Sun Cluster 3.0-3.1 With Sun StorEdge 3900 Series or Sun StorEdge 6900 Series System Manual

For Solaris OS

Adobe PostScript™

060117@13215

Contents

Preface

The *Sun Cluster 3.0-3.1 With Sun StorEdge 3900 Series or Sun StorEdge 6900 Series System Manual* provides procedures specific to Sun StorEdge™ 3900 or 6900 storage devices that are placed in a Sun™ Cluster environment.

Use this manual with any version of Sun Cluster 3.0 and 3.1 software on SPARC® based clusters and any supported version of Sun Cluster 3.1 software on x86 based clusters. Unless otherwise noted, procedures are the same for all supported Sun Cluster 3.0 and 3.1 versions. See the for a list of changes to this manual.

Note – In this document, the term "x86" refers to the Intel 32-bit family of microprocessor chips and compatible microprocessor chips made by AMD.

Sun Cluster software runs on two platforms, SPARC® and x86. The information in this document pertains to both platforms unless otherwise specified in a special chapter, section, note, bulleted item, figure, table, or example.

This book contains information on several Sun StorEdge 3900 and 6900 series storage devices. Only the Sun StorEdge 6920 array is supported with x86 based clusters.

Who Should Use This Book

This book is for Sun representatives who are performing the initial installation of a Sun 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 Sun 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 the following chapters.

Chapter 1 lists requirements specific to the procedures in this book.

Chapter 2 discusses how to install Sun StorEdge 3900 or 6900 series storage systems and how to configure logical units on these systems.

Chapter 3 describes how to maintain Sun StorEdge 3900 or 6900 series storage systems in a running cluster.

Chapter 4 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 Sun Cluster 3.0-3.1 With Sun StorEdge 3900 Series or Sun StorEdge 6900 Series System Manual

Revision Date	New Information
January 2006	Extended support for the SE 6920 array to x86 based clusters. Added information in support of single, dual-port HBA configurations.

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 Sun Cluster books support the Sun Cluster 3.1, 2004, and 2005 releases. If you are maintaining a different version of Sun Cluster software, refer to the appropriate documentation. All Sun 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
Sun StorEdge Traffic Manager Installation and Configuration Guide	816-1420
Sun StorEdge Traffic Manager Software Release Notes: For the Solaris Operating System	817–0385
Sun StorEdge T3 and T3+ Array Administrator's Guide	816-0776
Sun StorEdge T3 and T3+ Array Configuration Guide	816-0777
Sun StorEdge T3 and T3+ Array Installation, Operation, and Service Manual	816-0773
Sun StorEdge T3 Disk Tray Release Notes, Version 1.17b Controller Firmware	806-1497
Sun StorEdge 6920 System Host Installation Software Guide, Release 2.0	817–5831
Sun StorEdge 6920 Series Site Preparation Guide, Release 2.0	817–5224
Sun StorEdge 6920 System Getting Started Guide	817–5727
Sun StorEdge 3900 and 6900 Series 2.0 Reference and Service Manual	817-5229
Sun StorEdge 3900 and 6900 Series 20 Reference and Service Manual	816-5253

TABLE P–3 Sun Cluster Documentation

Application	Title
Concepts	*Sun Cluster Concepts Guide for Solaris OS*
Overview	*Sun Cluster Overview for Solaris OS*

TABLE P–3 Sun Cluster Documentation *(Continued)*

Application	Title
Hardware administration	*Sun Cluster 3.0-3.1 Hardware Administration Manual for Solaris OS*
	Individual hardware administration guides
Software installation	*Sun Cluster Software Installation Guide for Solaris OS*
Data service administration	*Sun Cluster Data Services Planning and Administration Guide for Solaris OS*
	Individual data service guides
Data service development	*Sun Cluster Data Services Developer's Guide for Solaris OS*
System administration	*Sun Cluster System Administration Guide for Solaris OS*
Error messages	*Sun Cluster Error Messages Guide for Solaris OS*
Command and function reference	*Sun Cluster Reference Manual for Solaris OS*
Release Notes	*Sun Cluster 3.1 8/05 Release Notes for Solaris OS*
	Sun Cluster 3.x Release Notes Supplement

