

# **Oracle® Solaris Cluster 3.3 With Sun StorEdge 6320 System Manual**

## **License Restrictions Warranty/Consequential Damages Disclaimer**

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.

### **Warranty Disclaimer**

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.

### **Restricted Rights Notice**

If this is software or related 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.

### **Hazardous Applications Notice**

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 that 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.

### **Trademark Notice**

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

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. AMD, Opteron, the AMD logo, and the AMD Opteron logo are trademarks or registered trademarks of Advanced Micro Devices. UNIX is a registered trademark of The Open Group in the United States and other countries.

### **Third Party Content, Products, and Services Disclaimer**

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.

---

Ce logiciel et la documentation qui l'accompagne sont protégés par les lois sur la propriété intellectuelle. Ils sont concédés sous licence et soumis à des restrictions d'utilisation et de divulgation. Sauf disposition de votre contrat de licence ou de la loi, vous ne pouvez pas copier, reproduire, traduire, diffuser, modifier, breveter, transmettre, distribuer, exposer, exécuter, publier ou afficher le logiciel, même partiellement, sous quelque forme et par quelque procédé que ce soit. Par ailleurs, il est interdit de procéder à toute ingénierie inverse du logiciel, de le désassembler ou de le décompiler, excepté à des fins d'interopérabilité avec des logiciels tiers ou tel que prescrit par la loi.

Les informations fournies dans ce document sont susceptibles de modification sans préavis. Par ailleurs, Oracle Corporation ne garantit pas qu'elles soient exemptes d'erreurs et vous invite, le cas échéant, à lui en faire part par écrit.

Si ce logiciel, ou la documentation qui l'accompagne, est concédé sous licence au Gouvernement des Etats-Unis, ou à toute entité qui délivre la licence de ce logiciel ou l'utilise pour le compte du Gouvernement des Etats-Unis, la notice suivante s'applique :

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.

Ce logiciel ou matériel a été développé pour un usage général dans le cadre d'applications de gestion des informations. Ce logiciel ou matériel n'est pas conçu ni n'est destiné à être utilisé dans des applications à risque, notamment dans des applications pouvant causer des dommages corporels. Si vous utilisez ce logiciel ou matériel dans le cadre d'applications dangereuses, il est de votre responsabilité de prendre toutes les mesures de secours, de sauvegarde, de redondance et autres mesures nécessaires à son utilisation dans des conditions optimales de sécurité. Oracle Corporation et ses affiliés déclinent toute responsabilité quant aux dommages causés par l'utilisation de ce logiciel ou matériel pour ce type d'applications.

Oracle et Java sont des marques déposées d'Oracle Corporation et/ou de ses affiliés. Tout autre nom mentionné peut correspondre à des marques appartenant à d'autres propriétaires qu'Oracle.

AMD, Opteron, le logo AMD et le logo AMD Opteron sont des marques ou des marques déposées d'Advanced Micro Devices. Intel et Intel Xeon sont des marques ou des marques déposées d'Intel Corporation. Toutes les marques SPARC sont utilisées sous licence et sont des marques ou des marques déposées de SPARC International, Inc. UNIX est une marque déposée concédée sous licence par X/Open Company, Ltd.

# Contents

---

- Preface .....5**
- 1 Installing and Maintaining a Sun StorEdge 6320 System ..... 11**
  - Installing Storage Systems ..... 11
    - ▼ How to Install Storage Systems in a New Cluster ..... 12
    - Adding Storage Systems to an Existing Cluster ..... 14
  - Configuring Storage Systems ..... 18
    - ▼ How to Create a Logical Volume ..... 19
    - ▼ How to Remove a Logical Volume ..... 21
  - Maintaining Storage Systems ..... 23
    - StorEdge 6320 System FRUs ..... 24
    - ▼ How to Upgrade Storage Array Firmware ..... 25
    - ▼ How to Remove a Storage System ..... 26
      - Replacing a Node-to-Switch Component ..... 26
    - ▼ How to Replace a Chassis ..... 28
    - ▼ How to Replace a Host Adapter ..... 29
- Index .....33**



# Preface

---

The *Oracle Solaris Cluster 3.3 With Sun StorEdge 6320 System Manual* provides procedures that are specific to Oracle's Sun StorEdge 6320 system that is placed in an Oracle Solaris Cluster environment.

