapter (MIA) in storage array that you are adding, as shown in [Figure 1–3](#).**

If you are not adding a StorEdge T3+ array, skip this step.

To install a media interface adapter (MIA), see the *Sun StorEdge T3 and T3+ Array Configuration Guide*.

- 5 If necessary, install gigabit interface converters (GBICs) or Small Form-Factor Pluggables (SFPs) in the FC hub/switch, as shown in [Figure 1–3](#).**

The GBICs or SFPs enables you to connect the FC hubs/switches to the storage array that you are adding.

To install an FC hub/switch GBIC or an SFP, see the documentation that shipped with your FC hub/switch hardware.

Note – If you are using two FC switches and Sun SAN software to create a storage area network (SAN), see “SAN Solutions in an Oracle Solaris Cluster Environment” in *Oracle Solaris Cluster 3.3 Hardware Administration Manual* for more information.

- 6 Install the Ethernet cable between the storage array and the Local Area Network (LAN), as shown in Figure 1–3.**

- 7 Power on the storage array array.**

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

To power on a storage array, see the *Sun StorEdge T3 and T3+ Array Installation, Operation, and Service Manual*.

- 8 Access the storage array that you are adding. If necessary, install the required controller firmware for the 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*.

- 9 If this new storage array does not yet have a unique target address, change the target address for this new storage array.**


If the target address for this array is already unique, skip this step.

To verify and assign a target address, see the *Sun StorEdge T3 and T3+ Array Configuration Guide*.

- 10 Install a fiber-optic cable between the FC hub/switch and the storage array, as shown in Figure 1–3.**

To install a fiber-optic cable, see the *Sun StorEdge T3 and T3+ Array Configuration Guide*.

FIGURE 1–3 Adding a Single-Controller Configuration: Part I



Note – Figure 1–3 shows how to cable two storage arrays to enable data sharing and host-based mirroring. This configuration prevents a single-point of failure.

- 11 Configure the new storage array.**

To create a logical volume, see the *Sun StorEdge T3 and T3+ Array Administrator's Guide*.

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

Record this information because you use this information in [Step 40](#) and [Step 41](#) of this procedure to return resource groups and device groups to these nodes.

Use the following command:

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

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

```
# clnode evacuate
```

- 14 If you need to install a host adapter in Node A, and if it is the first FC host adapter on Node A, determine whether the Fibre Channel support packages are already installed on these nodes.**

This product requires the following packages.

```
# pkginfo | egrep wLux
system SUNWluxd Sun Enterprise Network Array sf Device Driver
system SUNWluxdx Sun Enterprise Network Array sf Device Driver
(64-bit)
system SUNWluxl Sun Enterprise Network Array social Device Driver
system SUNWluxlx Sun Enterprise Network Array social Device Driver
(64-bit)
system SUNWluxop Sun Enterprise Network Array firmware and utilities
```

If this is not the first FC host adapter on Node A, skip to [Step 16](#). If you do not need to install a host adapter in Node A, skip to [Step 35](#).

- 15 If the Fibre Channel support packages are not installed, install the required support packages that are missing.**

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

- 16 Shut down Node A.**

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

- 17 Power off Node A.**

- 18 Install the host adapter in Node A.**

To install a host adapter, see the documentation that shipped with your host adapter and node.

- 19 If necessary, power on and boot Node A into noncluster mode by adding -x to your boot instruction.**

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

- 20 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.

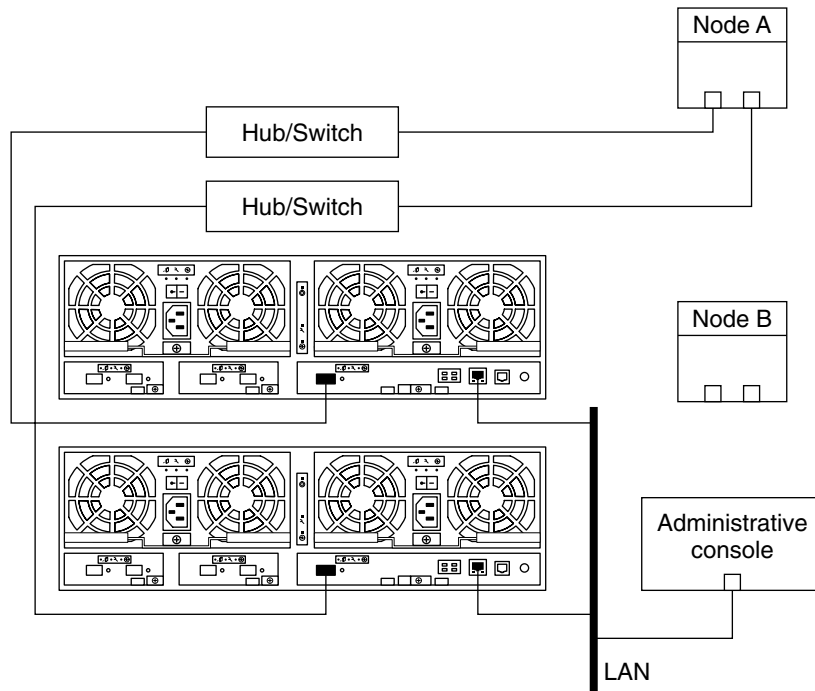
If you must apply a patch when a node is in noncluster mode, you can apply it in a rolling fashion, one node at a time, unless instructions for a patch require that you shut down the entire cluster. Follow the procedures in “[How to Apply a Rebooting Patch \(Node\)](#)” in *Oracle Solaris Cluster System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

21 Connect a fiber-optic cable between the FC hub/switch and Node A, as shown in [Figure 1–4](#).

To install an FC host adapter GBIC or an SFP, see your host adapter documentation. To install a fiber-optic cable, see the *Sun StorEdge T3 and T3+ Array Configuration Guide*.

FIGURE 1–4 Adding a Single-Controller Configuration: Part II



22 If necessary, install the required Oracle Solaris patches for array support on Node A.

For a list of required Oracle Solaris patches for storage array support, see the *Sun StorEdge T3 and T3+ Array Release Notes*.

23 Shut down Node A.

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

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

```
# boot -r
```

25 Label the new logical volume.

To label a logical volume, see the *Sun StorEdge T3 and T3+ Array Administrator's Guide*.

