

**Oracle® Solaris Cluster 3.3 With Sun
StorEdge 3900 Series or 6900 Series System
Manual**

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

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

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

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

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

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

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

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

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

Contents

| | |
|---|----|
| Preface | 5 |
| 1 Restrictions and Requirements | 11 |
| Restrictions | 11 |
| Requirements | 11 |
| 2 Installing and Configuring a Sun StorEdge 3900 or 6900 Series System | 13 |
| Installing Storage Systems | 13 |
| How to Install Storage Systems in a New Cluster | 14 |
| ▼ How to Add a Storage System to an Existing Cluster | 15 |
| Configuring Storage Systems | 16 |
| ▼ How to Create a Logical Volume | 16 |
| ▼ How to Remove a Logical Volume | 18 |
| 3 Maintaining a Sun StorEdge 3900 or 6900 Series System | 23 |
| Maintaining Storage Systems | 23 |
| FRUs That <i>Do Not</i> Require Oracle Solaris Cluster Maintenance Procedures | 25 |
| ▼ How to Remove a Storage System | 26 |
| ▼ SPARC: How to Replace a Virtualization Engine (<i>Sun StorEdge 6910 or Sun StorEdge 6960 Storage System Only</i>) | 27 |
| Replacing a Node-to-Switch Component | 27 |
| ▼ How to Upgrade Storage Array Firmware When Using Mirroring | 29 |
| ▼ How to Upgrade Storage Array Firmware When Not Using Mirroring | 30 |
| ▼ How to Replace a Disk Drive | 31 |
| ▼ How to Replace a Host Adapter | 32 |

| | |
|--|-----------|
| 4 Sun StorEdge 3900 and 6900 Series Storage System Cabling Diagrams | 37 |
| Installation Diagrams | 37 |
| | |
| Index | 41 |

Preface

The *Oracle Solaris Cluster 3.3 With Sun StorEdge 3900 Series or 6900 Series System Manual* provides procedures specific to Oracle's Sun StorEdge 3900 or 6900 storage devices that are placed in an Oracle Solaris Cluster environment.

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

Note – This Oracle Solaris Cluster release supports systems that use the SPARC and x86 families of processor architectures: UltraSPARC, SPARC64, and AMD64. In this document, the label x86 refers to systems that use the AMD64 family of processor architectures. The information in this document pertains to both platforms unless otherwise specified in a special chapter, section, note, bulleted item, figure, table, or example.

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

Who Should Use This Book

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

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

How This Book Is Organized

This book contains the following chapters.

[Chapter 1, “Restrictions and Requirements,”](#) lists requirements specific to the procedures in this book.

[Chapter 2, “Installing and Configuring a Sun StorEdge 3900 or 6900 Series System,”](#) discusses how to install Sun StorEdge 3900 or 6900 series storage systems and how to configure logical units on these systems.

[Chapter 3, “Maintaining a Sun StorEdge 3900 or 6900 Series System,”](#) describes how to maintain Sun StorEdge 3900 or 6900 series storage systems in a running cluster.

[Chapter 4, “Sun StorEdge 3900 and 6900 Series Storage System Cabling Diagrams,”](#) provides examples of supported cabling configurations.

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 3900 Series or 6900 Series System Manual

| Revision Date | New Information |
|----------------|--|
| January 2009 | Update links to other books in the Preface. |
| September 2010 | Update book with new product name and remove old CLI commands. |

Related Documentation

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

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

TABLE P-2 Hardware Documentation

| Title | Part Number |
|---|-------------|
| <i>Sun StorEdge Traffic Manager Installation and Configuration Guide</i> | 816-1420 |
| <i>Sun StorEdge Traffic Manager Software Release Notes</i> | 817-0385 |
| <i>Sun StorEdge T3 and T3+ Array Administrator's Guide</i> | 816-0776 |
| <i>Sun StorEdge T3 and T3+ Array Configuration Guide</i> | 816-0777 |
| <i>Sun StorEdge T3 and T3+ Array Installation, Operation, and Service Manual</i> | 816-0773 |
| <i>Sun StorEdge T3 Disk Tray Release Notes, Version 1.17b Controller Firmware</i> | 806-1497 |
| <i>Sun StorEdge 6920 System Host Installation Software Guide, Release 2.0</i> | 817-5831 |
| <i>Sun StorEdge 6920 Series Site Preparation Guide, Release 2.0</i> | 817-5224 |
| <i>Sun StorEdge 6920 System Getting Started Guide</i> | 817-5727 |
| <i>Sun StorEdge 6920 System Release Notes</i> | 817-5229 |
| <i>Sun StorEdge 3900 and 6900 Series 2.0 Reference and Service Manual</i> | 816-5253 |

