Sun Cluster 3.0-3.1 With SCSI JBOD Storage Device Manual for Solaris OS

Adobe PostScript™

060117@13215

Contents

Preface

The *Sun Cluster 3.0-3.1 With SCSI JBOD Storage Device Manual for Solaris OS* provides procedures specific to SCSI JBOD storage devices in a Sun™ Cluster environment. Use this manual with any version of Sun Cluster 3.0 and 3.1 software on SPARC® based systems and with any supported versions of Sun Cluster 3.1 software on x86 based systems.

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.

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

Who Should Use This Book

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

Revision History

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

TABLE P–1 Sun Cluster 3.0-3.1 With SCSI JBOD Storage Device Manual for Solaris OS

Revision Date	New Information
January 2006	Added information on single dual-port HBA configurations.
October 2005	Added Sun StorEdge 3320 as a supported array.
September 2005	Added information to support Sun StorEdge™ 3310 and 3120 JBOD devices on x86 based systems.

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.

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

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

TABLE P–2 Sun StorEdge 3310 and 3320 SCSI Storage Array

Title	Part Number
OpenBoot 4.x Command Reference Manual	817–1177
Available at `http://docs.sun.com/db/doc/817-1177-10`	
Sun StorEdge 3000 Family Installation, Operation, and Service Manual, 3320 SCSI Array	819–1274
Available at `http://docs.sun.com/db/doc/819-1274-11`	
Sun StorEdge 3000 Family Installation, Operation, and Service Manual, 3310 SCSI Array	816–7290
Sun StorEdge 3000 Family FRU Installation Guide	816–7326
Sun StorEdge 3000 Family CLI 2.0 User's Guide	817–4951
Sun StorEdge PCI Dual Ultra3 SCSI Host Adapter Release Notes	816-2157
Sun StorEdge 3310 SCSI Array Release Notes	816-7292
Sun StorEdge 3320 SCSI Array Release Notes	816-7660
Available at `http://docs.sun.com/db/doc/816-7660-10`	

TABLE P–3 Sun StorEdge 3120 Storage Array

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

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

Title	Part Number
Netra st A1000/D1000 Installation and Maintenance Manual	805-7147
Netra st A1000/D1000 Product Notes	805-7148
OpenBoot 4.x Command Reference Manual	806-2906
Available at `http://docs.sun.com/db/doc/806-2906-10`	
Sun StorEdge A1000 and D1000 Installation, Operations, and Service Manual	805-2624

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

Title	Part Number
Netra st D130 Installation and Maintenance Manual	806-1489

TABLE P–6 SPARC: Sun StorEdge D2 Storage Array

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

TABLE P–7 SPARC: Sun StorEdge Multipack Storage Array

Title	Part Number
OpenBoot 4.x Command Reference Manual	806-2906
Available at `http://docs.sun.com/db/doc/806-2906-10`	
SPARCstorage MultiPack 2 Service Manual	805–1667

TABLE P–8 Sun Cluster Documentation

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

TABLE P–8 Sun Cluster Documentation   *(Continued)*

Application	Title
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 3.0-3.1 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 9)
- The release number of Sun Cluster (for example, Sun Cluster 3.1 8/05)

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.

Documentation, Support, and Training

The Sun web site provides information about the following additional resources:

- Documentation (`http://www.sun.com/documentation/`)
- Support (`http://www.sun.com/support/`)
- Training (`http://www.sun.com/training/`)

Typographic Conventions

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

TABLE P–9 Typographic Conventions

Typeface	Meaning	Example
`AaBbCc123`	The names of commands, files, and directories, and onscreen computer output	Edit your `.login` file. Use `ls -a` to list all files. `machine_name% you have mail.`

TABLE P–9 Typographic Conventions   *(Continued)*

Typeface	Meaning	Example
AaBbCc123	What you type, contrasted with onscreen computer output	`machine_name%` **su** `Password:`
aabbcc123	Placeholder: replace with a real name or value	The command to remove a file is rm *filename*.
AaBbCc123	Book titles, new terms, and terms to be emphasized	Read Chapter 6 in the *User's Guide*. A *cache* is a copy that is stored locally. Do *not* save the file. **Note:** Some emphasized items appear bold online.

Shell Prompts in Command Examples

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

TABLE P–10 Shell Prompts

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

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 Training at `http://training.sun.com/`.

Installing a SCSI JBOD Storage Device

This chapter describes the procedures about how to install SCSI JBOD storage devices in a Sun™ Cluster environment.

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

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

Installing a Storage Array

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

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

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

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

Number of Devices[1]	Maximum Length
UltraSCSI (40MB/s: FAST-20, Wide)	
up to 4	3 meters
up to 8	1.5 meters
FastSCSI (20MB/s: FAST-10, Wide)	
up to 16	6 meters

[1] Devices include both targets and initiators.

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

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

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

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

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

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

TABLE 1–2 Task Map: Installing Storage Arrays

Task	Information
Install a storage array in a new cluster, before the OS and Sun Cluster software are installed.	"SPARC: How to Install a Storage Array in a New SPARC Based Cluster" on page 15
	"x86: How to Install a Storage Array in a New X86 Based Cluster" on page 20

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

Task	Information
Add a storage array to an existing cluster.	"SPARC: How to Add a Storage Array to an Existing SPARC Based Cluster" on page 26
	"x86: How to Add a Storage Array to an Existing X86 Based Cluster" on page 31

Setting SCSI Initiator IDs

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

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

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

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

This procedure assumes that you are installing one or more storage arrays at initial installation of a SPARC based cluster. If you are adding arrays to a running cluster, use the procedure in "SPARC: How to Add a Storage Array to an Existing SPARC Based Cluster" on page 26.

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

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

Before You Begin

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

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

  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`.

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

Steps

1. **Verify that the storage arrays are set up correctly for your planned configuration.**

2. **If necessary, install the host adapters in the nodes that you plan to connect to the storage array.**

   If possible, put each host adapter on a separate I/O board to ensure maximum redundancy.

3. **Cable the storage arrays.**

   For cabling diagrams, see Chapter 3.

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

For a full list of commands, see the *OpenBoot 2.x Command Reference Manual*.

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

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

f. **Store the changes.**

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

7. **Install the operating system software.**

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

See your Sun Cluster installation documentation for instructions.

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

See the host adapter documentation for driver installation procedures .

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

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. **If you are using Sun StorEdge 3310 JBOD arrays with the Sun StorEdge PCI/PCI-X Dual Ultra320 SCSI host adapter, you must throttle down the speed of**

the adapter to U160. Add the following entry to your `/kernel/drv/mpt.conf` file:

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

9. **Install the Sun Cluster software and volume management software.**

   For software installation procedures, see the Sun Cluster installation documentation.

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

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

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

▼ x86: How to Install a Storage Array in a New X86 Based Cluster

This procedure assumes that you are installing one or more storage arrays at initial installation of an x86 based cluster. If you are adding arrays to a running cluster, use the procedure in "SPARC: How to Add a Storage Array to an Existing SPARC Based Cluster" on page 26.

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

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

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

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

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

Steps 1. **Verify that the storage arrays are set up correctly for your planned configuration.**

2. **If necessary, install the host adapters in the nodes that you plan to connect to the storage array.**

   If possible, put each host adapter on a separate BUS to ensure maximum redundancy.

3. **Power on one node.**

4. **On the first node, ensure that each device in the SCSI chain has a unique SCSI address by configuring the initiator IDs in the `BIOS`.**

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

   Note – Perform these steps on only one cluster node.

   ---

   a. **Access your `BIOS` settings.**

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

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

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

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

   For cabling diagrams, see Chapter 3.

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

   ---

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

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

7. **Power on the storage array.**

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

8. **Install the operating system software on the node for which you configured the `BIOS` in Step 4.**

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

      See your Sun Cluster installation documentation for instructions.

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

      For driver installation procedures, see the host adapter documentation.

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

      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`.