26 (Optional) On Node A, verify that the device IDs (DIDs) are assigned to the new LUNs.

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

27 If you need to install a host adapter in Node B, and if the host adapter that you are installing the first FC host adapter on Node B, determine whether the Fibre Channel support packages are already installed on these nodes.

This product requires the following packages.

```
# pkginfo | egrep Wlux
system SUNWluxd Sun Enterprise Network Array sf Device Driver
system SUNWluxdx Sun Enterprise Network Array sf Device Driver
(64-bit)
system SUNWluxl Sun Enterprise Network Array social Device Driver
system SUNWluxlx Sun Enterprise Network Array social Device Driver
(64-bit)
system SUNWluxop Sun Enterprise Network Array firmware and utilities
```

If this is not the first FC host adapter on Node B, skip to [Step 29](#). If you do not need to install a host adapter, skip to [Step 34](#).

28 If the Fibre Channel support packages are not installed, install the required support packages that are missing.

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

29 Shut down Node B.

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

30 Power off Node B.

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

31 Install the host adapter in Node B.

To install a host adapter, see the documentation that shipped with your host adapter and node.

32 If necessary, power on and boot Node B into noncluster mode by adding -x to your boot instruction.

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

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

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

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

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

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

34 If necessary, install a GBIC or an SFP, as shown in [Figure 1–5](#).


To install an FC hub/switch GBIC or an SFP, see the documentation that shipped with your FC hub/switch hardware.

Note – If you are using two FC switches and Sun SAN software to create a storage area network (SAN), see “[SAN Solutions in an Oracle Solaris Cluster Environment](#)” in *Oracle Solaris Cluster 3.3 Hardware Administration Manual* for more information.

35 If necessary, connect a fiber-optic cable between the FC hub/switch and Node B, as shown in [Figure 1–5](#).

To install an FC host adapter GBIC or an SFP, see your host adapter documentation. To install a fiber-optic cable, see the *Sun StorEdge T3 and T3+ Array Configuration Guide*.

FIGURE 1-5 Adding a Single-Controller Configuration: Part III



- 36 If necessary, install the required Oracle Solaris patches for storage array support on Node B.**
For a list of required Oracle Solaris patches for storage array support, see the *Sun StorEdge T3 and T3+ Array Release Notes*.
- 37 Shut down Node B.**
To shut down and power off a node, see your Oracle Solaris Cluster system administration documentation.
- 38 Perform a reconfiguration boot to create the new Oracle Solaris device files and links on Node B.**
`# boot -r`
- 39 (Optional) On Node B, verify that the device IDs (DIDs) are assigned to the new LUNs.**
`# cldevice list -n NodeB -v`
- 40 (Optional) Restore the device groups to the original node.**
Do the following for each device group that you want to return to the original node.
`# cldevicegroup switch -n nodename devicegroup1[devicegroup2 ...]`
`-n nodename` The node to which you are restoring device groups.

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

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

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

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

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

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

42 Perform volume management administration to incorporate the new logical volumes into the cluster.

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

▼ How to Add a Storage Array to an Existing Cluster, Using a Partner-Group Configuration

This procedure contains instructions for adding new storage array partner groups to a running cluster. The following procedures contain instructions for other array-installation situations:

- [“How to Install a Storage Array in a New Cluster, Using a Partner-Group Configuration” on page 16](#)
- [“How to Install a Storage Array in a New Cluster Using a Single-Controller Configuration” on page 12](#)
- [“How to Add a Storage Array to an Existing Cluster, Using a Single-Controller Configuration” on page 20](#)

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 **Set up a Reverse Address Resolution Protocol (RARP) server on the network you want the new storage arrays to reside on. Afterward, assign an IP address to the new storage arrays.**


Note – Assign an IP address to the *master* controller unit only. The master controller unit is the storage array that has the interconnect cables attached to the second port of each interconnect card, as shown in [Figure 1-6](#).

This RARP server enables you to assign an IP address to the new storage arrays. Assign an IP address by using the storage array's unique MAC address. To set up a RARP server, see the *Sun StorEdge T3 and T3+ Array Installation, Operation, and Service Manual*.

- 3 Install the Ethernet cable between the storage arrays and the local area network (LAN), as shown in [Figure 1-6](#).
- 4 If they are not already installed, install interconnect cables between the two storage arrays of each partner group, as shown in [Figure 1-6](#).

To install interconnect cables, see the *Sun StorEdge T3 and T3+ Array Installation, Operation, and Service Manual*.

FIGURE 1-6 Adding a Partner-Group Configuration: Part I



- 5 Power on the storage arrays.

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

To power on storage arrays, see the *Sun StorEdge T3 and T3+ Array Installation, Operation, and Service Manual*.

6 Administer the storage arrays' network addresses and settings.

Use the `telnet` command to access the *master* controller unit and to administer the storage arrays.

To administer the network address and the settings of a storage array, see the *Sun StorEdge T3 and T3+ Array Installation, Operation, and Service Manual*.

7 Install any required storage array controller firmware upgrades.

For partner-group configurations, use the `telnet` command to the *master* controller unit. If necessary, install the required controller firmware for the storage array.

For the required revision number of the storage array controller firmware, see the *Sun StorEdge T3 Disk Tray Release Notes*.

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

To verify and assign a target address to a storage array, see the *Sun StorEdge T3 and T3+ Array Configuration Guide*.

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

For more information, see the *Sun StorEdge T3 and T3+ Array Administrator's Guide*.

10 Ensure that the `mp_support` parameter for each storage array is set to `mpxio`.

For more information, see the *Sun StorEdge T3 and T3+ Array Administrator's Guide*.

11 Configure the new storage arrays with the desired logical volumes.

To create and initialize a logical volume, see the *Sun StorEdge T3 and T3+ Array Administrator's Guide*. To mount a logical volume, see the *Sun StorEdge T3 and T3+ Array Installation, Operation, and Service Manual*.

12 Reset the storage arrays.

To reboot or reset a storage array, see the *Sun StorEdge T3 and T3+ Array Installation, Operation, and Service Manual*.

13 If you are adding Sun StorEdge T3+ arrays, install the media interface adapter (MIA) in the Sun StorEdge T3 arrays that you are adding, as shown in [Figure 1–6](#).