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

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

Documentation, Support, and Training

See the following web sites for additional resources:

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

Oracle Welcomes Your Comments

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

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

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

Typographic Conventions

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

TABLE P-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% you have mail.</code> |
| AaBbCc123 | What you type, contrasted with onscreen computer output | <code>machine_name% su</code> Password: |
| <i>aabbcc123</i> | Placeholder: replace with a real name or value | The command to remove a file is <i>rm filename</i> . |
| <i>AaBbCc123</i> | Book titles, new terms, and terms to be emphasized | Read Chapter 6 in the <i>User's Guide</i> . A <i>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 | machine_name% |
| C shell for superuser | machine_name# |

Restrictions and Requirements

This chapter includes only restrictions and requirements that have a direct impact on the procedures in this book. For general support information, contact your Oracle service provider.

Restrictions

Storage devices that do *not* support LUN masking cannot reside in the same fabric zone as the Sun StorEdge 6920 storage system.

When using storage-based replication with the StorEdge 6920 system, do not configure a replicated volume as a quorum device. Locate any quorum devices on an unreplicated volume. See “Using Storage-Based Data Replication Within a Cluster” in *Oracle Solaris Cluster System Administration Guide* “Using Storage-Based Data Replication Within a Cluster” in *Oracle Solaris Cluster System Administration Guide* for more information on storage-based replication.

Requirements

Storage devices that support logical unit number (LUN) masking need to be implemented so that LUNs are *not* visible to the Sun StorEdge 6920 storage system.

Installing and Configuring a Sun StorEdge 3900 or 6900 Series System

This chapter contains the procedures about how to install and configure Sun StorEdge 3900 and 6900 series systems in an Oracle Solaris Cluster environment. It contains the following major sections:

- “Installing Storage Systems” on page 13
- “Configuring Storage Systems” on page 16

The storage system configuration utilities can be run from a menu-driven interface or a command line interface. This chapter describes the menu-driven interface.

For information about storage system architecture, features, and configuration utilities, see your storage system documentation listed in “[Related Documentation](#)” on page 6.

Installing Storage Systems

This section contains instructions on installing storage systems both in new clusters and to existing clusters.

TABLE 2-1 Task Map: Installing Storage Systems

| Task | Information |
|---|---|
| Install a storage system in a new cluster, before the OS and Oracle Solaris Cluster software are installed. | “ How to Install Storage Systems in a New Cluster ” on page 14 |
| Add a storage system to an operational cluster. | “ How to Add a Storage System to an Existing Cluster ” on page 15 |

How to Install Storage Systems in a New Cluster

Use this procedure to install a storage system in an initial cluster. To add a storage system to an existing cluster, use the procedure in [“How to Add a Storage System to an Existing Cluster” on page 15](#).

This procedure relies on the following assumptions:

- You have *not* installed the Solaris Operating System.
- You have *not* installed the Oracle Solaris Cluster software.
- Your cluster has *not* been established.
- You have enough host adapters to connect the nodes and the storage system to the service processor panel.

If you need to install host adapters, see [“How to Replace a Host Adapter” on page 32](#). When this procedure asks you to replace the failed host adapter, simply install the new host adapter instead.

▼ Install and Cable the Hardware

1 Unpack, place, and level the storage system.

For instructions, see your storage system installation documentation. For a list of this documentation, see [“Related Documentation” on page 6](#).

2 (Optional) Install the Fibre Channel (FC) switch for the storage system if you do not have a switch installed.

For the procedure about how to install an FC switch, see the documentation that shipped with your FC switch hardware.

3 Connect the nodes to the service processor panel.

- *(SAN Configuration)* Connect the FC switches to the service processor panel.
- *(Direct-Attached Configuration)* Connect each node to the service processor panel directly.

For the procedure about how to cable the storage system, see your storage system installation documentation.

4 Install the system power cord and the system antistatic strap.

For instructions, see your storage system installation documentation.

5 Power on the storage system and the nodes.

For instructions about how to power on the storage system, see your storage system installation documentation.

6 Configure the service processor.

For more information, see your storage system installation documentation.