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

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

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

   To find this information on the V40z server with SG-XCPI2SCSI-LM320 Sun StorEdge PCI/PCI-X Dual Ultra320 SCSI Host Adapter cards, use the following command:

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

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

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

   c. **Include the following entries:**

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

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

   The entries in this example have the following meanings:

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

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

   `# boot`

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

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

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

    ■ **On systems that run Solaris 8 Update 7 or below, perform a reconfiguration boot.**

    ■ **For Solaris 9 and above, run the `update_drv -f sd` command and then the `devfsadm` command.**

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

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

       See your Sun Cluster installation documentation for instructions.

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

       For driver installation procedures, see the host adapter documentation.

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

       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`.

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

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

14. **To make the changes to the `/kernel/drv/sd.conf` file active, perform one of the following options.**

    ■ **On systems that run Solaris 8 Update 7 or below, perform a reconfiguration boot.**

    ■ **For Solaris 9 and above, run the `update_drv -f sd` command and then the `devfsadm` command.**

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

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

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

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

17. **To make the changes to the `/kernel/drv/sd.conf` file active, perform one of the following options.**

    ■ **On systems that run Solaris 8 Update 7 or below, perform a reconfiguration boot.**

    ■ **For Solaris 9 and above, run the `update_drv -f sd` command and then the `devfsadm` command.**

18. **Install the Sun Cluster software and volume management software on each node.**

    For software installation procedures, see the Sun Cluster installation documentation.

Example 1–1 A Completed `mpt.conf` File

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

■ The output of the format command is that shown in Step a.

■ You are using Sun StorEdge 3310 JBOD arrays with the Sun StorEdge PCI/PCI-X Dual Ultra320 SCSI host adapter.

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

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

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

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

This procedure contains instructions for adding storage arrays to an operational cluster. If you need to install storage arrays to a new cluster, use the procedure in "SPARC: How to Install a Storage Array in a New SPARC Based Cluster" on page 15 or "x86: How to Install a Storage Array in a New X86 Based Cluster" on page 20.

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

- You need to increase your storage pool.
- You need to upgrade to a higher-quality or to a larger storage array.

  To upgrade storage arrays, remove the old storage array and then add the new storage array.

  To replace a storage array with the same type of storage array, see "How to Replace the Chassis" on page 43.

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

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

- Your cluster is operational and all nodes are powered on.
- Your nodes are not configured with dynamic reconfiguration functionality.

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

Steps 1. **Install all software that is specific to the storage array or to any new host adapters.**

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

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`.

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

   ```
   # scswitch -S -h from-node
   ```

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

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

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

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

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

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

   - If necessary, terminate the ports that will connect to `Node B`.

     - If you have a Netra™ D130 array, always terminate the ports that connect to `Node B`.

     - If you have a StorEdge 3310 or 3320 SCSI array, terminate the ports that connect to `Node B` when using a split-bus configuration.

   - If your storage array is a StorEdge 3310 or 3320 SCSI array, do not power on the storage array until the storage array is cabled to `Node A`.

   For cabling diagrams, see Chapter 3.

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

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

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

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

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

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

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

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

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

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

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

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

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

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

      For a full list of commands, see the *OpenBoot 2.x Command Reference Manual*.

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

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

f. **Store the changes.**

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

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

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

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

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

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

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

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

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

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

6. **Perform a reconfiguration boot on `Node A` to create the new Solaris device files and links.**

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

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

   `# scdidadm -l`

9. **Move all resource groups and device groups off `Node B`.**

   `# scswitch -S -h` *from-node*

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

    a. **Shut down `Node B`.**

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

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

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

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

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

    If you added port terminators in Step 4, remove the terminator ports and connect the storage array to `Node B`.

    For cabling diagrams, see Chapter 3.

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

    ---

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

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

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

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

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

13. **Perform a reconfiguration boot to create the new Solaris device files and links.**

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

```
# scdidadm -L
```

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

For more information, see your Solstice DiskSuite/Solaris Volume Manager or VERITAS Volume Manager documentation.

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

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

▼ x86: How to Add a Storage Array to an Existing X86 Based Cluster

This procedure contains instructions for adding storage arrays to an operational cluster. If you need to install storage arrays to a new cluster, use the procedure in "SPARC: How to Install a Storage Array in a New SPARC Based Cluster" on page 15 or "x86: How to Install a Storage Array in a New X86 Based Cluster" on page 20.

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

- You need to increase your storage pool.
- You need to upgrade to a higher-quality or to a larger storage array.

  To upgrade storage arrays, remove the old storage array and then add the new storage array.

  To replace a storage array with the same type of storage array, see "How to Replace the Chassis" on page 43.

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

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

- Your cluster is operational and all nodes are powered on.
- Your nodes are not configured with dynamic reconfiguration functionality.

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

Steps   **1. Install all software that is specific to the storage array or to any new host adapters.**

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

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.

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

```
# scswitch -S -h from-node
```

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

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

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

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

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

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

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

---

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

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

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

c. **Select a unique value for each of the new host adapter's ports.**

d. **Set each controller's `scsi-initiator-id` to that value.**

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

- If necessary, terminate the ports that will connect to `Node B`.

  - If you have a Netra™ D130 array, always terminate the ports that connect to `Node B`.
  - If you have a StorEdge 3310 or 3320 SCSI array, terminate the ports that connect to `Node B` when using a split-bus configuration.

- If your storage array is a StorEdge 3310 or 3320 SCSI array, do not power on the storage array until the storage array is cabled to `Node A`.

For cabling diagrams, see Chapter 3.

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

---

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

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

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

      To find this information on the V40z server with SG-XCPI2SCSI-LM320 Sun StorEdge PCI/PCI-X Dual Ultra320 SCSI Host Adapter cards, use the following command:

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

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

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

   c. **Include the following entries:**

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

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

      The entries in this example have the following meanings:

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

7. **Perform a reconfiguration boot on `Node A` to create the new Solaris device files and links.**

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

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

   `# scdidadm -l`

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

    `# scswitch -S -h` *from-node*

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

    a. **Shut down `Node B`.**

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

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

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

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

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

    If you added port terminators in Step 5, remove the terminator ports and connect the storage array to `Node B`.

    For cabling diagrams, see Chapter 3.

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

    ---

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

    a. **Access your `BIOS` settings.**

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

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

14. **Perform a reconfiguration boot to create the new Solaris device files and links.**

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

```
# scdidadm -L
```

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

For more information, see your Solstice DiskSuite/Solaris Volume Manager or VERITAS Volume Manager documentation.

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

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

Maintaining a SCSI JBOD Storage Device

This chapter describes the procedures about how to maintain SCSI JBOD storage devices in a Sun™ Cluster environment.

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

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

This chapter contains the following procedures:

Maintaining Storage Arrays

The maintenance procedures in "FRUs That *Do Not* Require Sun Cluster Maintenance Procedures" on page 39 are performed the same as in a noncluster environment. Table 2–1 lists the procedures that require cluster-specific steps.

TABLE 2–1 Task Map: Maintaining a Storage Array

Task	Information
Remove a storage array	"How to Remove a Storage Array" on page 41
Replace a storage array To replace a storage array, remove the storage array. Add a new storage array to the configuration.	"SPARC: How to Add a Storage Array to an Existing SPARC Based Cluster" on page 26 "How to Remove a Storage Array" on page 41
Add a JBOD as an expansion unit. Follow the same procedure that you use in a noncluster environment.	*Sun StorEdge 3000 Family Installation, Operation, and Service Manual*
Remove a JBOD expansion unit. Follow the same procedure that you use in a noncluster environment.	*Sun StorEdge 3000 Family Installation, Operation, and Service Manual*
Replace a chassis	"How to Replace the Chassis" on page 43
Replace a SCSI cable	"How to Replace a SCSI Cable" on page 48
Replace a host adapter	"Replacing a Host Adapter" on page 49
Add a disk drive	"How to Add a Disk Drive" on page 61
Remove a disk drive	"How to Remove a Disk Drive" on page 62
Replace a disk drive	"How to Replace a Disk Drive *Without* Oracle Real Application Clusters" on page 64 "How to Replace a Disk Drive *With* Oracle Real Application Clusters" on page 67
Upgrade disk drive firmware	"How to Upgrade Disk Drive Firmware" on page 70
Upgrade host adapter firmware	"How to Upgrade Host Adapter Firmware" on page 71

FRUs That *Do Not* Require Sun Cluster Maintenance Procedures

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

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

Sun StorEdge 3120 Storage Array FRUs

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

- Air management blank sled
- Fan module
- Power supply

Sun StorEdge 3310 and 3320 SCSI Storage Array FRUs

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

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

SPARC: Sun StorEdge D1000 Storage Array FRUs

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

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

SPARC: Sun StorEdge Netra D130/S1 Storage Array FRUs

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

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

SPARC: Sun StorEdge D2 Storage Array FRUs

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

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

SPARC: Sun StorEdge Multipack Storage Array FRUs

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

- Backplane
- Power supply

▼ How to Remove a Storage Array

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

Before You Begin

This procedure relies on the following prerequisites and assumptions.