Using UNIX Commands

This document contains information about commands that are used to install, configure, or upgrade a Sun 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 Solaris™ Operating System (Solaris OS)
- Other software documentation that you received with your system
- Solaris Operating System man pages

Getting Help

If you have problems installing or using Sun 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, Solaris 8)
- The release number of Sun Cluster (for example, Sun Cluster 3.0)

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

Command	Function
`prtconf -v`	Displays the size of the system memory and reports information about peripheral devices
`psrinfo -v`	Displays information about processors
`showrev -p`	Reports which patches are installed
`prtdiag -v`	Displays system diagnostic information
`/usr/cluster/bin/scinstall -pv`	Displays Sun Cluster release and package version information

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

Accessing Sun Documentation Online

The docs.sun.com^{SM} Web site enables you to access Sun technical documentation online. You can browse the docs.sun.com archive or search for a specific book title or subject. The URL is http://docs.sun.com.

Ordering Sun Documentation

Sun Microsystems offers select product documentation in print. For a list of documents and how to order them, see "Buy printed documentation" at http://docs.sun.com.

Typographic Conventions

The following table describes the typographic changes used in this book.

TABLE P–4 Typographic Conventions

Typeface or Symbol	Meaning	Example
AaBbCc123	The names of commands, files, and directories; on-screen computer output	Edit your `.login` file. Use `ls -a` to list all files. `machine_name% you have mail.`
AaBbCc123	What you type, contrasted with on-screen computer output	`machine_name%` **su** `Password:`
AaBbCc123	Command-line placeholder: replace with a real name or value	To delete a file, type **rm** *filename*.
AaBbCc123	Book titles, new words, or terms, or words to be emphasized.	Read Chapter 6 in *User's Guide*. These are called *class* options. You must be *root* to do this.

Shell Prompts in Command Examples

The following table shows the default system prompt and superuser prompt for the C shell, Bourne shell, and Korn shell.

TABLE P–5 Shell Prompts

Shell	Prompt
C shell prompt	`machine_name%`
C shell superuser prompt	`machine_name#`
Bourne shell and Korn shell prompt	`$`
Bourne shell and Korn shell superuser prompt	`#`

Product Training

Sun Microsystems offers training in many Sun technologies through a variety of instructor-led courses and self-paced courses. For information about the training courses that Sun offers and to enroll in a class, visit Sun Microsystems Training at `http://training.sun.com/`.

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 Sun 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" in *Sun Cluster 3.0-3.1 Hardware Administration Manual for Solaris OS* 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 a Sun Cluster environment. It contains the following major sections:

- "Installing Storage Systems" on page 15
- "Configuring Storage Systems" on page 18

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 7.

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 Sun Cluster software are installed.	"How to Install Storage Systems in a New Cluster" on page 16
Add a storage system to an operational cluster.	"How to Add a Storage System to an Existing Cluster" on page 17

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 17.

This procedure relies on the following assumptions:

- You have *not* installed the Solaris Operating System.
- You have *not* installed the Sun 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

Steps   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 7.

   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 Solaris Operating System and Configure Multipathing

Steps 1. **On all nodes, install the Solaris operating system and apply the required Solaris patches for Sun Cluster software and storage array support.**

For the procedure about how to install the Solaris operating environment, see "How to Install Solaris Software" in *Sun Cluster Software Installation Guide for Solaris OS*.

2. **On all nodes, install any required patches or software for Sun StorEdge Traffic Manager software support to nodes and enable multipathing.**

For the procedure about how to install the Sun StorEdge Traffic Manager software, see "How to Install Sun Multipathing Software" in *Sun Cluster Software Installation Guide for Solaris OS*.

See Also ■ To create a logical volume, see "How to Create a Logical Volume" on page 19.