▼ Install the Oracle Solaris Operating System and Configure Multipathing

- **On all nodes, install the Oracle Solaris operating system and apply the required Oracle Solaris patches for Oracle Solaris Cluster software and storage array support.**

For the procedure about how to install the Oracle Solaris operating environment, see [“How to Install Solaris Software”](#) in *Oracle Solaris Cluster Software Installation Guide*.

- See Also**
- To create a logical volume, see [“How to Create a Logical Volume”](#) on page 16.
 - To continue with Oracle Solaris Cluster software installation tasks, see your Oracle Solaris Cluster software installation documentation.

▼ How to Add a Storage System to an Existing Cluster

Use this procedure to add a new storage system to a *running* cluster. To install a new storage system in a Oracle Solaris Cluster configuration that is not running, use the procedure in [“How to Install Storage Systems in a New Cluster”](#) on page 14

Before You Begin This procedure relies on the following assumptions:

- You have enough host adapters to connect the nodes and the storage system to the service processor panel.
If you need to install host adapters, see [“How to Replace a Host Adapter”](#) on page 32. When this procedure asks you to replace the failed host adapter, simply install the new host adapter instead.
- You do *not* need to install a new node in your cluster.
If you need to add a node to your cluster, see your Oracle Solaris Cluster system administration documentation. Ensure that you install the required Oracle Solaris patches for storage system support.

1 Unpack, place, and level the storage system.

For instructions, see your storage system installation documentation. For a list of this documentation, see [“Related Documentation”](#) on page 6.

2 (Optional) Install the Fibre Channel (FC) switch for the storage system if you do not have a switch installed.

For the procedure about how to install an FC switch, see the documentation that shipped with your FC switch hardware.

3 Connect the nodes to the service processor panel.

- *(SAN Configuration)* **Connect the FC switches to the service processor panel.**
- *(Direct-Attached Configuration)* **Connect each node to the service processor panel directly.**

For the procedure about how to cable the storage system, see your storage system installation documentation.

4 Install the system power cord and the system antistatic strap.

For instructions, see your storage system installation documentation.

5 Power on the storage system and the nodes.

For instructions about how to power on the storage system, see your storage system installation documentation.

6 Configure the service processor.

For more information, see your storage system installation documentation.

- See Also**
- To create a logical volume, see [“How to Create a Logical Volume” on page 16](#).
 - If you added a new node, perform the steps in [“Install the Oracle Solaris Operating System and Configure Multipathing” on page 15](#) as appropriate. Apply the correct patches for storage system support and enable multipathing.

Configuring Storage Systems

This section contains the procedures about how to configure a storage system in a running cluster. [Table 2–2](#) lists these procedures.

TABLE 2–2 Task Map: Configuring a Storage System

| Task | Information |
|--------------------------|---|
| Create a logical volume. | “How to Create a Logical Volume” on page 16 |
| Remove a logical volume. | “How to Remove a Logical Volume” on page 18 |