- You plan to remove the references to the disk drives in the array.
- You do not need to *replace* the storage array's chassis.

  If you need to replace your storage array's chassis, see "How to Replace the Chassis" on page 43.

- Your nodes are not configured with dynamic reconfiguration functionality.

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

Steps

1. **Does the storage array that you want to remove contain a quorum device?**

   ```
   # scstat -q
   ```

   - **If no, proceed to Step 2.**

   - **If yes, add a new quorum device that will not be impacted by this procedure. Remove the old quorum device.**

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

2. **If necessary, back up the metadevice or volume.**

   For more information, see your Solstice DiskSuite/Solaris Volume Manager or VERITAS Volume Manager documentation.

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

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

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

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

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

   ```
   # cfgadm -al
   ```

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

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

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

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

   ```
   # devfsadm -C
   ```

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

   ```
   # scdidadm -C
   ```

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

   For the procedure about how to power off a storage array, see your storage documentation. For a list of storage documentation, see "Related Documentation" on page 6.

10. **Remove the storage array.**

    For the procedure about how to remove a storage array, see your storage documentation. For a list of storage documentation, see "Related Documentation" on page 6.

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

    Note – If other parallel SCSI devices are connected to the nodes, you can delete the contents of the nvramrc script. Then, at the OpenBoot PROM, set setenv use-nvramrc? false. Afterward, reset the scsi-initiator-id to 7 as outlined in "Installing a Storage Array" on page 13.

    ---

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

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

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

    ```
    # scdidadm -L
    ```

▼ How to Replace the Chassis

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

If you need to replace the entire storage array, see "How to Remove a Storage Array" on page 41.

Before You Begin
This procedure relies on the following assumptions.

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

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

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

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

Steps
1. **Does the chassis contain a disk drive that is a quorum device?**

   `# scstat -q`

   - **If no, proceed to Step 2.**

   - **If yes, add a new quorum device that will not be impacted by this procedure. Remove the old quorum device.**

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

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

   For more information, see your Solstice DiskSuite/Solaris Volume Manager or VERITAS Volume Manager documentation.

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

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

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

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

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

   You can disconnect the cables in any order.

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

   For more information, see your storage documentation. For a list of storage documentation, see "Related Documentation" on page 6.

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

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

   You can connect the cables in any order.

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

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

   Move your other components to the new chassis as well. For the procedures about how to replace your storage array's components, see your storage documentation. For a list of storage documentation, see "Related Documentation" on page 6.

9. **Power on the storage array.**

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

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

    ```
    # devfsadm
    ```

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

    ```
    # scgdevs
    ```

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

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

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

    For more information, see your Solstice DiskSuite/Solaris Volume Manager or VERITAS Volume Manager documentation.

How to Disconnect and Reconnect a Node from Shared Storage

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

This procedure relies on the following assumptions.

- You intend to reconnect the shared storage to the same host adapter that the storage array was connected before you disconnected the SCSI cable.

- This procedure defines Node A as the node you want to disconnect from the shared storage. Node B is the remaining node.

- This procedure assumes that you might have more than one storage array connected to Node A.

▼ Disconnect the Node from Shared Storage

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

Note – To determine the specific terminator that your storage array supports, see your storage documentation. For a list of storage documentation, see "Related Documentation" on page 6.

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 6 to return resource groups and device groups to these nodes.

   ```
   # scstat
   ```

2. **Identify the submirrors on the storage array that is connected to `Node A`.**

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

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

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

   For more information, see your Solstice DiskSuite/Solaris Volume Manager or VERITAS Volume Manager documentation.

5. **Shut down `Node A`.**

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

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

   ■ **Disconnect the SCSI cable between `Node A` and the storage array. Attach this SCSI cable to `Node B` on the storage array. For an illustration, see Figure 2–1.**
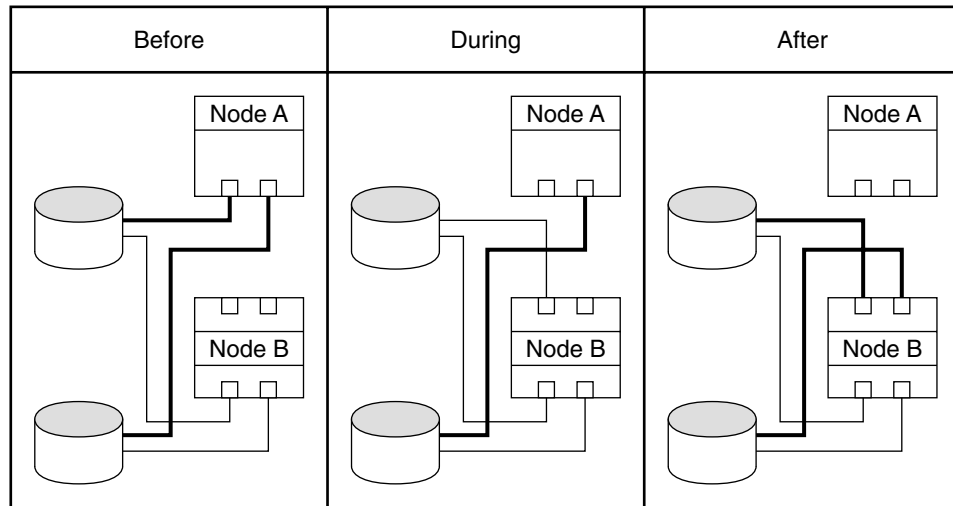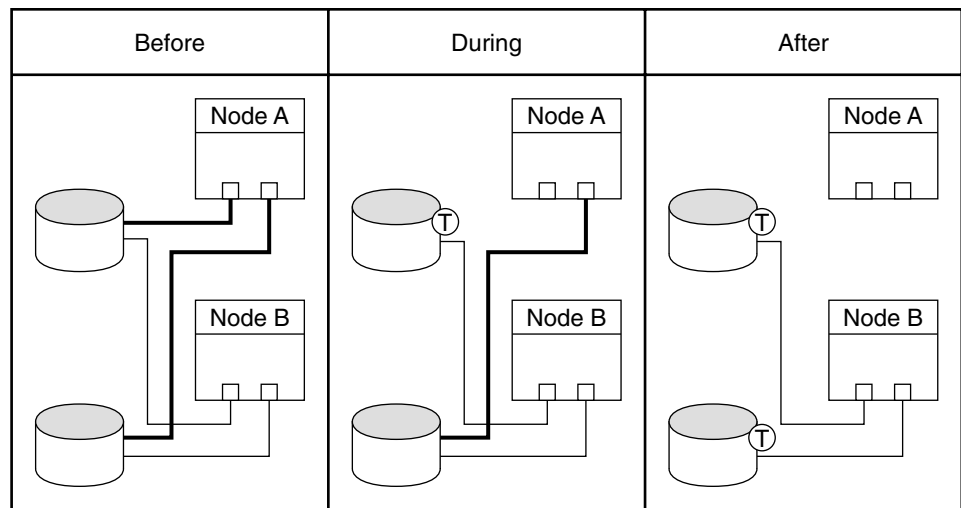

   FIGURE 2–1 Disconnecting the Node from Shared Storage by Using Host Adapters on `Node B`

   ■ **Disconnect the SCSI cable between `Node A` and the storage array. Install an appropriate SCSI terminator to this SCSI connector on the storage array. For an illustration, see Figure 2–2.**

(T) Terminator

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

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

8. **If you detached the submirrors in Step 4, reattach the submirrors. Wait for the submirrors to resynchronize.**

   For more information, see your Solstice DiskSuite/Solaris Volume Manager or VERITAS Volume Manager documentation.