■ To continue with Sun Cluster software installation tasks, see your Sun 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 Sun Cluster configuration that is not running, use the procedure in "How to Install Storage Systems in a New Cluster" on page 16

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 Sun Cluster system administration documentation. Ensure that you install the required Solaris patches for storage system support.

Steps 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 7.

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 19.
- If you added a new node, perform the steps in "Install the Solaris Operating System and Configure Multipathing" on page 17 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 19
Remove a logical volume.	"How to Remove a Logical Volume" on page 20

▼ How to Create a Logical Volume

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

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 storage device is installed and configured. If you are using multipathing, the storage devices is configured as described in the installation procedure.
- If you are using Sun StorEdge Traffic Manager, it is installed and configured and the path between Traffic Manager and the storage device is functioning.

Steps 1. **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 7.**

- 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.

2. **Are you using multipathing?**

   - **If yes, proceed to Step 3.**

   - **If no, skip to Step 4.**

3. **Are any devices that are associated with the volume you created at an `unconfigured` state?**

   `# cfgadm -al | grep disk`

   - **If no, proceed to Step 4.**

   - **If yes, configure the Traffic Manager paths on each node that is connected to the storage device.**

     `cfgadm -o force_update -c configure` *controllerinstance*

For the procedure about how to configure Traffic Manager paths, see the *Sun StorEdge Traffic Manager Installation and Configuration Guide*.

4. On one node that is connected to the storage device, use the `format` command to label the new logical volume.

5. From any node in the cluster, update the global device namespace.

```
# scgdevs
```

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

---

6. To manage this volume with volume management software, use the appropriate Solstice DiskSuite/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 Solstice DiskSuite/Solaris Volume Manager or VERITAS Volume Manager documentation.

See Also
- To configure a logical volume as a quorum device, see Chapter 5, "Administering Quorum," in *Sun Cluster System Administration Guide for Solaris OS*.
- To create a new resource or configure a running resource to use the new logical volume, see Chapter 2, "Administering Data Service Resources," in *Sun Cluster Data Services Planning and Administration Guide for Solaris OS*.

▼ 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.

Steps 1. **Identify the logical volume that you are removing.**

Refer to your Solstice DiskSuite/Solaris Volume Manager or VERITAS Volume Manager documentation for more information.

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

3. **Check if the logical volume that you are removing is a quorum device.**

   ```
   # scstat -q
   ```

   If yes, choose and configure another device as the quorum device. Then remove the old quorum device.

   For procedures about how to add and remove quorum devices, see Chapter 5, "Administering Quorum," in *Sun Cluster System Administration Guide for Solaris OS*.

4. **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 Solstice DiskSuite/Solaris Volume Manager or VERITAS Volume Manager documentation.

5. **If you are using volume management software, run the appropriate Solstice DiskSuite/Solaris Volume Manager or VERITAS Volume Manager commands to remove the logical volume from any diskset or disk group.**

   For more information, see your Solstice DiskSuite/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 Sun 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

---

6. **If you are using multipathing, unconfigure the volume in Sun StorEdge Traffic Manager.**

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

7. **Access the storage device and remove the logical volume.**

   For the procedure about how to remove the volume, see your storage documentation. For a list of storage documentation, see "Related Documentation" on page 7.

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

   Record this information because you use it in Step 13 of this procedure to return resource groups and device groups to these nodes.

   ```
   # scstat
   ```

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

   ```
   # scswitch -s -h from-node
   ```

10. **Shut down and reboot Node A.**

    For the procedure about how to shut down and power off a node, see Chapter 3, "Shutting Down and Booting a Cluster," in *Sun Cluster System Administration Guide for Solaris OS*.

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

    ```
    # devfsadm -C
    # scdidadm -C
    ```

12. **For each additional node that is connected to the shared storage that hosted the logical volume, repeat Step 8 to Step 11.**

13. **(Optional) Return the resource groups and device groups that you identified in Step 8 to all cluster 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 a Sun 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 7.