▼ 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.

```
# cfdm -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 a Sun StorEdge 3900 or 6900 Series System

This chapter contains the procedures about how to maintain Sun StorEdge 3900 and 6900 series systems in an Oracle Solaris Cluster environment. It contains the following procedures:

- “How to Remove a Storage System” on page 26
- “SPARC: How to Replace a Virtualization Engine (*Sun StorEdge 6910 or Sun StorEdge 6960 Storage System Only*)” on page 27
- “Replacing a Node-to-Switch Component” on page 27
- “How to Upgrade Storage Array Firmware When Using Mirroring” on page 29
- “How to Upgrade Storage Array Firmware When Not Using Mirroring” on page 30
- “How to Replace a Disk Drive” on page 31
- “How to Replace a Host Adapter” on page 32

For information about storage system architecture, features, and configuration utilities, see your storage documentation listed in [“Related Documentation”](#) on page 6.

Maintaining Storage Systems

This section contains the procedures about how to maintain storage systems in a running cluster. [Table 3–1](#) lists these procedures. This section does not include procedures about how to add or remove disk drives. Storage arrays in your storage system operate only when fully configured with disk drives.



Caution – If you remove any field replaceable unit (FRU) from the storage arrays for an extended period of time, thermal complications might result. To prevent these complications, 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, the storage array, and all attached storage arrays, shut down and power off.

This caution does not apply to the StorEdge 6920 system.

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 3-1 Task Map: Maintaining a Storage System

| Task | Information |
|---|---|
| Remove a storage system. | “How to Remove a Storage System” on page 26 |
| Replace a virtualization engine. This procedure applies to a Sun StorEdge 6910 or a Sun StorEdge 6960 storage system <i>only</i> . | “SPARC: How to Replace a Virtualization Engine (<i>Sun StorEdge 6910 or Sun StorEdge 6960 Storage System Only</i>)” on page 27 |
| Replace a node-to-switch fiber-optic cable. | “Replacing a Node-to-Switch Component” on page 27 |
| Replace a gigabit interface converter (GBIC) or Small Form-Factor Pluggable (SFP) on a node's host adapter. | “Replacing a Node-to-Switch Component” on page 27 |
| Replace a GBIC or an SFP on an FC switch, connecting to a node. | “Replacing a Node-to-Switch Component” on page 27 |
| Upgrade storage array firmware. | “How to Upgrade Storage Array Firmware When Using Mirroring” on page 29 or “How to Upgrade Storage Array Firmware When Not Using Mirroring” on page 30 |

TABLE 3-1 Task Map: Maintaining a Storage System (Continued)

| Task | Information |
|---------------------------------------|--|
| Replace a disk drive. | “How to Replace a Disk Drive” on page 31 |
| Replace a node's host adapter. | “How to Replace a Host Adapter” on page 32 |
| Add a node to the storage array. | Oracle Solaris Cluster system administration documentation |
| Remove a node from the storage array. | Oracle Solaris Cluster system administration documentation |

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

This section contains lists of administrative tasks that require no cluster-specific procedures.

SPARC: FRUs for the Sun StorEdge 3900 Series, StorEdge 6910, and StorEdge 6960 Systems

SPARC: See the *Sun StorEdge T3 and T3+ Array Installation, Operation, and Service Manual* for the following procedures.

- SPARC: Replacing an Ethernet cable.
- SPARC: Replacing a power cable on the storage array.
- SPARC: Replacing a unit interconnect card (UIC).
- SPARC: Replacing a power and cooling unit (PCU).

SPARC: For the following procedures, see the *Sun StorEdge 3900 and 6900 Series 20 Reference and Service Manual*. For a URL to this storage documentation, see [“Related Documentation” on page 6](#).

- SPARC: Replacing an Ethernet hub.
- SPARC: Replacing a storage service processor.
- SPARC: Replacing a GBIC or an SFP on an FC switch that connects to a storage array.
- SPARC: Replacing a storage array-to-switch fiber-optic cable.
- SPARC: Replacing an FC switch.
- SPARC: Removing a virtualization engine.
- SPARC: Upgrading virtualization engine firmware.

FRUs for Sun StorEdge 6920 Storage Systems

- Replacing cards in a DSP
- Replacing a DSP
- Adding a disk drive
- Removing a disk drive
- Replacing a storage array's chassis
- Replacing an Ethernet cable
- Replacing an Ethernet hub in a cabinet
- Replacing an expansion cabinet.
- Replacing a fan assembly.
- Replacing a power cable in the storage system.
- Replacing a power and cooling unit (PCU).
- Replacing a service processor.
- Replacing a service processor accessory tray.
- Replacing a service processor panel.
- Replacing a USB relay panel.
- Replacing a USB flash disk.
- Extending a dual-controller configuration.
- Removing a dual-controller configuration.
- Replacing a controller.
- Replacing a battery.

▼ How to Remove a Storage System

Use this procedure to permanently remove a storage system from a running cluster.

This procedure defines Node N as the node that is connected to the storage system you are removing and the node with which you begin working.



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

Before You Begin

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

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

- 1 **If necessary, back up all database tables, data services, and volumes that are associated with each partner group that is affected.**

- 2 **Remove references to the volumes that reside on the storage system that you are removing.**
For more information, see your Solaris Volume Manager or Veritas Volume Manager documentation.
- 3 **Disconnect the cables that connected Node N to the FC switches in your storage system.**
- 4 **On all nodes, remove the obsolete Oracle Solaris links and device IDs.**

```
# devfsadm -C
# cldevice clear
```
- 5 **Repeat [Step 3](#) through [Step 4](#) for each node that is connected to the storage system.**

▼ **SPARC: How to Replace a Virtualization Engine** (*Sun StorEdge 6910 or Sun StorEdge 6960 Storage System Only*)

Use this procedure to replace a virtualization engine in a storage system in a running cluster.

- 1 **Replace the virtualization engine hardware.**
For instructions, see the *Sun StorEdge 3900 and 6900 Series 20 Reference and Service Manual*. For a URL to this storage documentation, see [“Related Documentation”](#) on page 6.
- 2 **On any node, view the virtualization engine controller status and enable the virtualization engine controllers.**