To install an MIA, see the *Sun StorEdge T3 and T3+ Array Configuration Guide*.

14 If necessary, install GBICs or SFPs in the FC switches, as shown in Figure 1–6.

To install a GBIC or an SFP to an FC switch, see the documentation that shipped with your FC switch hardware.

15 Install a fiber-optic cable between each FC switch and both new storage arrays of the partner group, as shown in Figure 1–6.

To install a fiber-optic cable, see the *Sun StorEdge T3 and T3+ Array Installation, Operation, and Service Manual*.

Note – If you are using two FC switches and Sun SAN software to create a storage area network (SAN), see “SAN Considerations” on page 67 for more information.

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

Record this information because you use it in [Step 30](#) of this procedure to return resource groups and device groups to these nodes.

Use the following command:

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

17 Move all resource groups and device groups off each node in the cluster.

```
# clnode evacuate
```

18 If you need to install host adapters in the node, and if the host adapter you are installing is the first adapter on the node, determine whether the required support packages are already installed on this node.

The following packages are required.

```
# pkginfo | egrep WLux
system SUNWluxd Sun Enterprise Network Array sf Device Driver
system SUNWluxdx Sun Enterprise Network Array sf Device Driver
(64-bit)
system SUNWluxl Sun Enterprise Network Array social Device Driver
system SUNWluxlx Sun Enterprise Network Array social Device Driver
(64-bit)
system SUNWluxop Sun Enterprise Network Array firmware and utilities
system SUNWluxox Sun Enterprise Network Array libraries (64-bit)
```

If this is not the first host adapter on the node, skip to [Step 20](#).

19 If the required support packages are not present, install them.

The support packages are located in the Product directory of the Oracle Solaris DVD. Add any missing packages.

- 20 If you need to install host adapters in the node, shut down and power off the node, and then install them in the node.**

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

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

- 21 If you installed host adapters in the node, power on and boot the node into noncluster mode.**

To boot nodes, see [Chapter 3, “Shutting Down and Booting a Cluster,”](#) in *Oracle Solaris Cluster System Administration Guide*.

- 22 If you installed host adapters in the node, and if necessary, upgrade the host adapter firmware on the node.**

- 23 If necessary, install the required Oracle Solaris patches for storage array support on the node.**

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*.


- 24 If you installed host adapters in the node, reboot the node in cluster mode.**

- 25 Connect fiber-optic cables between the node and the FC switches, as shown in [Figure 1–7](#).**

To install a fiber-optic cable, see the *Sun StorEdge T3 and T3+ Array Installation, Operation, and Service Manual*.

Note – If you are using two FC switches and Sun SAN software to create a storage area network (SAN), see [“SAN Considerations”](#) on [page 67](#) for more information.

FIGURE 1-7 Adding a Partner-Group Configuration: Part II



- 26 On the current node, update the `/devices` and `/dev` entries.


```
# devfsadm
```
- 27 From any node in the cluster, update the global device namespace.


```
# cldevice populate
```
- 28 Label the new storage array logical volume.

To label a logical volume, see the *Sun StorEdge T3 and T3+ Array Administrator's Guide*.
- 29 (Optional) On the current node, verify that the device IDs (DIDs) are assigned to the new LUNs.


```
# cldevice list -n CurrentNode -v
```
- 30 (Optional) Restore the resource groups to the original node.

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

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

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

- 31 For each of the other nodes in the cluster, repeat [Step 17](#) through [Step 30](#).
- 32 Perform volume management administration to incorporate the new logical volumes into the cluster.

For more information, see your Veritas Volume Manager documentation.

Configuring Sun StorEdge T3 and T3+ Arrays

This section contains the procedures about how to configure a storage array in a running cluster. The following table lists these procedures.

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

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

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

TABLE 1-2 Task Map: Configuring a Storage Array

| Task | Information |
|--------------------------|---|
| Create a logical volume. | “How to Create a Logical Volume” on page 34 |
| Remove a logical volume. | “How to Remove a Logical Volume” on page 36 |

▼ How to Create a Logical Volume

Use this procedure to create a logical volume from unassigned storage capacity.

Note – Oracle's Sun storage documentation uses the following terms:

- Logical volume
- Logical device
- Logical unit number (LUN)

This manual uses *logical volume* to refer to all such logical constructs.

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

- All nodes are booted in cluster mode and attached to the storage device.
- The storage device is installed and configured. If you are using multipathing, the storage device is configured as described in the installation procedure.
- If you are using Solaris I/O multipathing (MPxIO) for the Oracle Solaris 10 OS, previously called Sun StorEdge Traffic Manager in the Solaris 9 OS, verify that the paths to the storage device are functioning. To configure multipathing, see the *Solaris Fibre Channel Storage Configuration and Multipathing Support Guide*.

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

- 1 Become superuser or assume a role that provides `soLarIS.cLuster.mOdiFY` role-based access control (RBAC) authorization.**
- 2 Follow the instructions in your storage device's documentation to create and map the logical volume. For a URL to this storage documentation, see [“Related Documentation” on page 6](#).**
 - Completely set up the logical volume. When you are finished, the volume must be created, mapped, mounted, and initialized.
 - If necessary, partition the volume.
 - To allow multiple clusters and nonclustered nodes to access the storage device, create initiator groups by using LUN masking.
- 3 If you are not using multipathing, skip to [Step 5](#).**
- 4 If you are using multipathing, and if any devices that are associated with the volume you created are at an unconfigured state, configure the multipathing paths on each node that is connected to the storage device.**

To determine whether any devices that are associated with the volume you created are at an unconfigured state, use the following command.

```
# cfgadm -al | grep disk
```

Note – To configure the Oracle Solaris I/O multipathing paths on each node that is connected to the storage device, use the following command.

```
# cfmadm -o force_update -c configure controllerinstance
```

To configure multipathing, see the *Solaris Fibre Channel Storage Configuration and Multipathing Support Guide*.

- 5** On one node that is connected to the storage device, use the `format` command to label the new logical volume.
- 6** From any node in the cluster, update the global device namespace.

```
# cldevice populate
```

Note – You might have a volume management daemon such as `vold` running on your node, and have a DVD drive connected to the node. Under these conditions, a device busy error might be returned even if no disk is inserted in the drive. This error is expected behavior. You can safely ignore this error message.