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 `scdidadm -c` command, the following error message appears on your console if the device ID changed unexpectedly.

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

To fix device IDs that report this error, run the `scdidadm -R` 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 *only*.	"SPARC: How to Replace a Virtualization Engine *(Sun StorEdge 6910 or Sun StorEdge 6960 Storage System Only)*" 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

TABLE 3–1 Task Map: Maintaining a Storage System	*(Continued)*
Task	**Information**
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
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.	Sun Cluster system administration documentation
Remove a node from the storage array.	Sun Cluster system administration documentation

FRUs That *Do Not* Require Sun 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 7.

- 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.

Steps

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 Solstice DiskSuite/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.**

   ```
   # devfsadm -C
   # scdidadm -C
   ```

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.

Steps   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 7.

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

For the procedure about how 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

Steps   **1. Is your configuration active-passive?**

If yes, and 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

Steps   **1. Check if the physical data path failed.**

If no, proceed to Step 2.

If yes:

a. **Replace the component.**

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

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.

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

   ```
   # scstat
   ```

3. **Move all resource groups and device groups to another node.**

   ```
   # scswitch -s -h from-node
   ```

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

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

5. **(Optional) If necessary, return the resource groups and device groups that you identified in Step 2 to Node A.**

   ```
   # scswitch -z -g resource-group -h nodename
   # swswitch -z -D device-group -h nodename
   ```

▼ 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 procedures 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.

Steps 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 Solstice DiskSuite/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 for the procedure about how to verify the firmware level, 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 a URL to this storage documentation, see "Related Documentation" on page 7.

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 Solstice DiskSuite/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.

Steps   1. **Shut down the entire cluster.**

   For the procedure about how to shut down a cluster, see your Sun 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 7.

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 7.

4. **Boot all nodes back into the cluster.**

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

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 – 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.

Steps 1. **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 Solstice DiskSuite/Solaris Volume Manager or VERITAS Volume Manager documentation.

2. **Check if the logical volume (in Step 1) is a quorum device.**

   ```
   # scstat -q
   ```

   If yes, choose and configure another device to be the new quorum device. Remove the old quorum device. For procedures about how to add and remove quorum devices, see your Sun Cluster system administration documentation.

3. **Replace the failed disk drive.**

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

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

   For the procedure about how to add a quorum device, see your Sun Cluster system administration documentation.

5. **Did you detach a submirror or plex in Step 1?**

   - If no, you are finished with this procedure.

   - If yes, use volume manager commands to reattach the submirror or plex.

     For more information, see your Solstice DiskSuite/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 *Sun Cluster Hardware Administration Manual for Solaris OS* 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 *Sun Cluster 3.0-3.1 Hardware Administration Manual for Solaris OS*.

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

Steps   1. **Determine the resource groups and device groups that are running on `Node A`.**

   Record this information because you use this information in Step 9 of this procedure to return resource groups and device groups to `Node A`.

   ```
   # scstat
   ```

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

   ```
   # scswitch -S -h nodename
   ```

3. **Shut down `Node A`.**

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

4. **Power off `Node A`.**

5. **Replace the failed host adapter.**

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

6. **Do you need to upgrade the node's host adapter firmware?**

   - **If yes, boot `Node A` into noncluster mode. Proceed to Step 7.**

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

   - **If no, proceed to Step 8.**

7. **Upgrade the host adapter firmware on `Node A`.**

   PatchPro is a patch-management tool that eases the selection and download of patches required for installation or maintenance of Sun Cluster software. PatchPro provides an Interactive Mode tool especially for Sun Cluster. The Interactive Tool makes the installation of patches easier. PatchPro's Expert Mode tool helps you to maintain your configuration with the latest set of patches. Expert Mode is especially useful for obtaining all of the latest patches, not just the high availability and security patches.

   To access the PatchPro tool for Sun Cluster software, go to http://www.sun.com/PatchPro/, click Sun Cluster, then choose either Interactive Mode or Expert Mode. Follow the instructions in the PatchPro tool to describe your cluster configuration and download the patches.