```
# cfgadm -al
# cfgadm -c configure c::controller id
```

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 32.

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 28.
- If your cluster does *not* use multipathing, see “[How to Replace a Node-to-Switch Component in a Cluster Without Multipathing](#)” on page 28.

▼ **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 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 and to verify the firmware level, see the *Sun StorEdge 3900 and 6900 Series Reference Manual*.
- To apply firmware patches, see the firmware patch README file.

For a URL to this storage documentation, see “[Related Documentation](#)” on page 6.

- 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 Storage Array Firmware When Not Using Mirroring**

In a partner-pair configuration, you can have nonmirrored data. However, this configuration requires you to shut down the cluster when upgrading firmware.

1 Shut down the entire cluster.

For the procedure about how to shut down a cluster, see your Oracle Solaris Cluster system administration documentation.

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

- For the list of required storage array patches, see the *Sun StorEdge 3900 and 6900 Series Reference Manual*.
- For the procedure about how to apply firmware patches, see the firmware patch README file.
- For the procedure about how to verify the firmware level, see the *Sun StorEdge 3900 and 6900 Series Reference Manual*.
- For a URL to this storage documentation, see [“Related Documentation” on page 6](#).

3 If you have not already done so, reset the storage arrays.

- For the procedure about how to reset a storage array, see the *Sun StorEdge T3 and T3+ Array Installation, Operation, and Service Manual*.
- For a URL to this storage documentation, see [“Related Documentation” on page 6](#).

4 Boot all nodes back into the cluster.

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

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