- 7** To manage this volume with volume management software, use Solaris Volume Manager or Veritas Volume Manager commands to update the list of devices on all nodes that are attached to the new volume that you created.

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

- See Also**
- To configure a logical volume as a quorum device, see Chapter 6, “Administering Quorum,” in *Oracle Solaris Cluster System Administration Guide*.
 - To create a new resource or configure a running resource to use the new logical volume, see Chapter 2, “Administering Data Service Resources,” in *Oracle Solaris Cluster Data Services Planning and Administration Guide*.

▼ How to Remove a Logical Volume

Use this procedure to remove a logical volume. This procedure defines Node A as the node with which you begin working.

Note – Sun storage documentation uses the following terms:

- Logical volume
- Logical device
- Logical unit number (LUN)

This manual uses *logical volume* to refer to all such logical constructs.

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

- All nodes are booted in cluster mode and attached to the storage device.
- The logical volume and the path between the nodes and the storage device are both operational.

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

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

2 Identify the logical volume that you are removing.

Refer to your Solaris Volume Manager or Veritas Volume Manager documentation for more information.

3 (Optional) Migrate all data off the logical volume that you are removing. Alternatively, back up that data.

4 If the LUN that you are removing is configured as a quorum device, choose and configure another device as the quorum device. Then remove the old quorum device.

To determine whether the LUN is configured as a quorum device, use the following command.

```
# clquorum show
```

For procedures about how to add and remove quorum devices, see [Chapter 6, “Administering Quorum,”](#) in *Oracle Solaris Cluster System Administration Guide*.

5 If you are using volume management software, use that software to update the list of devices on all nodes that are attached to the logical volume that you are removing.

For instructions about how to update the list of devices, see your Solaris Volume Manager or Veritas Volume Manager documentation.

- 6 If you are using volume management software, run the appropriate Solaris Volume Manager or Veritas Volume Manager commands to remove the logical volume from any diskset or disk group.**

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

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

```
# vxdisk offline Accessname
# vxdisk rm Accessname

Accessname    Disk access name
```

- 7 If you are using multipathing, unconfigure the volume in Solaris I/O multipathing.**

```
# cfgadm -o force_update -c unconfigure Logical_Volume
```

- 8 Access the storage device and remove the logical volume.**

To remove the volume, see your storage documentation. For a list of storage documentation, see [“Related Documentation”](#) on page 6.

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

Record this information because you use it in [Step 14](#) and [Step 15](#) of this procedure to return resource groups and device groups to these nodes.

Use the following command:

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

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

```
# clnode evacuate nodename
```

- 11 Shut down and reboot Node A.**

To shut down and boot a node, see [Chapter 3, “Shutting Down and Booting a Cluster,”](#) in *Oracle Solaris Cluster System Administration Guide*.

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

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

- 13** For each additional node that is connected to the shared storage that hosted the logical volume, repeat [Step 9](#) to [Step 12](#).

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

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

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

-n nodename

The node to which you are restoring device groups.

devicegroup1[devicegroup2 ...]

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

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

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

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

nodename

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

resourcegroup1[resourcegroup2 ...]

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

Maintaining and Upgrading a Sun StorEdge T3 or T3+ Array

This chapter contains the procedures about how to maintain Sun StorEdge T3 and Sun StorEdge T3+ arrays in a single-controller (noninterconnected) or partner-group (interconnected) configuration in an Oracle Solaris Cluster environment.

This book contains the following topics.

- “Maintaining Sun StorEdge T3 or T3+ Storage Array Components” on page 41
- “Upgrading Sun StorEdge T3 or T3+ Storage Arrays” on page 60
- “SAN Considerations” on page 67

For information about how to use a storage array in a storage area network (SAN), see “SAN Solutions in an Oracle Solaris Cluster Environment” in *Oracle Solaris Cluster 3.3 Hardware Administration Manual*.

Maintaining Sun StorEdge T3 or T3+ Storage Array Components

This section contains the procedures about how to maintain a storage array. The following table lists these procedures. This section does not include a procedure about how to add a disk drive and a procedure about how to remove a disk drive, because a storage array only operates when fully configured.



Caution – If you remove any field-replaceable unit (FRU) for an extended period of time, thermal complications might result. To prevent this complication, the storage array is designed so that an orderly shutdown occurs. This shutdown occurs when you remove a component for longer than 30 minutes. Therefore, a replacement part must be immediately available before you start an FRU replacement procedure. You must replace an FRU within 30 minutes. If you do not replace the FRU within this time, the storage array, and all attached storage array, shut down and power off.

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

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

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

TABLE 2-1 Task Map: Maintaining a Storage Array

| Task | Information |
|---|--|
| Replace a disk drive. | “How to Replace a Disk Drive” on page 44 |
| Replace a host-to-hub/switch fiber-optic cable. | “Replacing a Node-to-Switch Component” on page 45 |
| Replace an FC host adapter GBIC or an SFP. | “Replacing a Node-to-Switch Component” on page 45 |
| Replace an FC hub/switch GBIC or an SFP that connects an FC hub/switch to a host. | “Replacing a Node-to-Switch Component” on page 45 |
| Replace a hub/switch-to-storage array fiber-optic cable. | “How to Replace a Hub/Switch or Hub/Switch-to-Array Component in a Single-Controller Configuration” on page 47 |
| Replace an FC hub/switch GBIC or an SFP that connects the FC hub/switch to a storage array. | “How to Replace a Hub/Switch or Hub/Switch-to-Array Component in a Single-Controller Configuration” on page 47 |