   For third-party firmware patches, see the SunSolve[SM] Online site at `http://sunsolve.ebay.sun.com`.

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

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

9. **Return the resource groups and device groups you identified in Step 1 to `Node A`.**

   ```
   # scswitch -z -g resource-group -h nodename
   # scswitch -z -D device-group-name -h nodename
   ```

   For more information, see your Sun Cluster system administration documentation.

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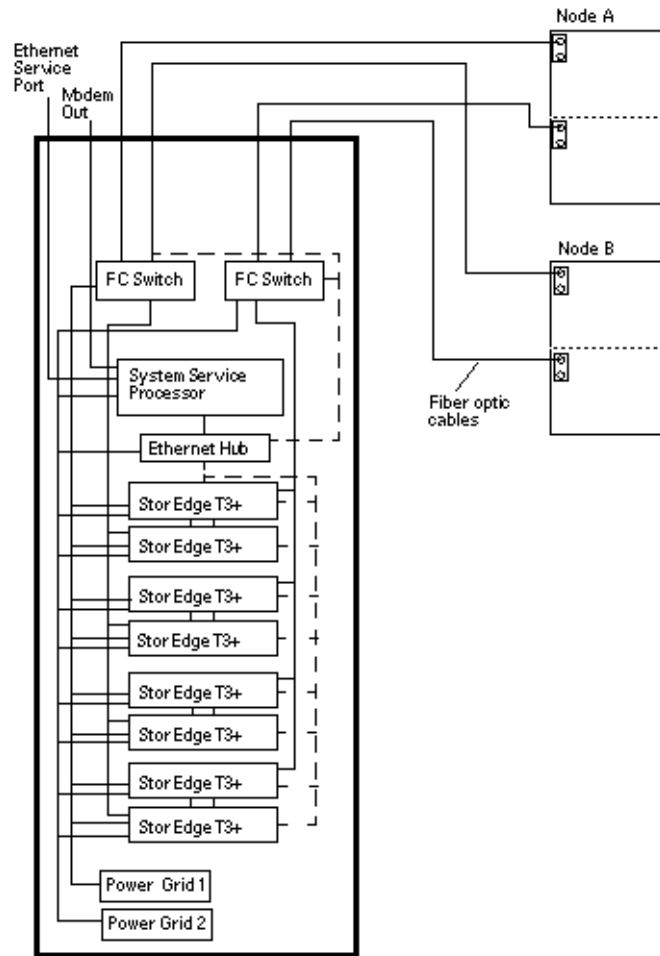
FIGURE 4–1 SPARC: Installing a Sun StorEdge 3900 Series Storage System