9. **Repeat Step 4 through Step 8 for each remaining storage array that is connected to Node A.**

10. **Reconnect the node to the shared storage as outlined in "Reconnect the Node to Shared Storage" on page 47.**

▼ Reconnect the Node to Shared Storage

Steps   1. **(Optional) If necessary, detach the submirrors on the storage array.**

   For more information, see your Solstice DiskSuite/Solaris Volume Manager or VERITAS Volume Manager documentation.

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

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

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

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

   For more information, see your Solstice DiskSuite/Solaris Volume Manager or VERITAS Volume Manager documentation.

5. **Repeat Step 1 through Step 4 for the remaining storage array that you want to reconnect to `Node A`.**

6. **(Optional) Return the resource groups and device groups that you identified in Step 1 to `Node A`.**

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

▼ How to Replace a SCSI Cable

Use this procedure to replace a failed SCSI cable in a running cluster.

Steps 1. **Is one of the disks in the storage array a quorum device? This array is the array that is attached to the cable you are replacing.**

   ```
   # scstat -q
   ```

   - If no, proceed to Step 2.
   - If yes, relocate that quorum device to another suitable storage array

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

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

   For more information, see your Solstice DiskSuite/Solaris Volume Manager or VERITAS Volume Manager documentation.

3. **Replace the SCSI cable.**

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

   For more information, see your Solstice DiskSuite/Solaris Volume Manager or VERITAS Volume Manager documentation.

5. **(Optional) Did you relocate the quorum device in Step 1?**

   - If no, this step completes the procedure.

- If yes, relocate the quorum device function to this storage array.

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

Replacing a Host Adapter

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

Choose the procedure that corresponds to your cluster configuration.

---

x86 only – If your cluster is x86 based, Oracle RAC services are not supported. Follow the instructions in "How to Replace a Host Adapter When Using Failover and Scalable Data Services Only" on page 50.

Cluster Configuration	Instructions
Sun Cluster failover and scalable data services only, using the recommended HBA configuration	"How to Replace a Host Adapter When Using Failover and Scalable Data Services Only" on page 50
Oracle Parallel Server/Real Application Clusters (OPS/RAC) only, using the recommended HBA configuration	"SPARC: How to Replace a Host Adapter When Using Oracle Real Application Clusters Only" on page 52
Both failover and scalable data services and OPS/RAC, using the recommended HBA configuration	"SPARC: How to Replace a Host Adapter When Using Both Failover and Scalable Data Services and Oracle Real Application Clusters" on page 55
All clusters using a single, dual-port HBA to provide both paths to shared data	"How to Replace a Host Adapter When Using a Single, Dual-Port HBA to Provide Both Paths to Shared Data" on page 59

Note – The first three procedures in this section assume that you are using the recommended HBA configuration: two redundant hardware paths to shared data. If you choose to use a single HBA configuration, see "Configuring Cluster Nodes With a Single, Dual-Port HBA" in *Sun Cluster 3.0-3.1 Hardware Administration Manual for Solaris OS* for the risks and restrictions of that configuration and use "How to Replace a Host Adapter When Using a Single, Dual-Port HBA to Provide Both Paths to Shared Data" on page 59.

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

Before You Begin

This procedure relies on the following prerequisites and assumptions.

- You are *not* using Oracle Real Application Clusters. If you are using Oracle Real Application Clusters, follow the instructions in "SPARC: How to Replace a Host Adapter When Using Oracle Real Application Clusters Only" on page 52 or "SPARC: How to Replace a Host Adapter When Using Both Failover and Scalable Data Services and Oracle Real Application Clusters" on page 55.

- Except for the failed host adapter, your cluster is fully operational and all nodes are powered on.

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

Steps

1. **Does your cluster use Oracle Real Application Clusters?**

   - If yes, follow the instructions in "SPARC: How to Replace a Host Adapter When Using Oracle Real Application Clusters Only" on page 52 or "SPARC: How to Replace a Host Adapter When Using Both Failover and Scalable Data Services and Oracle Real Application Clusters" on page 55.
   - If no, proceed to Step 2.

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

   Record this information because you use this information in Step 17 of this procedure to return resource groups and device groups to `Node A`.

   ```
   # scstat
   ```

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

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

4. **Determine if the storage device that is attached to the failed host adapter contains a quorum device.**

   ```
   # scstat -q
   ```

   - If no, skip this step.
   - If yes, add a new quorum device on a different storage device and remove the old quorum device.

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

5. **Detach the Solstice DiskSuite/Solaris Volume Manager submirrors or VERITAS Volume Manager plexes on the storage array.**

   For more information, see your Solstice DiskSuite/Solaris Volume Manager or VERITAS Volume Manager documentation.

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

   Record this information because you use it in Step 16 of this procedure to reattach submirrors on the storage array. For more information about how to determine which submirrors or plexes are affected, see your Solstice DiskSuite/Solaris Volume Manager or VERITAS Volume Manager documentation.

7. **Is `Node A` enabled with the Solaris DR feature?**

   - If yes, perform any DR-specific steps and skip to Step 10.
   - If no, proceed to Step 8.

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

8. **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.

9. **Power off `Node A`.**

10. **Replace the failed host adapter.**

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

11. **Is `Node A` enabled with the Solaris DR feature?**

    - If yes, perform any DR-specific steps and skip to Step 15.
    - If no, proceed to Step 12.

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

12. **Power on `Node A`.**

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

    For instructions on setting SCSI initiator IDs in x86 based systems, see "x86: How to Install a Storage Array in a New X86 Based Cluster" on page 20.

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

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

15. **If necessary, 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`.

16. **Reattach the Solstice DiskSuite/Solaris Volume Manager submirrors or
    VERITAS Volume Manager plexes on the storage array.**

    For more information, see your Solstice DiskSuite/Solaris Volume Manager or
    VERITAS Volume Manager documentation.

17. **(Optional) If necessary, return the resource groups and device groups you
    identified in Step 2 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.

18. **Did you relocate the quorum device in Step 4?**

    - If no, this step completes the procedure.
    - If yes, and you want the cluster configured as it was before replacing the host
      adapter, relocate the quorum device function back to this storage array.

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

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

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

- You are using Oracle Real Application Clusters *only*. If you are using other failover
  or scalable data services, follow the instructions in "How to Replace a Host
  Adapter When Using Failover and Scalable Data Services Only" on page 50 or
  "SPARC: How to Replace a Host Adapter When Using Both Failover and Scalable
  Data Services and Oracle Real Application Clusters" on page 55.

- Except for the failed host adapter, your cluster is fully operational and all nodes are
  powered on.

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

Steps 1. **Does your cluster use failover or scalable data services (other than Oracle Real
    Application Clusters)?**

- If yes, follow the instructions in "How to Replace a Host Adapter When Using Failover and Scalable Data Services Only" on page 50 or "SPARC: How to Replace a Host Adapter When Using Both Failover and Scalable Data Services and Oracle Real Application Clusters" on page 55.
- If no, proceed to Step 2.

2. **Determine the Oracle instance that is running on `Node A`.**

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

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

   For information about how to shut down and restart an Oracle instance in the RAC environment, refer to your Oracle documentation.

4. **Determine if the storage device that are attached to the failed host adapter contains a quorum device.**

   # **`scstat -q`**

   - If no, skip to Step 5.
   - If yes, add a new quorum device on a different storage device and remove the old quorum device.

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

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

   For more information, see your VERITAS Volume Manager documentation.

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

   Record this information because you use it in Step 15 of this procedure to reattach plexes on the storage array. For more information about how to determine which plexes are affected, see your VERITAS Volume Manager documentation.

7. **Is `Node A` enabled with the Solaris dynamic reconfiguration (DR) feature?**

   - If yes, perform any DR-specific steps and skip to Step 10.
   - If no, proceed to Step 8.

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

8. **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.

9. **Power off `Node A.`**

10. **Replace the failed host adapter.**

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

11. **Is `Node A` enabled with the Solaris dynamic reconfiguration feature?**

    - If yes, perform any DR-specific steps and skip to Step 14.
    - If no, proceed to Step 12.

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

12. **Power on `Node A`.**

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

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

14. **If necessary, 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`.