Use this manual with any version of Oracle Solaris Cluster 3.3 software on SPARC based clusters or x86 based clusters. See the [“Revision History” on page 6](#) for a list of changes to this manual.

---

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

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

---

## Who Should Use This Book

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

This document is intended for experienced system administrators with extensive knowledge of Oracle software and hardware. Do not use this document as a planning or presales 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 one chapter that consists of three major sections.

Section 1 discusses how to install Sun StorEdge 6320 storage arrays.

Section 2 discusses how to configure logical units on Sun StorEdge 6320 storage arrays.

Section 3 describes how to maintain Sun StorEdge 6320 storage arrays in a running cluster.

# Revision History

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

TABLE P-1 Oracle Solaris Cluster 3.3 With Sun StorEdge 6320 System Manual

Revision Date	New Information
January 2009	Updated links to Preface for different versions of Sun Cluster documentation.
September 2010	Updated product name to Oracle Solaris Cluster throughout the book and removed old CLI commands for the 3.3 release.
May 2011	Updated links to Oracle site for the 3.3 5/11 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://www.oracle.com/technetwork/indexes/documentation/index.html>.

TABLE P-2 Hardware Documentation

Title	Part Number
<i>Sun StorEdge 6320 System Reference and Service Manual</i>	816-7879
<i>Sun StorEdge 6320 System Installation Guide</i>	816-7878
<i>Sun StorEdge 6320 Series Site Preparation Guide</i>	816-7877
<i>Sun StorEdge 6020 and 6120 Array System Manual</i>	817-0200
<i>Sun StorEdge 6120 Array Installation Guide</i>	817-0199
<i>Sun StorEdge 6000 Family Host Installation Software Guide</i>	817-1739

TABLE P-3 Oracle Solaris Cluster and Sun Cluster Documentation

Documentation
<a href="#">Oracle Solaris Cluster 3.3</a>
<a href="#">Sun Cluster 3.2</a>

TABLE P-3 Oracle Solaris Cluster and Sun Cluster Documentation (Continued)

Documentation

[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 and Support

See the following web sites for additional resources:

- Documentation (<http://www.oracle.com/technetwork/indexes/documentation/index.html>)
- Support (<http://www.oracle.com/us/support/systems/index.html>)

# Oracle Software Resources

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>).
- Get hands-on step-by-step tutorials with Oracle By Example (<http://www.oracle.com/technetwork/tutorials/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.
<b>AaBbCc123</b>	What you type, contrasted with onscreen computer output	<code>machine_name%</code> <b>su</b> Password:
<i>aabbcc123</i>	Placeholder: replace with a real name or value	The command to remove a file is <i>rm filename</i> .



TABLE P-4    Typographic Conventions    (Continued)

Typeface	Meaning	Example
<i>AaBbCc123</i>	Book titles, new terms, and terms to be emphasized	Read Chapter 6 in the <i>User's Guide</i> .  <i>A cache</i> is a copy that is stored locally.  Do <i>not</i> save the file.  <b>Note:</b> Some emphasized items appear bold online.

# Shell Prompts in Command Examples

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

TABLE P-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#



# Installing and Maintaining a Sun StorEdge 6320 System

---

This chapter contains the procedures about how to install, configure, and maintain Oracle's Sun StorEdge 6320 system. These procedures are specific to an Oracle Solaris Cluster environment.

This chapter contains the following main topics:

- “Installing Storage Systems” on page 11
- “Configuring Storage Systems” on page 18
- “Maintaining Storage Systems” on page 23

For detailed information about storage system architecture, features, and configuration utilities, see the *Sun StorEdge 6320 System Reference and Service Manual* and the *Sun StorEdge 6320 System Installation Guide*.

## Installing Storage Systems

This section contains the procedures listed in [Table 1–1](#).