TABLE 2-1 Task Map: Maintaining a Storage Array (Continued)

| Task | Information |
|---|--|
| Replace an FC hub/switch. | "How to Replace a Hub/Switch or Hub/Switch-to-Array Component in a Single-Controller Configuration" on page 47 |
| Replace an FC switch. | "How to Replace a Hub/Switch or Hub/Switch-to-Array Component in a Single-Controller Configuration" on page 47 |
| This procedure applies to SAN-configured clusters only. | "How to Replace a Hub/Switch or Hub/Switch-to-Array Component in a Single-Controller Configuration" on page 47 |
| Replace an FC hub/switch power cord. | "How to Replace a Hub/Switch or Hub/Switch-to-Array Component in a Single-Controller Configuration" on page 47 |
| Replace a media interface adapter (MIA) on a storage array. | "How to Replace a Hub/Switch or Hub/Switch-to-Array Component in a Single-Controller Configuration" on page 47 |
| This procedure does not apply to storage arrays. | "How to Replace a Hub/Switch or Hub/Switch-to-Array Component in a Single-Controller Configuration" on page 47 |
| Replace a storage array controller. | "How to Replace a Storage Array Controller" on page 50 |
| Replace a chassis. | "How to Replace a Chassis" on page 51 |
| Replace a host adapter. | "How to Replace a Host Adapter" on page 51 |
| Replace a host adapter. | "How to Replace a Host Adapter" on page 51 |
| Remove a storage array. | "How to Remove a Storage Array in a Single-Controller Configuration" on page 54 |
| Remove a storage array. | "How to Remove a Storage Array in a Single-Controller Configuration" on page 54 |
| Remove a partner group. | "How to Remove a Partner Group" on page 57 |
| Add a node to the storage device. | Oracle Solaris Cluster system administration documentation |
| Remove a node from the storage device. | Oracle Solaris Cluster system administration documentation |

Sun StorEdge T3 and T3+ Array FRUs

The following is a list of administrative tasks that require no cluster-specific procedures. See the *Sun StorEdge T3 and T3+ Array Installation, Operation, and Service Manual* for the following procedures.

- Replacing an Ethernet cable.
- Replacing a power cable on the storage array.
- Replacing a power and cooling unit (PCU).
- Replacing a unit interconnect card (UIC).
- Adding a disk drive.
- Removing a disk drive.

Upgrading a Sun StorEdge T3 array controller to a Sun StorEdge T3+ array controller requires no cluster-specific procedures. See the *Sun StorEdge T3 Array Controller Upgrade Manual* for the following procedures.

▼ How to Replace a Disk Drive

Use this procedure to replace a failed disk drive in a storage array in a running cluster.

Note – Oracle storage documentation uses the following terms:

- Logical volume
- Logical device
- Logical unit number (LUN)

This manual uses *logical volume* to refer to all such logical constructs.

Before You Begin This procedure 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` RBAC authorization.**
- 2 If the failed disk drive affect the storage array logical volume's availability, If yes, use volume manager commands to detach the submirror or plex.**

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

- 3 If the logical volume (in [Step 1](#)) is configured as a quorum device, choose another volume to configure as the quorum device. Then remove the old quorum device.**

To determine whether the LUN is configured as a quorum device, use one of the following commands.

```
# clquorum show
```

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

- 4 Replace the failed disk drive.**

For instructions, refer to the *Sun StorEdge T3 and T3+ Array Installation, Operation, and Service Manual*.

- 5 (Optional) If the new disk drive is part of a logical volume that you want to be a quorum device, add the quorum device.**

To add a quorum device, see your Oracle Solaris Cluster system administration documentation.

- 6 If you detached a submirror or plex in [Step 1](#), use volume manager commands to reattach the submirror or plex.**

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

Replacing a Node-to-Switch Component

Use this procedure to replace a node-to-switch component that has failed or that you suspect might be contributing to a problem.

Note – Node-to-switch components that are covered by this procedure include the following components:

- Node-to-switch fiber-optic cables
- Gigabit interface converters (GBICs) or small form-factor pluggables (SFPs) on an FC switch
- FC switches

To replace a host adapter, see [“How to Replace a Host Adapter” on page 51](#).

This procedure defines Node A as the node that is connected to the node-to-switch component that you are replacing. This procedure assumes that, except for the component you are replacing, your cluster is operational.

Ensure that you are following the appropriate instructions:

- If your cluster uses multipathing, see “[How to Replace a Node-to-Switch Component in a Cluster That Uses Multipathing](#)” on page 46.
- If your cluster does *not* use multipathing, see “[How to Replace a Node-to-Switch Component in a Cluster Without Multipathing](#)” on page 46.

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

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

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

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

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

You have completed this procedure.

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

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

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

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

```
# clnode evacuate nodename
```

5 Replace the node-to-switch component.

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

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

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

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

```
-n nodename
```

The node to which you are restoring device groups.

```
devicegroup1 [devicegroup2 ...]
```

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

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

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

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

```
nodename
```

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

```
resourcegroup1 [resourcegroup2 ...]
```

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

▼ How to Replace a Hub/Switch or Hub/Switch-to-Array Component in a Single-Controller Configuration

Use this procedure to replace a hub/switch or the following hub/switch-to-array components for an array in a single-controller configuration. To replace these components for an array in a partner-group configuration, see [“How to Replace an FC Switch or Storage Array-to-Switch Component in a Partner-Group Configuration”](#) on page 49. You can use storage array in single-controller configuration with FC switches when you create a SAN.

- Fiber-optic cable that connects a hub/switch to a storage array
- FC hub/switch GBIC or an SFP that connects a hub/switch to a storage array
- FC hub/switch
- FC hub/switch power cord
- Media interface adapter (MIA) on a StorEdge T3 array

This component does not apply to StorEdge T3+ arrays.

1 Detach the submirrors on the storage array. This array is connected to the hub/switch-to-array fiber-optic cable that you are replacing. Detach the submirrors to stop all I/O activity to this storage array.

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

2 Replace the hub/switch or hub/switch-to-array component.

- For the procedure about how to replace a fiber-optic cable, see the *Sun StorEdge T3 and T3+ Array Configuration Guide*.
- For the procedure about how to replace an FC hub/switch GBIC or an SFP, an FC hub/switch, or an FC hub/switch power cord, see the documentation that shipped with your FC hub/switch hardware.
- For the procedure about how to replace an MIA, see the *Sun StorEdge T3 and T3+ Array Installation, Operation, and Service Manual*.
- If you are replacing FC switches in a SAN, follow the hardware installation and SAN configuration instructions in the documentation that shipped with your switch hardware.

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

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

3 Reattach the submirrors to resynchronize the submirrors.

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

▼ How to Replace an FC Switch or Storage Array-to-Switch Component in a Partner-Group Configuration

Use this procedure to replace components for a storage array in a partner-group configuration in a running cluster. To replace components for arrays in a single-controller configuration, see “[How to Replace a Hub/Switch or Hub/Switch-to-Array Component in a Single-Controller Configuration](#)” on page 47. Use this procedure to replace the following storage array-to-switch components.