```
# format
```

▼ 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 `soLar.is.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.

▼ 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.

◆ ◆ ◆ 4 CHAPTER 4

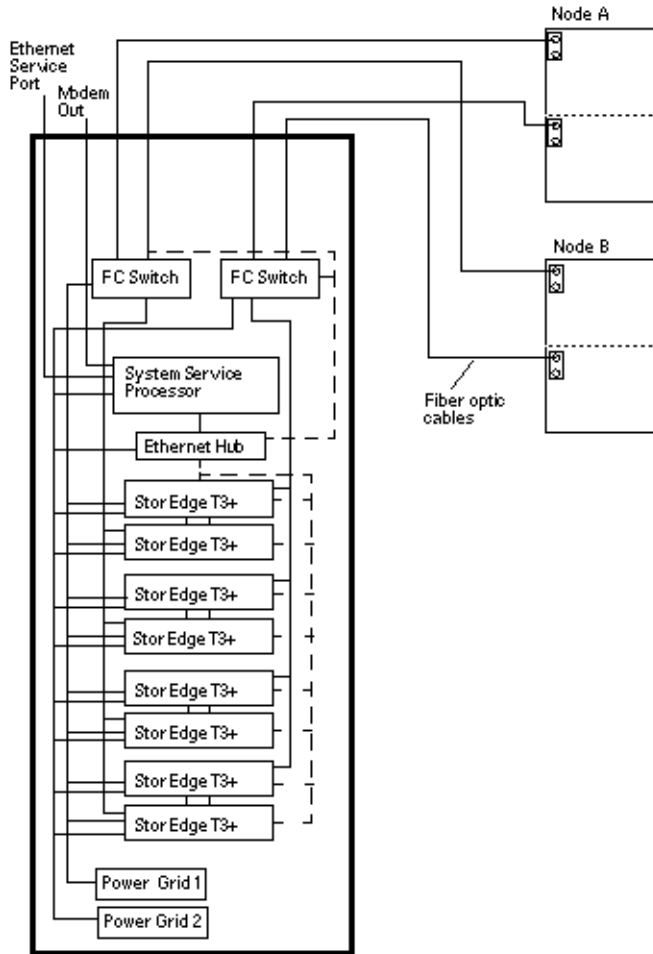
Sun StorEdge 3900 and 6900 Series Storage System Cabling Diagrams

The diagrams in this chapter provide examples of supported cabling configurations. Refer to the storage system documentation for additional configurations.

Installation Diagrams

SPARC: [Figure 4-1](#) illustrates one way of installing a Sun StorEdge 3900 series storage system.

FIGURE 4-1 SPARC: Installing a Sun StorEdge 3900 Series Storage System



You can install your Sun StorEdge 6920 storage system in several different configurations. [Figure 4-2](#) and [Figure 4-3](#) are two examples.

FIGURE 4-2 Sun StorEdge 6920 Direct-Connect Configuration





FIGURE 4-3 Sun StorEdge 6920 System Switched Configuration



Index

A

- accessory trays, replacing, 26
- adding
 - See also* installing
 - disk drives, 26
 - nodes, 25
 - storage systems, 15–16
- arrays, *See* storage arrays, storage systems

B

- batteries, replacing, 26

C

- chassis
 - replacing, 26, 29–30
- controllers
 - extending dual-controller configurations, 26
 - removing dual-controller configurations, 26
 - replacing, 26
- cooling units
 - replacing, 25, 26
- creating, logical volumes, 16–18

D

- deleting logical volumes, 18–21
- disk drives
 - adding, 26

disk drives (*Continued*)

- removing, 26
 - replacing, 31–32
- ## DSPs
- replacing, 26
 - replacing cards, 26
- ## dual-controller configurations
- extending, 26
 - removing, 26

E

- Ethernet cables
 - replacing, 25, 26
- Ethernet hubs
 - replacing, 25, 26
- expansion cabinets, replacing, 26
- extending, dual-controller configurations, 26

F

- fan assemblies, replacing, 26
- field replaceable units, *See* FRUs
- firmware, upgrading, 30–31
- flash disks, replacing, 26
- FRUs
 - SE 3900 systems, 25
 - SE 6910 systems, 25
 - SE 6920 systems, 26
 - SE 6960 systems, 25

G

GBICs, replacing, 25

H

HBAs, replacing, 32–35

help, 8

host adapters, replacing, 32–35

I

installing

See also adding

storage systems, 14–15

L

logical devices, *See* logical volumes

logical unit numbers, *See* logical volumes

logical volumes

creating, 16–18

removing, 18–21

LUN masking, 11

LUNs, *See* logical volumes

M

maintaining, storage systems, 23–35

N

node-to-switch components, replacing, 27–29

nodes

adding, 25

removing, 25

P

panels

replacing service processor panels, 26

replacing USB relay panels, 26

PCUs, *See* power units

power cables

replacing, 25, 26

power units

replacing, 25, 26

R

relay panels, replacing, 26

removing

disk drives, 26

dual-controller configurations, 26

logical volumes, 18–21

nodes, 25

storage systems, 26–27

virtualization engines, 25

replacing

accessory trays, 26

batteries, 26

cards in DSPs, 26

chassis, 26, 29–30

controllers, 26

cooling units, 25, 26

disk drives, 31–32

DSPs, 26

Ethernet cables, 25

Ethernet cables, 26

Ethernet hubs, 25, 26

expansion cabinets, 26

fan assemblies, 26

GBICs, 25

host adapters, 32–35

node-to-switch components, 27–29

power cables, 25, 26

power units, 25, 26

service processor panels, 26

services processors, 26

SFPs, 25

storage array -to-switch fiber-optic cables, 25

storage service processor, 25

replacing (*Continued*)

- switches, 25
 - unit interconnect cards, 25
 - USB flash disks, 26
 - USB relay panels, 26
 - virtualization engines, 27
- requirements, SE 6920 systems, 11
- restrictions, SE 6920 systems, 11

S

- SE 3900 systems, FRUs, 25
- SE 6910 systems, FRUs, 25
- SE 6920 systems
 - FRUs, 26
 - requirements, 11
 - restrictions, 11
- SE 6960 systems, FRUs, 25
- service processor panels, replacing, 26
- service processors
 - replacing, 26
 - replacing accessory trays, 26
- SFPs, replacing, 25
- storage array -to-switch fiber-optic cables,
 - replacing, 25
- storage array firmware, upgrading, 30–31
- storage arrays, *See* storage systems
- storage-based replication, 11
- storage service processor, replacing, 25
- storage systems, 13–21
 - adding, 15–16
 - installing, 14–15
 - maintaining, 23–35
 - removing, 26–27
- switches, replacing, 25
- systems, *See* storage systems

T

- technical support, 8

U

- UIC, *See* unit interconnect cards
- unit interconnect cards, replacing, 25
- upgrading firmware
 - storage arrays, 30–31
 - virtualization engines, 25
- USB flash disks, replacing, 26
- USB relay panels, replacing, 26

V

- virtualization engines
 - removing, 25
 - replacing, 27
 - upgrading firmware, 25