**TABLE 1–1** Task Map: Installing Storage Systems

Task	Information
Install storage systems in a new cluster, before the OS and Oracle Solaris Cluster software is installed.	<a href="#">“How to Install Storage Systems in a New Cluster” on page 12</a>
Add storage systems to an existing cluster.	<a href="#">“Adding Storage Systems to an Existing Cluster” on page 14</a>

You can install your storage system in several different configurations. Evaluate your needs. Determine which configuration is best for your situation. See the *Sun StorEdge 6320 System Installation Guide*, and “Installing Storage Arrays” in *Oracle Solaris Cluster 3.3 With Sun StorEdge 6120 Array Manual*.

## ▼ How to Install Storage Systems in a New Cluster

Use this procedure to install a storage system *before* you install the Oracle Solaris operating environment and Oracle Solaris Cluster software on your nodes. To add a storage system to an existing cluster, use the procedure in “[Adding Storage Systems to an Existing Cluster](#)” on [page 14](#).

**1 If necessary, install host adapters in the nodes to be connected to the storage system.**

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

**2 (StorEdge 6320SL storage system ONLY) Install the Fibre Channel (FC) switch for the storage system if you do not have a switch installed.**

---

**Note** – In a StorEdge 6320SL storage system, the customer provides the switch.

---

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

**3 Unpack, place, and level the storage system.**

For instructions, see the *Sun StorEdge 6320 System Installation Guide*.

**4 Install the system power cord and the system grounding strap.**

For instructions, see the *Sun StorEdge 6320 System Installation Guide*.

**5 (StorEdge 6320SL storage system ONLY) Connect the storage arrays to the FC switches by using fiber-optic cables.**



---

**Caution** – Do not connect the switch's Ethernet port to the storage system's private LAN.

---

For the procedure about how to cable the storage system, see the *Sun StorEdge 6320 System Installation Guide*.

**6 Power on the storage system and the nodes.**

For instructions about how to power on the storage system, see the *Sun StorEdge 6320 System Installation Guide*. For instructions about how to power on a node, see the documentation that shipped with your node hardware.

**7 Configure the service processor.**

For more information, see the *Sun StorEdge 6320 System Installation Guide*.

**8 Create a volume.**

For the procedure about how to create a volume, see the *Sun StorEdge 6320 System Reference and Service Manual*.

**9 (Optional) Specify initiator groups for the volume.**

For the procedure about how to specify initiator groups, see the *Sun StorEdge 6320 System Reference and Service Manual*.

**10 If necessary, reconfigure the storage system's FC switches to ensure that all nodes can access each storage array.**

The following configurations might prevent some nodes from accessing each storage array in the cluster.

- Zone configuration
- Multiple clusters that use same switch
- Unconfigured ports or misconfigured ports

**11 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*.

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 for the Oracle Solaris 10 OS, see the *Solaris Fibre Channel Storage Configuration and Multipathing Support Guide*.

**12 Update the Oracle Solaris device files and links.**

```
# devfsadm
```

---

**Note** – You can wait for the `devfsadm` daemon to automatically update the Oracle Solaris device files and links, or you can run the `devfsadm` command to immediately update the Oracle Solaris device files and links.

---

**13 Confirm that all storage arrays that you installed are visible to all nodes.**

```
# luxadm probe
```

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

## Adding Storage Systems to an Existing Cluster

Use this procedure to add a new storage system to a running cluster. To install systems to a new Oracle Solaris Cluster configuration that is not running, use the procedure in [“How to Install Storage Systems in a New Cluster” on page 12](#).

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

### ▼ How to Perform Initial Configuration Tasks on the Storage Array

- 1 *(StorEdge 6320SL storage system ONLY)* **Install the Fibre Channel (FC) switch for the storage system if you do not have a switch installed.**

---

**Note** – In a StorEdge 6320SL storage system, the customer provides the switch.

---

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

- 2 **Configure the service processor.**

For more information, see the *Sun StorEdge 6320 System Installation Guide*.

- 3 **Create a volume.**

For the procedure about how to create a volume, see the *Sun StorEdge 6320 System Reference and Service Manual*.

- 4 **(Optional) Specify initiator groups for the volume.**

For the procedure about how to specify initiator groups, see the *Sun StorEdge 6320 System Reference and Service Manual*.

- 5 **Unpack, place, and level the storage system.**

For instructions, see the *Sun StorEdge 6320 System Installation Guide*.