- Fiber-optic cable that connects an FC switch to a storage array.
- GBIC or an SFP on an FC switch, connecting to a storage array.
- FC switch.
- Media Interface Adapter (MIA) on a Sun StorEdge T3 storage array. This component does not apply to Sun StorEdge T3+ storage arrays.
- Interconnect cables between two interconnected storage arrays of a partner group.

1 Access the storage array. This storage array is connected to the FC switch or component that you are replacing.

2 View the controller status for the two storage arrays in the partner group.

For more information, see the *Sun StorEdge T3 and T3+ Array Administrator's Guide*.

3 If you replacing an FC switch or component that is attached to a storage array controller that is ONLINE, disable the controller.

If the controller is already DISABLED, skip to [Step 5](#).

For more information, see the *Sun StorEdge T3 and T3+ Array Administrator's Guide*.

4 Verify that the controller's state has been changed to DISABLED.

5 Replace the component by using the following references.

- For the procedure about how to replace a fiber-optic cable between a storage array and an FC switch, see the documentation that shipped with your switch hardware.
- For the procedure about how to replace a GBIC or an SFP on an FC switch, see the documentation that shipped with your FC switch hardware.
- For the procedure about how to replace an FC switch, see the documentation that shipped with your switch hardware.

Note – If you are replacing an FC switch and you intend to save the switch IP configuration for restoration to the replacement switch, wait to connect the cables to the replacement switch. Connect the cables to the replacement switch *after* you recall the Fabric configuration to the replacement switch. For more information about how to save and recall switch configurations, see the documentation that shipped with your switch hardware.

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

- **For the procedure about how to replace an MIA, see the *Sun StorEdge T3 and T3+ Array Installation, Operation, and Service Manual*.**
 - **For the procedure about how to replace interconnect cables, see the *Sun StorEdge T3 and T3+ Array Installation, Operation, and Service Manual*.**
- 6 **If necessary, access the storage array. This storage array is the storage array in the partner group that is still online.**
 - 7 **Enable the storage array that you disabled in [Step 3](#).**
For more information, see the *Sun StorEdge T3 and T3+ Array Administrator's Guide*.
 - 8 **Verify that the controller's state has been changed to `ONLINE`.**

▼ **How to Replace a Storage Array Controller**

- 1 **Detach the submirrors on the storage array. This array is connected to the controller that you are replacing. Detach the submirrors to stop all I/O activity to this storage array.**
For more information, see your Solaris Volume Manager or Veritas Volume Manager documentation.
- 2 **Replace the controller.**
For the procedure about how to replace a storage array controller, see the *Sun StorEdge T3 and T3+ Array Installation, Operation, and Service Manual*.

3 Reattach the submirrors to resynchronize the submirrors.

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

▼ How to Replace a Chassis

Use this procedure to replace a storage array chassis. This procedure assumes that you are retaining all FRUs other than the chassis and the midplane. To replace the chassis, you must replace both the chassis and the midplane because these components are manufactured as one part.



Caution – You must be an Oracle service provider to perform this procedure. If you need to replace a chassis, contact your Oracle service provider.

1 Detach the submirrors on the storage array. This array is connected to the chassis that you are replacing. Detach the submirrors to stop all I/O activity to this storage array.

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

2 Replace the chassis and the midplane.

For the procedure about how to replace a storage array chassis, see the *Sun StorEdge T3 and T3+ Array Field Service Manual*.

3 Reattach the submirrors to resynchronize the submirrors.



Caution – The world wide names (WWNs) change as a result of this procedure. You must reconfigure your volume manager software to recognize the new WWNs.

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

▼ How to Replace a Host Adapter

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

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

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

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

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

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

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

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

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

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

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

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

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

```
# clnode evacuate nodename
```

4 Shut down Node A.

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

5 Power off Node A.

6 Replace the failed host adapter.

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

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

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

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

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

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

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

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

- 9 Boot Node A into cluster mode.**

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

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

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

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

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

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

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

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

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

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

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

▼ How to Remove a Storage Array in a Single-Controller Configuration

Use this procedure to remove a storage array and its submirrors from a running cluster. This procedure provides the flexibility to remove the host adapters from the nodes for the storage array that you are removing. To remove a partner group from the cluster, see [“How to Remove a Partner Group” on page 57](#).

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

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

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

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



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

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

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

- 1 **Back up all database tables, data services, and volumes that are associated with the storage array. This storage array is the storage array you are removing.**

- 2 Detach the submirrors from the storage array that you are removing. Detach the submirrors to stop all I/O activity to the storage array.**

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

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

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

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

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

Use the following command:

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

- 5 Shut down Node A.**

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

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

If this is not the last storage array, skip this step.

For the procedure about how to remove a fiber-optic cable, see the *Sun StorEdge T3 and T3+ Array Configuration Guide*.

Note – If you use your storage array in a SAN-configured cluster, you must keep two FC switches configured in parallel. This configuration maintains cluster availability. See “[SAN Solutions in an Oracle Solaris Cluster Environment](#)” in *Oracle Solaris Cluster 3.3 Hardware Administration Manual* for more information.

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

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

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

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

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

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

10 Boot Node A into cluster mode.

For more information about how to boot nodes, see your Oracle Solaris Cluster system administration documentation.

11 Shut down Node B.

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

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

For the procedure about how to remove a fiber-optic cable, see the *Sun StorEdge T3 and T3+ Array Configuration Guide*.

Note – If you use your storage array in a SAN-configured cluster, you must keep two FC switches configured in parallel. This configuration maintains cluster availability. See “SAN Solutions in an Oracle Solaris Cluster Environment” in *Oracle Solaris Cluster 3.3 Hardware Administration Manual* for more information.

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

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

14 Remove the host adapter from Node B.

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

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

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

16 Boot Node B into cluster mode.

For more information about how to boot nodes, see your Oracle Solaris Cluster system administration documentation.

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

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


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

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

```
# cldevicegroup switch -n nodenamedevicegroup1[ devicegroup2 ...]
```

-n nodename

The node to which you are restoring device groups.

devicegroup1[devicegroup2 ...]