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

    For more information, see your VERITAS Volume Manager documentation.

16. **(Optional) Bring your Oracle Real Application Clusters instance online. This is the instance you identified in Step 2.**

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

17. **(Optional) Did you relocate a quorum device in Step 4?**

    - If no, this step completes the procedure.
    - If yes, and you want your configuration to use the same quorum structure after the host adapter replacement, relocate the quorum device function back to this storage array.

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

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

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

- You are using both Oracle Real Application Clusters and failover or scalable data services.

  If you are using only Oracle Real Application Clusters, follow the instructions in "SPARC: How to Replace a Host Adapter When Using Oracle Real Application Clusters Only" on page 52. If you are using only failover and scalable data services, follow the instructions in "How to Replace a Host Adapter When Using Failover and Scalable Data Services Only" on page 50.

- Except for the failed host adapter, your cluster is operational and all nodes are powered on.

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

Steps   1. **Does your cluster use both failover or scalable data services and Oracle Real Application Clusters?**

   - If no, follow the instructions in "SPARC: How to Replace a Host Adapter When Using Oracle Real Application Clusters Only" on page 52 or "How to Replace a Host Adapter When Using Failover and Scalable Data Services Only" on page 50.
   - If yes, proceed to Step 2.

2. **Determine the Oracle instance that is running on `Node A`.**

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

3. **Shut down the Oracle Real Application Clusters instance identified in Step 2.**

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

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

   Record this information because you use it in Step 18 of this procedure to return resource groups and device groups to `Node A`.

   ```
   # scstat
   ```

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

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

6. **Determine if the storage device that is connected to the failed host adapter contains a quorum device.**

   ```
   # scstat -q
   ```

7. **If Node A contains a quorum device, add a new quorum device on a different storage device and remove the old quorum device.**

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

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

   For more information, see your VERITAS Volume Manager documentation.

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

   Record this information because you use it in Step 18 of this procedure to reattach plexes on the storage array. For more information about how to determine which plexes are affected, see your VERITAS Volume Manager documentation.

10. **Is Node A enabled with the Solaris dynamic reconfiguration (DR) feature?**

    - If yes, perform any DR-specific steps and skip to Step 13.
    - If no, proceed to Step 11.

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

11. **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.

12. **Power off Node A.**

13. **Replace the failed host adapter.**

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

14. **Is Node A enabled with the Solaris DR feature?**

    - If yes, perform any DR-specific steps and skip to Step 17.
    - If no, proceed to Step 15.

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

15. **Power on Node A.**

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

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

17. **If necessary, 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`.

18. **Reattach the VERITAS Volume Manager plexes on the storage array to their respective volumes.**

    For more information, see your VERITAS Volume Manager documentation.

19. **(Optional) If necessary, return the resource groups and device groups you identified in Step 5 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.

20. **(Optional) Bring your Oracle Real Application Clusters instance online. This is the instance you identified in Step 2.**

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

21. **(Optional) If you relocated a quorum device in Step 7 and you want your configuration to use the same quorum structure after the host adapter replacement, relocate the quorum device function back to this storage array.**

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

Example 2–1   Replacing a Host Adapter in a Running Cluster

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

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

```
.
.
oracle 14434     1  0 13:05:43 ?        0:00 ora_pmon_tpcc1
oracle 14458     1  0 13:05:50 ?        0:00 ora_d000_tpcc1
```

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

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

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

SQL*Plus: Release 9.2.0.4.0 - Production on Mon Jan 5 14:12:28 2004

Copyright (c) 1982, 2002, Oracle Corporation.  All rights reserved.


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

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

SQL> exit
Disconnected from Oracle9i Enterprise Edition Release 9.2.0.4.0 - 64bit
Production With the Partitioning, Real Application Clusters, OLAP and
Oracle Data Mining options
JServer Release 9.2.0.4.0 - Production
$ lsnrctl

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

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

Welcome to LSNRCTL, type "help" for information.

LSNRCTL> stop
Connecting to (DESCRIPTION=(ADDRESS=(PROTOCOL=IPC)(KEY=EXTPROC)))
The command completed successfully
LSNRCTL>
```

1. After you have stopped the Oracle Real Application Clusters instance, check for a
   quorum device and, if necessary, reconfigure the quorum device.

2. When you are certain that the node with the failed adapter does not contain a
   quorum device, proceed to determine the affected plexes.

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

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

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

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

1. After the plexes are detached, you can safely shut down the node, if necessary.

2. Proceed with replacing the failed host adapter, following instructions that accompanied that device.

3. After you replace the failed host adapter, and Node A is in cluster mode, reattach the plexes and replace any quorum device to reestablish your original cluster configuration.

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

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

■ You are using a single, dual-port HBA to provide the redundant paths to your shared data.

   For an explanation of the limitations and risks of using this configuration, see "Configuring Cluster Nodes With a Single, Dual-Port HBA" in *Sun Cluster 3.0-3.1 Hardware Administration Manual for Solaris OS*.

■ Except for the failed host adapter, your cluster is operational and all nodes are powered on.

■ Your nodes are not configured for dynamic reconfiguration.

   You cannot use dynamic reconfiguration for this procedure when using a single HBA.

■ Node A is the node with the failed host adapter.

Steps 1. If you are using scalable and failover services, determine the resource groups and device groups that are running on `Node A`.

Record this information because you use it in Step 18 of this procedure to return resource groups and device groups to Node A.

```
# scstat
```

2. Record the details of metadevices that are affected by the failed host adapter.

3. SPARC: If you are using Oracle Real Application Clusters, shut down all RAC instances running in your cluster.

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

4. Shut down the cluster.

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

5. Power off `Node A`.

6. Replace the failed host adapter.

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

7. Power on `Node A`.

8. Boot all nodes into cluster mode.

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

9. If necessary, 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`.

10. Perform any volume management maintenance procedures that are necessary to fix any metadevices affected by this procedure.

For more information, see your volume manager software documentation.

11. (Optional) If necessary, return the resource groups and device groups you identified in Step 5 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.

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

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

▼ How to Add a Disk Drive

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

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

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

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

- Your storage array contains an empty disk slot.
- Your nodes are not configured with dynamic reconfiguration functionality.

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

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

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

2. Install the disk drive.

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

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

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

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

   `# devfsadm`

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

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

   `# ls -l /dev/rdsk`

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

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

   If a volume management daemon such as vold is running on your node, and you have a CD-ROM drive that is connected to the node, a device busy error might be returned even if no disk is in the drive. This error is an expected behavior.

   `# scgdevs`

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

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

   `# scdidadm -L`

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

   For more information, see your Solstice DiskSuite/Solaris Volume Manager or VERITAS Volume Manager documentation.

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

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

▼ How to Remove a Disk Drive

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

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

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

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

- You do not need to remove the entire storage array.

  If you need to remove the storage array, see "How to Remove a Storage Array" on page 41.

- You do not need to replace the storage array's chassis.

  If you need to replace your storage array's chassis, see "How to Replace the Chassis" on page 43.

- Your nodes are not configured with dynamic reconfiguration functionality.

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

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

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

   ```
   # scdidadm -l deviceID
   # cfgadm -al
   ```

2. **Determine if the disk you want to remove is a quorum device.**

   ```
   # scstat -q
   ```

   - **If no, skip this step.**

   - **If yes, add a new quorum device that will not be affected by this procedure. Remove the old quorum device.**

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

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

   For more information, see your Solstice DiskSuite/Solaris Volume Manager or VERITAS Volume Manager documentation.

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

   For more information, see your Solstice DiskSuite/Solaris Volume Manager or VERITAS Volume Manager documentation.

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

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

6. **Remove the disk drive.**

   For the procedure about how to remove a disk drive, see your storage
   documentation. For a list of storage documentation, see "Related Documentation"
   on page 6.

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

   ```
   # devfsadm -C
   ```

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

   ```
   # scdidadm -C
   ```

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

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

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

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

- Your cluster is operational.
- Your system is *not* running Oracle Real Application Clusters.

  If your system is running Oracle Real Application Clusters, see "How to Replace a
  Disk Drive *With* Oracle Real Application Clusters" on page 67.
- *(Solstice DiskSuite/Solaris Volume Manager Only)* If the disk drive failure prevents
  Solstice DiskSuite/Solaris Volume Manager from reading the disk label, you have a
  backup of the disk-partitioning information.
- Your nodes are not configured with dynamic reconfiguration functionality.

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

Steps   1. **Identify the failed disk drive.**

   ```
   # scdidadm -o diskid -l cNtXdY
   ```

2. **Is the disk drive that you want to remove a quorum device?**

   ```
   # scstat -q
   ```

   - If no, proceed to the next step.
   - If yes, add a new quorum device on a different storage device. Remove the old
     quorum device.

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

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

   For more information, see your Solstice DiskSuite/Solaris Volume Manager or VERITAS Volume Manager documentation.

4. *(Solstice DiskSuite/Solaris Volume Manager Only)* **Does your disk drive failure prevent Solstice DiskSuite/Solaris Volume Manager from reading the disk label?**

   - If yes, use the disk-partitioning information that you saved. You saved this disk-partitioning information when you performed one of the following tasks.

     - Installed your storage array in an initial cluster as outlined in Step 10 of "SPARC: How to Install a Storage Array in a New SPARC Based Cluster" on page 15
     - Added the storage array to an operational cluster.

   - If no, save the disk-partitioning information now if you have not already done so. Use this information when you partition the new disk drive.

   ```
   # prtvtoc /dev/rdsk/cNtXdYsZ > filename
   ```

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

   ---

5. **Replace the failed disk.**

   a. **Determine which node owns the device group.**

      ```
      # scstat -D
      ```

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

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

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

      ```
      # cfgadm -x replace_device cX::disk/cXtYdZ
      ```

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

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

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

e. Switch the primary for the device group.