You can install your Sun StorEdge 6920 storage system in several different configurations. Figure 4–2 and Figure 4–3 are two examples.
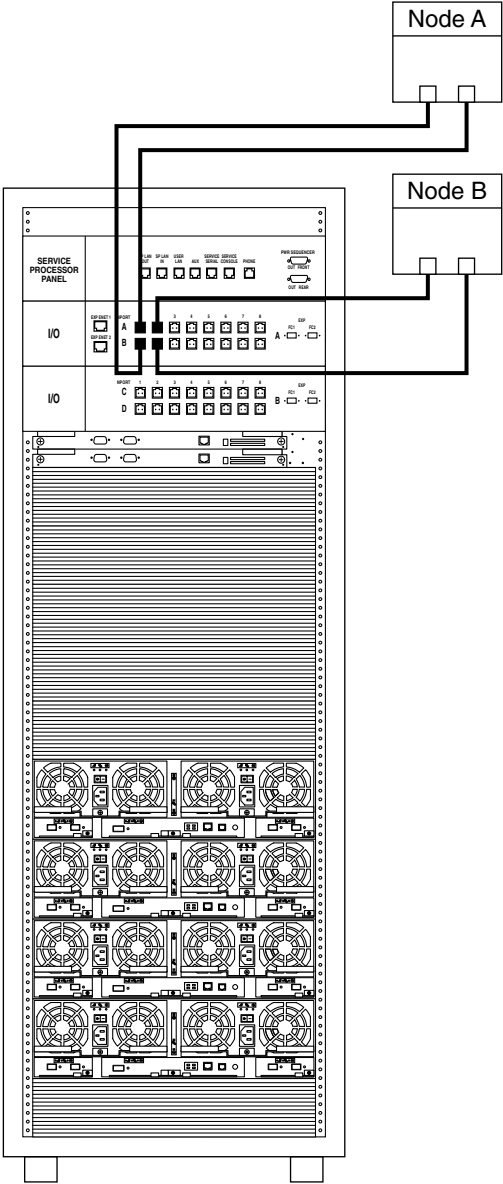
FIGURE 4–2 Sun StorEdge 6920 Direct-Connect Configuration
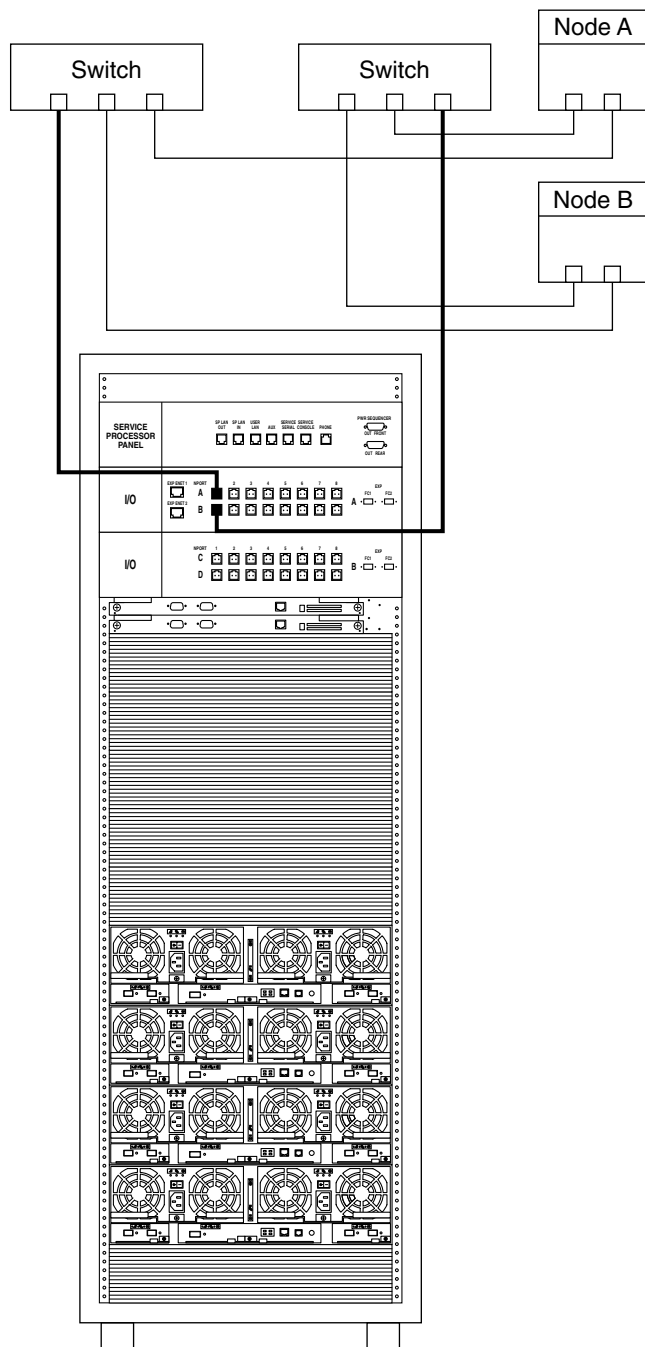
FIGURE 4–3 Sun StorEdge 6920 System Switched Configuration

Index