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

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

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

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

nodename

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

resourcegroup1[resourcegroup2 ...]

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

▼ How to Remove a Partner Group

Use this procedure to permanently remove storage array partner groups and their submirrors from a running cluster. To remove a storage array in single-controller configuration, see [“How to Remove a Storage Array in a Single-Controller Configuration”](#) on page 54.

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

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

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

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



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

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

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

- 1 If necessary, back up all database tables, data services, and volumes associated with each partner group that you are removing.**
- 2 If necessary, detach the submirrors from each storage array or partner group that you are removing. Detach the submirrors to stop all I/O activity to the storage array or partner group.** For more information, see your Veritas Volume Manager documentation.
- 3 Remove references to each LUN. Each LUN belongs to the storage array or partner group that you are removing.**

For more information, see your Veritas Volume Manager documentation.

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

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

Use the following command:

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

- 5 Shut down Node A.**

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

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

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

For the procedure about how to remove a fiber-optic cable, see the *Sun StorEdge T3 and T3+ Array Configuration Guide*.

Note – If you are using your storage arrays in a SAN-configured cluster, you must keep two FC switches configured in parallel. This configuration maintains cluster availability. See “[SAN Considerations](#)” on page 67 for more information.

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

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

9 Remove the host adapters from Node A.

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

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

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

11 Boot Node A into cluster mode.

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

12 Shut down Node B.

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

13 If any storage array that you are removing is the last storage array connected to an FC switch that is on Node B, disconnect the fiber-optic cable that connects this FC switch to Node B.

If array that you are removing is the last array connected to the node, , proceed to [Step 14](#).

For the procedure about how to remove a fiber-optic cable, see the *Sun StorEdge T3 and T3+ Array Configuration Guide*.

Note – If you are using your storage arrays in a SAN-configured cluster, you must keep two FC switches configured in parallel. This configuration maintains cluster availability. See “[SAN Considerations](#)” on page 67 for more information.

14 If you want to remove the host adapters from Node B, power off the node.

If you do not want to remove host adapters, skip to [Step 17](#).

15 Remove the host adapters from Node B.

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

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

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

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

TABLE 2-2 Task Map: Maintaining a Storage Array

| Task | Information |
|--|---|
| Upgrade storage array firmware. | “How to Upgrade Storage Array Firmware (No Submirrors)” on page 61 |
| Upgrade a StorEdge T3 array controller to a StorEdge T3+ array controller. | “How to Upgrade a StorEdge T3 Controller to a StorEdge T3+ Controller” on page 63 |
| Migrate to a partner group | “How to Migrate From a Single-Controller Configuration to a Partner-Group Configuration” on page 64 |

▼ How to Upgrade Storage Array Firmware (No Submirrors)

Use this procedure to upgrade storage array firmware in a running cluster, when your arrays are not configured to support submirrors. To upgrade firmware when you are using submirrors, see [“How to Upgrade Storage Array Firmware When Using Mirroring” on page 62](#). Firmware includes controller firmware, unit interconnect card (UIC) firmware, and disk drive firmware.



Caution – Perform this procedure on one storage array at a time. This procedure requires that you reset the storage arrays that you are upgrading. If you reset more than one storage array at a time, your cluster loses access to data.

Note – For all firmware installations, always read any README files that accompany the firmware patch for the latest information and special notes.

- 1 On one node that is attached to the storage array, detach the submirrors. This storage array is the storage array that you are upgrading.**

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

- 2 Apply the controller, disk drive, and UIC firmware 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*.

- 3 Reset the storage array, if you have not already done so.**

For the procedure about how to reboot a storage array, see the *Sun StorEdge T3 and T3+ Array Installation, Operation, and Service Manual*.

- 4 Reattach the submirrors to resynchronize the submirrors.**

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

▼ **How to Upgrade Storage Array Firmware When Using Mirroring**

Use this procedure to upgrade out-of-date controller firmware, disk drive firmware, or unit interconnect card (UIC) firmware. This procedure assumes that your cluster is operational. This procedure defines Node A as the node on which you are upgrading firmware. Node B is another node in the cluster.



Caution – Perform this procedure on one storage array at a time. This procedure requires that you reset the storage arrays that you are upgrading. If you reset more than one storage array at a time, your cluster loses access to data.

- 1 On the node that currently owns the disk group or diskset to which the mirror belongs, detach the storage array logical volume. This storage array is the storage array on which you are upgrading firmware.**

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

- 2 Apply the controller, disk drive, and UIC firmware patches.**

For the list of required storage array patches, see the *Sun StorEdge T3 Disk Tray Release Notes*. To apply firmware patches, see the firmware patch README file. To verify the firmware level, see the *Sun StorEdge T3 Disk Tray Release Notes*.

- 3 Disable the storage array controller that is attached to Node B. Disable the controller so that all logical volumes are managed by the remaining controller.**

For more information, see the *Sun StorEdge T3 and T3+ Array Administrator's Guide*.

- 4 On one node that is connected to the partner group, verify that the storage array controllers are visible to the node.**

```
# format
```

- 5 Enable the storage array controller that you disabled in [Step 3](#).**

- 6 Reattach the mirrors that you detached in [Step 1](#) to resynchronize the mirrors.**

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

▼ How to Upgrade a StorEdge T3 Controller to a StorEdge T3+ Controller

Use the following procedure to upgrade a StorEdge T3 storage array controller to a StorEdge T3+ storage array controller in a running cluster.



Caution – Perform this procedure on one storage array at a time. This procedure requires you to reset the storage arrays that you are upgrading. If you reset more than one storage array at a time, your cluster loses access to data.

- 1 **On one node that is attached to the StorEdge T3 storage array in which you are upgrading the controller, detach that storage array's submirrors.**

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

- 2 **Upgrade the StorEdge T3 storage array controller to a StorEdge T3+ storage array controller.**

For instructions, see the *Sun StorEdge T3 Array Controller Upgrade Manual*.

- 3 **Reattach the submirrors to resynchronize the submirrors.**

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

▼ How to Migrate From a Single-Controller Configuration to a Partner-Group Configuration

Use this procedure to migrate your storage arrays from a single-controller (noninterconnected) configuration to a partner-group (interconnected) configuration. This procedure assumes that the two storage arrays in the partner-group configuration are correctly isolated from each other on separate FC switches. Do not disconnect the cables from the FC switches or the nodes.