```
# scswitch -z -D disk-device-group -h node
```

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

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

g. Suspend activity on the device group's original primary.

```
# cfgadm -x replace_device cX::disk/cXtYdZ
```

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

h. When prompted to proceed with the disk replacement, type y to proceed.

Warning messages might be displayed. Ignore these messages.

6. **On all nodes that are attached to the storage device, run the `devfsadm`(1M) command to probe all devices and to write the new disk drive to the `/dev/rdsk` directory.**

```
# devfsadm
```

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

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

8. *(Solstice DiskSuite/Solaris Volume Manager Only)* **From any node that is connected to the storage device, partition the new disk drive by using the partitioning you saved when you installed or added the storage array.**

```
# fmthard -s filename /dev/rdsk/cNtXdYsZ
```

9. **From all nodes that are connected to the disk drive, update the DID database and driver.**

```
# scdidadm -R deviceID
```

Note – After running scdidadm –R on the first node, each subsequent node that you run the command on might display a warning. Ignore this warning.

---

10. **From any node, confirm that the failed disk drive has been replaced by comparing the new physical DID to the physical DID that was identified at the beginning of this procedure.**

If the new physical DID is different, you successfully replaced the failed disk drive with a new disk drive.

```
# scdidadm -o diskid -l cNtXdY
```

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

    For more information, see your Solstice DiskSuite/Solaris Volume Manager or VERITAS Volume Manager documentation.

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

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

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

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

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

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

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

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

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

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

  If your system is *not* running Oracle Real Application Clusters, see "How to Replace a Disk Drive *Without* Oracle Real Application Clusters" on page 64.

- Your nodes are not configured with dynamic reconfiguration functionality.

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

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

    # **vxdisk list**

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

   You use this instance number in Step 13.

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

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

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

   If you are using Solaris 10, the `node@2` portion of this output is not present. Solaris
   10 does not add this prefix for cluster nodes.

   ---

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

   You will use the device ID number in Step 15 and Step 16.

   ```
   # scdidadm -L | grep cXtYdZ
   ```

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

   You will use the disk ID number in Step 17.

   ```
   # scdidadm -o diskid -l cXtYdZ
   ```

5. **Determine if the disk drive that you want to replace is a quorum device.**

   ```
   # scstat -q
   ```

   - **If the disk drive is *not* a quorum device, skip this step.**

   - **If the disk drive that you want to replace is a quorum device, add a new
     quorum device on a different storage device and remove the old quorum
     device.**

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

6. **On each node, remove the failed disk from VERITAS Volume Manager control.**

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

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

   ```
   # vxdisk list
   ```

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

   For the procedure about how to remove a disk drive, see your storage documentation. For a list of storage documentation, see "Related Documentation" on page 6.

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

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

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

    ```
    # devfsadm -C
    ```

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

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

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

    For the procedure about how to add a disk drive, see your storage documentation. For a list of storage documentation, see "Related Documentation" on page 6.

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

    Use the instance number that you identified in Step 2.

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

14. **Verify that you added the disk.**

    ```
    # ls /dev/dsk/cXtYdZ
    ```

15. **On one node, update the device ID numbers.**

    Use the device ID number that you identified in Step 3.

    ```
    # scdidadm -R DID_number
    ```

16. **Verify that you updated the disk's device ID number.**

    Use the device ID number that you identified in Step 3.

    ```
    # scdidadm -L | grep DID_number
    ```

17. **Verify that the disk ID number is different than the disk ID number that you identified in Step 4.**

18. **On each node, add the new disk to the VERITAS Volume Manager database.**

    ```
    # vxdctl enable
    ```

19. **Verify that you added the new disk.**

    ```
    # vxdisk list |grep cXtYdZ
    ```

20. **Determine the master node.**

    ```
    # vxdctl -c mode
    ```

21. Perform disk recovery tasks on the master node.

Depending on your configuration and volume layout, select the appropriate VERITAS Volume Manager menu item to recover the failed disk.

```
# vxdiskadm
# vxtask list
```

▼ How to Upgrade Disk Drive Firmware

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

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

- Your cluster is operational.
- Your nodes are not configured with dynamic reconfiguration functionality.

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

Steps **1. Is the disk drive on which you want to upgrade the firmware a quorum device?**