- 6 **Install the system power cord and the system grounding strap.**

For instructions, see the *Sun StorEdge 6320 System Installation Guide*.

- 7 *(StorEdge 6320SL storage system ONLY)* **Connect the storage arrays to the FC switches by using fiber-optic cables.**



---

**Caution** – Do not connect the switch's Ethernet port to the storage system's private LAN.

---

For the procedure about how to cable the storage system, see the *Sun StorEdge 6320 System Installation Guide*.

## 8 Power on the new storage system.

---

**Note** – The storage arrays in your system might require several minutes to boot.

---

For the procedure about how to power on the storage system, see the *Sun StorEdge 6320 System Installation Guide*.

## 9 If necessary, reconfigure the storage system's FC switches to ensure that all nodes can access each storage array.

The following configurations might prevent some nodes from accessing each storage array in the cluster:

- Zone configuration
- Multiple clusters that use the same switch
- Unconfigured ports or misconfigured ports

## ▼ How to Connect the Node to the FC Switches

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` role-based access control (RBAC) authorization.

### 1 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 +
```

### 2 Move all resource groups and device groups off Node N.

```
# clnode evacuate nodename
```

### 3 If you do not need to install one or more host adapters in Node N, skip to [Step 10](#).

To install host adapters, proceed to [Step 4](#).

### 4 If the host adapter that you that are installing is the first host adapter on Node N, determine whether the required drivers for the host adapter are already installed on this node.

For the required packages, see the documentation that shipped with your host adapters.

If this is not the first host adapter, skip to [Step 6](#).

**5 If the required support packages not already installed, install them.**

The support packages are located in the Product directory of the Oracle Solaris CD-ROM.

**6 Shut down and power off Node N.**

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

**7 Install one or more host adapters in Node N.**

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

**8 Power on and boot Node N into noncluster mode by adding -x to your boot instruction.**

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

**9 If necessary, upgrade the host adapter firmware on Node N.**

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* on the [Oracle Technology Network](#). 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 [Oracle Technology Network](#).

**10 If necessary, install a GBIC or an SFP in the FC switch or the storage array.**

For the procedure about how to install a GBIC or an SFP, see the documentation that shipped with your FC switch hardware or the *Sun StorEdge 6320 System Installation Guide*.

**11 Connect a fiber-optic cable between the FC switch and Node N.**

For the procedure about how to install a fiber-optic cable, see the *Sun StorEdge 6320 System Installation Guide*.



**12 Install the required Oracle Solaris patches for storage array support on Node N.**

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* on the [Oracle Technology Network](#). 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 [Oracle Technology Network](#).

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

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

**14 Update the Oracle Solaris device files and links.**

```
# devfsadm
```

---

**Note** – You can wait for the `devfsadm` daemon to automatically update the Oracle Solaris device files and links, or you can run the `devfsadm` command to immediately update the Oracle Solaris device files and links.

---

**15 On Node N, update the paths to the DID instances.**

```
# cldevice populate
```

**16 If necessary, label the new storage array logical volume.**

For the procedure about how to label a logical volume, see the *Sun StorEdge 6320 System Reference and Service Manual*.

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

```
# cldevice list -n NodeN -v
```

**18 Repeat [Step 2](#) through [Step 17](#) for each remaining node that you plan to connect to the storage array.**

- 19

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

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

# **cldevicegroup switch -n** *nodename* *devicegroup1* *devicegroup2* ...]

-n *nodename*

The node to which you are restoring device groups.

*devicegroup1* [*devicegroup2* ...]

The device group or groups that you are restoring to the node.
- 20

**(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.

Use the following command:

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

*nodename*

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

*resourcegroup1* [*resourcegroup2* ...]

The resource group or groups that you are returning to the node or nodes.
- 21

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

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

# Configuring Storage Systems

This section contains the procedures about how to configure a storage system in a running cluster. [Table 1–2](#) lists these procedures.

TABLE 1–2 Task Map: Configuring Storage Arrays

Task	Information
Create a volume.	<a href="#">“How to Create a Logical Volume” on page 19</a>
Remove a volume.	<a href="#">“How to Remove a Logical Volume” on page 21</a>

The following is a list of administrative tasks that require no cluster-specific procedures. See the *Sun StorEdge 6320 System Reference and Service Manual* for the following procedures.

- Creating a storage pool.
- Removing a storage pool.

- Creating a volume group.
- Removing a volume group.
- Creating an initiator group.
- Adding an initiator group.
- Removing an initiator group.

## ▼ 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.