Caution – You must be an Oracle service provider to perform this procedure. If you need to migrate from a single-controller configuration to a partner-group configuration, contact your Oracle service provider.

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

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

- 1 **Back up all data on the storage arrays before you remove the storage arrays from the Oracle Solaris Cluster configuration.**

- 2 **Remove the noninterconnected storage arrays to be in your partner group from the cluster configuration.**

Follow the procedure in “[How to Remove a Storage Array in a Single-Controller Configuration](#)” on page 54.

- 3 **Connect and configure the single storage arrays to form a partner group.**

Follow the procedure in the *Sun StorEdge T3 and T3+ Array Field Service Manual*.

4 Ensure that each storage array has a unique target address.

For the procedure about how to verify and assign a target address to a storage array, see the *Sun StorEdge T3 and T3+ Array Configuration Guide*.

5 Ensure that the cache and mirror settings for each storage array are set to auto.**6 Ensure that the mp_support parameter for each storage array is set to mp_xio.****7 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*.

8 If necessary, install the required Oracle Solaris patches for storage array support on Node A.

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

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

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

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

9 Shut down Node A.

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

10 Perform a reconfiguration boot on Node A to create the new Oracle Solaris device files and links.

```
# boot -r
```

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

```
# devfsadm -C
```

12 On Node A, update the paths to the DID instances.

```
# cldevice clear
```

13 Label the new storage array logical volume.

For the procedure about how to label a logical volume, see the *Sun StorEdge T3 and T3+ Array Administrator's Guide*.

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

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

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

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

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

15 If necessary, install the required Oracle Solaris patches for storage array support on Node B.

For a list of required Oracle Solaris patches for storage array support, see the *Sun StorEdge T3 Disk Tray Release Notes*.

16 Shut down Node B.

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

-
- 17 **Perform a reconfiguration boot to create the new Oracle Solaris device files and links on Node B.**

```
# boot -r
```

- 18 **On Node B, update the /devices and /dev entries.**

```
# devfsadm -C
```

- 19 **On Node B, update the paths to the DID instances.**

```
# cldevice clear
```

- 20 **(Optional) On Node B, verify that the DIDs are assigned to the new LUNs.**

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

- 21 **On one node that is attached to the new storage arrays, reset the SCSI reservation state.**

```
# cldevice repair
```

Note – Repeat this command on the same node for each storage array LUN that you are adding to the cluster.

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

For more information, see your Veritas Volume Manager documentation.

SAN Considerations

This section contains information about how to use storage array in a SAN. This information is specific to a SAN in an Oracle Solaris Cluster environment. Use the cluster-specific procedures in this chapter to install and maintain a storage array in your cluster.

For instructions about how to create and maintain a SAN, see the documentation that shipped with your switch hardware. For information on switch ports, zoning, and required software and firmware, also see the documentation that shipped with your switch hardware.

SAN hardware includes the following components.

- FC switches
- FC host adapters
- Storage devices
- Enclosures

SAN software includes the following components.

- Drivers that are packaged with the operating system
- Firmware for the switches
- Management tools for the switches and storage devices

- Volume management software
- Other administration tools

SAN Clustering Considerations

If you are replacing an FC switch and you intend to save the switch IP configuration for restoration to the replacement switch, wait to connect the cables to the replacement switch. Connect the cables to the replacement switch *after* you recall the Fabric configuration to the replacement switch. For more information about how to save and recall switch configurations, see the documentation that shipped with your switch hardware.

Index

A

adding

See also installing

disk drive, 44

nodes, 43

partner groups, 28–34

single-controller storage arrays, 20–28

array-to-switch components, replacing, 49–50

arrays

See storage array s

installing partner groups, 16–19

C

chassis

replacing, 51, 62–63

configuring, storage arrays, 34–39

controller, upgrading, 44

controllers

replacing, 50–51

upgrading, 63–64

cooling units, replacing, 44

creating, logical volumes, 34–36

D

deleting logical volumes, 36–39

disk drive

adding, 44

removing, 44

disk drives, replacing, 44–45

E

Ethernet cables, replacing, 44

F

FC switches, replacing, 49–50

FRUs, 44

H

HBAs, replacing, 51–54

help, 7–8

host adapters, replacing, 51–54

hub-to-array components, replacing, 47–48

I

installing

See also adding

single-controller storage arrays, 12–15

storage arrays, 11–34

L

logical devices, *See* logical volumes

logical unit numbers, *See* logical volumes

logical volumes

 creating, 34–36

 removing, 36–39

LUNs, *See* logical volumes

M

maintaining

 storage arrays, 41–68

midplane, replacing, 51

migration, single-controller to partner-group

 configuration, 64–67

N

node-to-switch components, replacing, 45–47

nodes

 adding, 43

 removing, 43

P

partner groups

 adding, 28–34

 removing, 57–60

power cables, replacing, 44

power units, replacing, 44

R

removing

 disk drive, 44

 logical volumes, 36–39

 nodes, 43

 partner groups, 57–60

 single-controller arrays, 54–57

replacing

 array-to-switch components, 49–50

 chassis, 51, 62–63

 controllers, 50–51

replacing (*Continued*)

 cooling units, 44

 disk drives, 44–45

 Ethernet cables, 44

 FC switches, 49–50

 host adapters, 51–54

 hub-to-array components, 47–48

 midplane, 51

 node-to-switch components, 45–47

 power cables, 44

 power units, 44

 switch-to-array components, 47–48

 unit interconnect cards, 44

S

SAN

 clustering considerations, 67–68

 description, 67–68

single-controller to partner-group configuration

 migration, 64–67

storage area networks, *See* SAN

storage array firmware, upgrading, 61–62

storage arrays

 adding single controller, 20–28

 configuring, 34–39

 installing, 11–34

 installing single controller, 12–15

 maintaining, 41–68

 removing single controller, 54–57

 upgrading, 60–67

switch-to-array components, replacing, 47–48

T

technical support, 7–8

U

unit interconnect cards, replacing, 44

upgrading

 controller, 44

upgrading (*Continued*)

controllers, 63–64

storage array firmware, 61–62

storage arrays, 60–67