```
# scstat -q
```

- **If no, proceed to Step 2.**

- **If yes and you need a quorum device to maintain quorum, choose another disk drive to configure as the quorum device. Then remove the old quorum device.**

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

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

For more information, see your Solstice DiskSuite/Solaris Volume Manager or VERITAS Volume Manager documentation.

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

For more information, see your Solstice DiskSuite/Solaris Volume Manager or VERITAS Volume Manager documentation.

4. Install the firmware.

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

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

   For more information, see your Solstice DiskSuite/Solaris Volume Manager or VERITAS Volume Manager documentation.

▼ How to Upgrade Host Adapter Firmware

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

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

- Your cluster is operational.
- This procedure defines `Node A` as the node with which you begin working. `Node B` is the remaining node.

Steps   1. **Determine the resource groups and device groups that are running on the node on which you plan to upgrade the host adapter firmware.**

   ```
   # scstat
   ```

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

   ```
   # scswitch -S -h from-node
   ```

3. **Perform the firmware upgrade.**

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

4. **(Optional) Return the resource groups and device groups to `Node B`.**

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

Cabling Diagrams

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

- Sun StorEdge 3120 Storage Array
- Sun StorEdge 3310 SCSI Storage Array
- SPARC: Sun StorEdge D1000 Storage Array
- SPARC: Sun StorEdge Multipack Storage Array
- SPARC: Sun StorEdge Netra D130/S1 Storage Array
- SPARC: Sun StorEdge D2 Storage Array

Sun StorEdge 3120 Cabling Diagrams

Installing a Sun StorEdge 3120 Storage Array

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

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

Bus cable
Host cable
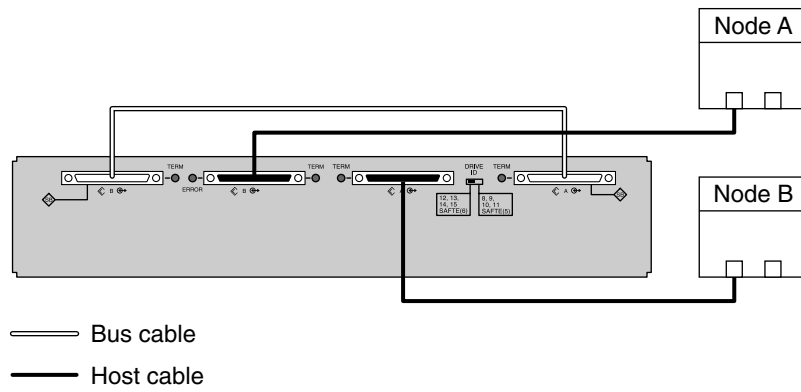
FIGURE 3–1 Cabling a Single-Bus, Multi-Initiator Sun StorEdge 3120 Storage Array
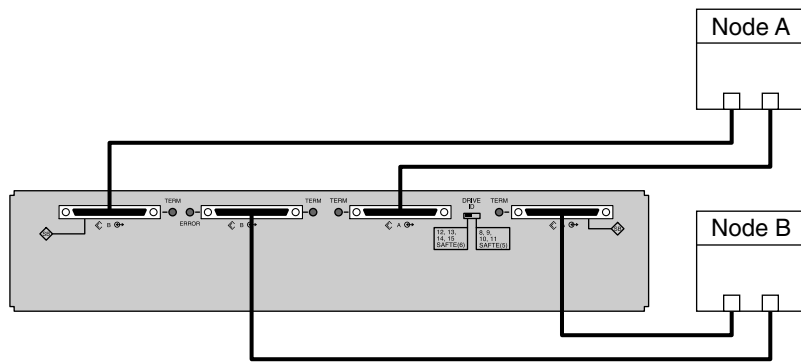
FIGURE 3–2 Cabling a Dual-Bus, Multi-Initiator Sun StorEdge 3120 Storage Array

Sun StorEdge 3310 and 3320 SCSI Cabling Diagrams

Installing a Sun StorEdge 3310 or 3320 Storage Array

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

For the procedure about how to cable the storage arrays, see the *Sun StorEdge 3000 Family Installation, Operation, and Service Manual*. For a URL to this storage documentation, see "Related Documentation" on page 6.
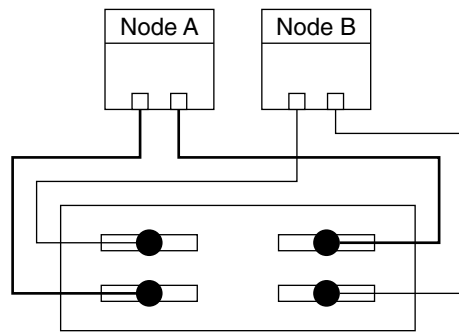
FIGURE 3–3 Cabling a Dual-Bus, Multi-Initiator Sun StorEdge 3310 or 3320 Storage Array

Adding a Sun StorEdge 3310 or 3320 Storage Array

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

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

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

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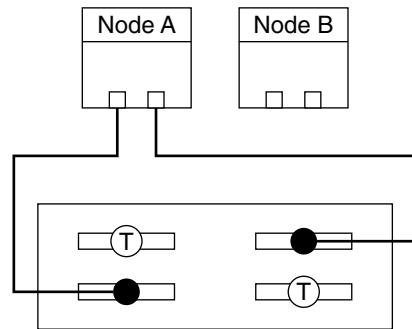
(T) Terminator

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

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

For the procedure about how to cable the storage arrays, see the *Sun StorEdge 3000 Family Installation, Operation, and Service Manual*. For a URL to this storage documentation, see "Related Documentation" on page 6.