```
# cfmadm -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 Storage Systems

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

TABLE 1–3 Task Map: Maintaining Storage Systems

| Task                            | Information                                                        |
|---------------------------------|--------------------------------------------------------------------|
| Remove a storage system.        | <a href="#">“How to Remove a Storage System” on page 26</a>        |
| Upgrade storage array firmware. | <a href="#">“How to Upgrade Storage Array Firmware” on page 25</a> |

**TABLE 1–3** Task Map: Maintaining Storage Systems (Continued)

| Task                                                                                                                                                                                            | Information                                                                                              |
|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|----------------------------------------------------------------------------------------------------------|
| Replace a node-to-switch component. <ul style="list-style-type: none"> <li>■ Node-to-switch fiber-optic cable</li> <li>■ FC host adapter</li> <li>■ FC switch</li> <li>■ GBIC or SFP</li> </ul> | <a href="#">“How to Replace a Node-to-Switch Component in a Cluster Without Multipathing” on page 27</a> |
| Replace a node's host adapter.                                                                                                                                                                  | <a href="#">“How to Replace a Host Adapter” on page 29</a>                                               |
| 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                                               |

## StorEdge 6320 System FRUs

The following is a list of administrative tasks that require no cluster-specific procedures. See the *Sun StorEdge 6320 System Reference and Service Manual* for the following procedures.

- Adding a disk drive.
- Replacing a failed disk drive.
- Removing a disk drive.
- Replacing an Ethernet cable.
- Replacing an Ethernet hub in a cabinet.
- Replacing the power cable on the storage array.
- Replacing a power and cooling unit (PCU).
- Replacing a storage service processor.
- Replacing a storage 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 chassis.
- Replacing a controller.
- Replacing a storage array-to-switch component in a dual-controller configuration.



A storage array-to-switch component includes all of the following components:

- Fiber-optic cable that connects an FC switch to a storage array
- FC switch GBIC or an SFP that connects an FC switch to a storage array
- FC switch that connects an FC switch to a storage array
- FC switch power cord
- Interconnect cables that connect two storage arrays
- Interconnect cards

## ▼ How to Upgrade Storage Array Firmware

Use this procedure to upgrade storage array firmware in a running cluster. Storage array firmware includes controller firmware, unit interconnect card (UIC) firmware, EPROM firmware, and disk drive firmware.

---

**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 Stop all I/O to the storage arrays you are upgrading.

### 2 Apply the controller, disk drive, and loop-card firmware patches by using the arrays' GUI tools.

For the list of required patches, see the *Sun StorEdge 6320 System Reference and Service 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 6320 System Reference and Service Manual*.

For specific instructions, see your storage array's documentation.

### 3 Confirm that all storage arrays that you upgraded are visible to all nodes.

```
# luxadm probe
```

### 4 Restart all I/O to the storage arrays.

You stopped I/O to these storage arrays in [Step 1](#).

## ▼ 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.

---

- 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 Solaris links and device IDs (DIDs).**  

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

## Replacing a Node-to-Switch Component

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

---

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

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

To replace a host adapter, see “[How to Replace a Host Adapter](#)” on page 29.

---

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

Ensure that you are following the appropriate instructions:

- If your cluster uses multipathing, see [“How to Replace a Node-to-Switch Component in a Cluster That Uses Multipathing” on page 27](#).
- If your cluster does *not* use multipathing, see [“How to Replace a Node-to-Switch Component in a Cluster Without Multipathing” on page 27](#).

## ▼ 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 Chassis

Use this procedure to replace a storage array chassis in a running cluster. This procedure assumes that you want to retain all FRUs other than the chassis and the backplane. To replace the chassis, you must replace both the chassis and the backplane. 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 storage array chassis or a storage array midplane, contact your Oracle service provider.

---

- 1 **Detach the submirrors on the storage array that is connected to the chassis and the midplane 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 **Are the storage arrays in your partner-group configuration redundant as a result of host-based mirroring?**
  - If yes, proceed to [Step 3](#).
  - If no, shut down the cluster.  
For the procedure about how to shut down a cluster, see your Oracle Solaris Cluster system administration documentation.
- 3 **Replace the chassis and the midplane.**  
For the procedure about how to replace a storage array chassis and a storage array midplane, see the *Sun StorEdge 6320 System Reference and Service Manual*.
- 4 **Did you shut down the cluster in [Step 2](#)?**
  - If no, proceed to [Step 5](#).
  - If yes, boot the cluster back into cluster mode.  
For the procedure about how to boot a cluster, see your Oracle Solaris Cluster system administration documentation.
- 5 **Reattach the submirrors that you detached in [Step 1](#) to resynchronize the submirrors.**




---

**Caution** – The world wide numbers (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* on the [Oracle Technology Network](#). 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 [Oracle Technology Network](#).

- 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.





# Index

---

## A

### adding

*See also* installing

disk drive, 24

initiator groups, 19

nodes, 24

partner groups, 24

storage systems, 14–18

arrays, *See* storage systems

## C

### chassis

replacing, 24, 28–29

configuring, storage systems, 18–23

controllers, replacing, 24

cooling units, replacing, 24

### creating

initiator groups, 19

logical volumes, 19–20

storage pools, 18

volume groups, 19

## D

deleting logical volumes, 21–23

disk drive, adding, 24

### disk drives

removing, 24

replacing, 24

## E

Ethernet cables, replacing, 24

Ethernet hubs, replacing, 24

## F

firmware, upgrade storage array firmware, 25

FRUs, 24–25

## H

HBAs, replacing, 29–31

help, 7–8

host adapters, replacing, 29–31

## I

### initiator groups

creating, 19

modifying, 19

removing, 19

### installing

*See also* adding

storage systems, 12–13

interconnect cables, replacing, 24

interconnect cards, replacing, 24

**L**

- logical devices, *See* logical volumes
- logical unit numbers
  - See* logical volumes
  - See* volumes
- logical volumes
  - creating, 19–20
  - removing, 21–23
- loop cables, *See* interconnect cables
- loop cards, *See* interconnect cards
- LUNs, *See* logical volumes

**M**

- maintaining, storage systems, 23–31
- midplane
  - replacing, 24, 28–29

**N**

- node-to-switch components, replacing, 26–28
- nodes
  - adding, 24
  - removing, 24

**P**

- partner groups
  - adding, 24
  - removing, 24
- pools, *See* storage pools
- power cables, replacing, 24
- power units, replacing, 24

**R**

- removing
  - disk drives, 24
  - initiator groups, 19
  - logical volumes, 21–23
  - nodes, 24

- removing (*Continued*)

- partner groups, 24
  - storage pools, 18
  - storage systems, 26
  - volume groups, 19

- replacing

- chassis, 24, 28–29
  - controllers, 24
  - cooling units, 24
  - disk drives, 24
  - Ethernet cables, 24
  - Ethernet hubs, 24
  - host adapters, 29–31
  - midplane, 24, 28–29
  - node-to-switch components, 26–28
  - power cables, 24
  - power units, 24
  - service processor accessory tray, 24
  - service processor panel, 24
  - storage service processor, 24
  - switch-to-storage array components, 24
  - USB flash disk, 24
  - USB relay panel, 24

**S**

- service processor, replacing, 24
- service processor accessory tray, replacing, 24
- service processor panel, replacing, 24
- SPA tray, *See* service processor accessory tray
- SSP, *See* service processor
- storage array firmware, upgrading, 25
- storage pools
  - creating, 18
  - removing, 18
- storage systems
  - adding, 14–18
  - configuring, 18–23
  - installing, 12–13
  - maintaining, 23–31
  - removing, 26
- switch-to-storage array components, replacing, 24
- systems, *See* storage systems

**T**

technical support, 7–8

**U**

upgrading, storage array firmware, 25

USB flash disk, replacing, 24

USB relay panel, replacing, 24

**V**

volume groups

    creating, 19

    removing, 19