Ensure that the cable does not exceed bus-length limitations. For more information about bus-length limitations, see the documentation that shipped with your hardware.
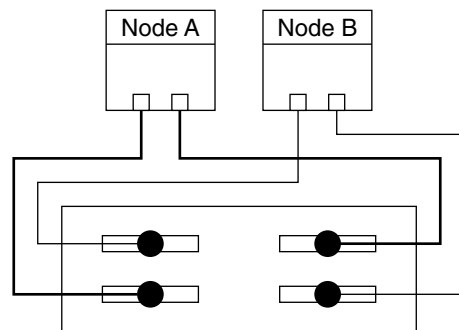
FIGURE 3–5 Cabling a Dual-Bus, Multi-Initiator Sun StorEdge 3310 or 3320 Storage Array: Part 2

SPARC: Sun StorEdge D1000 Cabling Diagrams

SPARC: Installing a Sun StorEdge D1000 Storage Array

Connect the cables to the storage arrays, as shown in Figure 3–6.

Ensure that the *entire* bus length that is connected to each storage array is less than 25 m. This measurement includes the cables to both nodes, as well as the bus length that is internal to each storage array, node, and the host adapter.
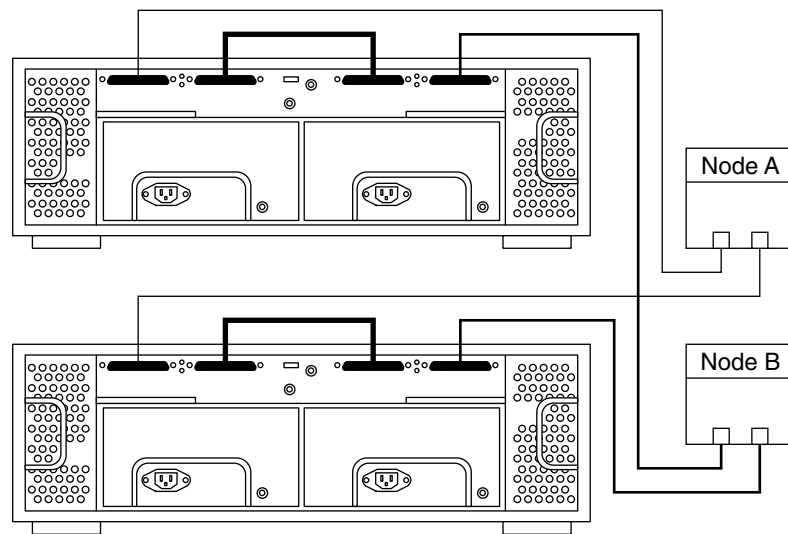
FIGURE 3–6 SPARC: Installing a Sun StorEdge D1000 Mirrored-Pair Configuration: An Example

SPARC: Adding a Sun StorEdge D1000 Storage Array

Ensure that the *entire* SCSI bus length to each enclosure is less than 25 m. This measurement includes the cables to both nodes, as well as the bus length that is internal to each enclosure, node, and host adapter. Refer to the documentation that shipped with the enclosure for other restrictions about SCSI operation.
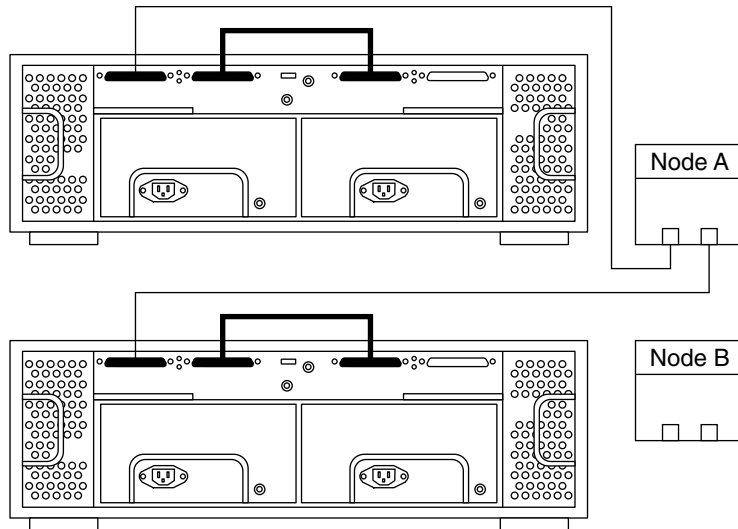
FIGURE 3–7 SPARC: Adding a Sun StorEdge D1000 Mirrored-Pair Configuration: Part 1

Connect the storage array to the host adapters by using differential SCSI cables as shown in see Figure 3–8.
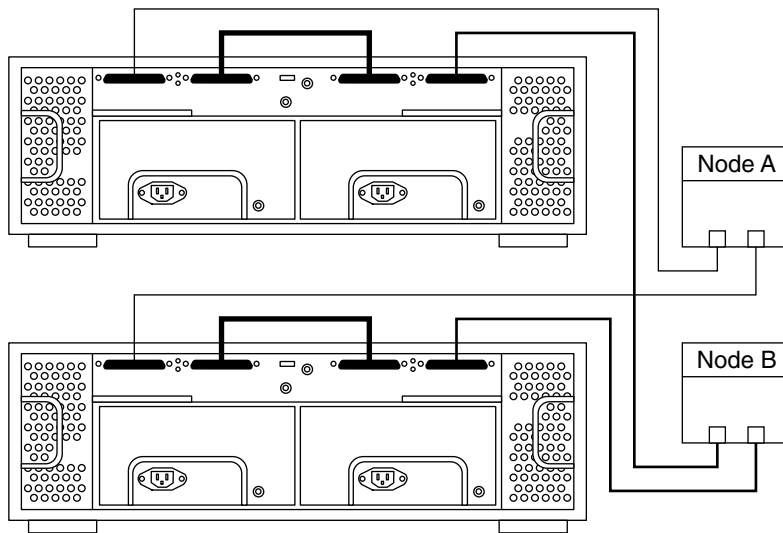
FIGURE 3–8 SPARC: Adding a Sun StorEdge D1000 Mirrored-Pair Configuration: Part 2

SPARC: Replacing a Sun StorEdge D1000 Storage Array

Disconnect the SCSI cables from the storage array. Remove the SCSI cables in the order that is shown in Figure 3–9.
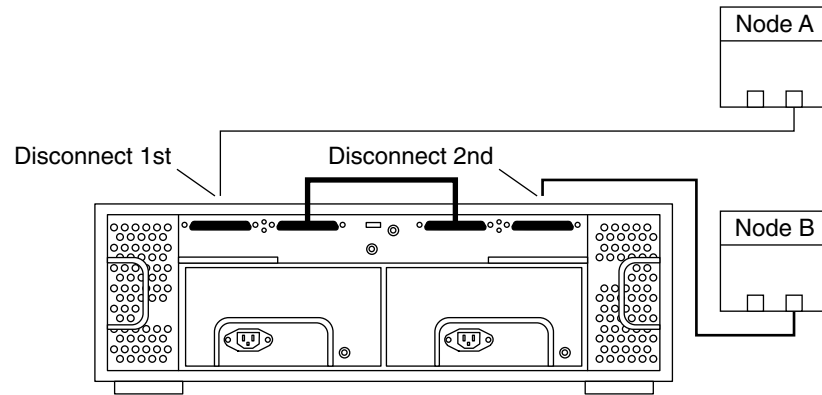
FIGURE 3–9 SPARC: Disconnecting the SCSI Cables From a Sun StorEdge D1000 Storage Array

SPARC: Removing a Sun StorEdge D1000 Storage Array

Disconnect the SCSI cables from the storage array. Remove the SCSI cables in the order that is shown in Figure 3–10.
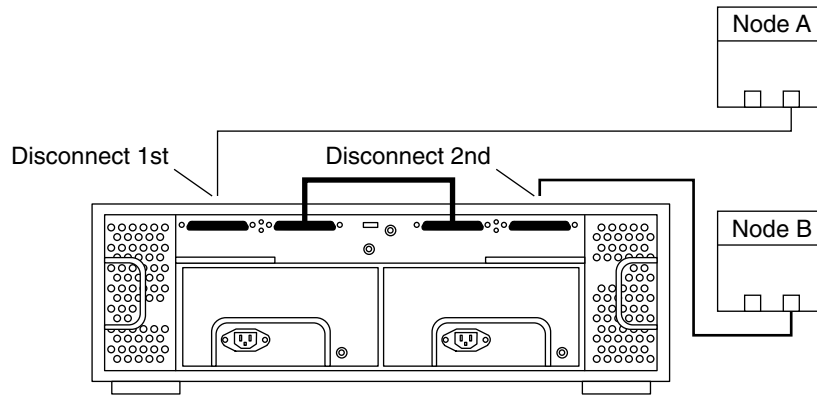
FIGURE 3–10 SPARC: Disconnecting the SCSI Cables From a Sun StorEdge D1000 Storage Array

SPARC: Sun StorEdge Multipack Cabling Diagrams

SPARC: Installing a Sun StorEdge Multipack Storage Array

Connect the cables to the storage array, as shown in Figure 3–11.

Ensure that the *entire* SCSI bus length to each storage array conforms to the distances in Table 1–1. This measurement includes the cables to both nodes, as well as the bus length that is internal to each storage array, node, and host adapter. Refer to the documentation that shipped with the storage array for other restrictions about SCSI operation.
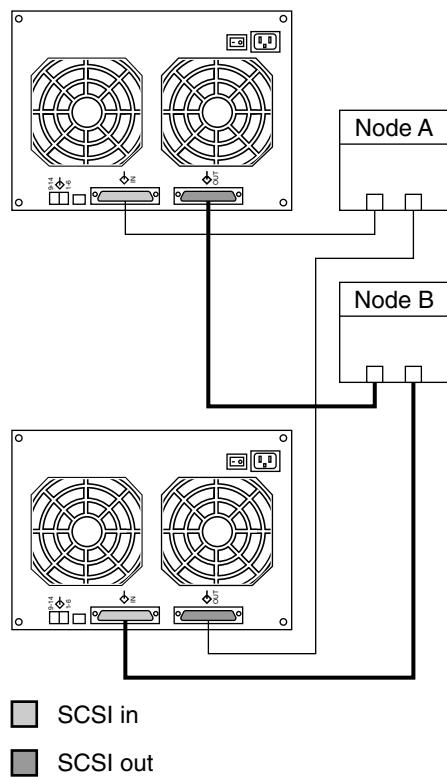
FIGURE 3–11 SPARC: Installing a Sun StorEdge Multipack Mirrored-Pair Configuration: An Example

SPARC: Adding a Sun StorEdge Multipack Storage Array

Connect the single-ended SCSI cable between the node and the storage arrays, as shown in Figure 3–12.

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

□ SCSI in

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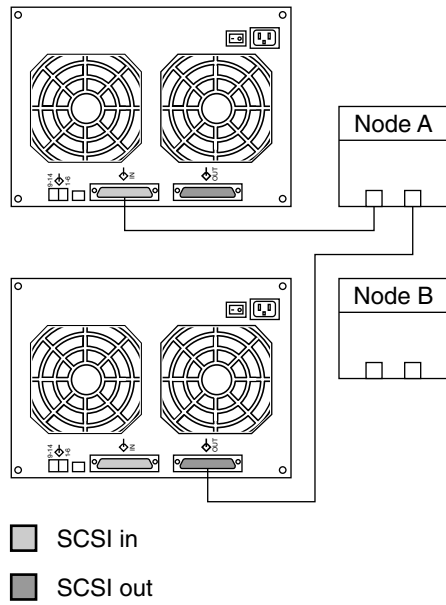
FIGURE 3–12 SPARC: Adding a Sun StorEdge Multipack Mirrored-Pair Configuration: Part 1

Connect the storage arrays to the host adapters by using single-ended SCSI cables, as shown in Figure 3–13.
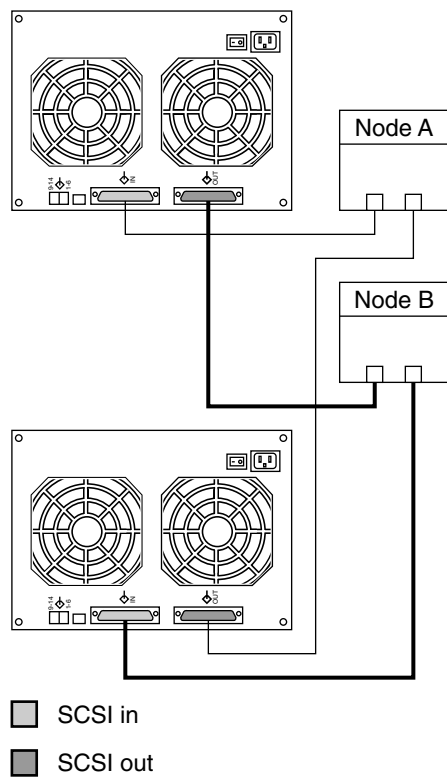
FIGURE 3–13 SPARC: Adding a Sun StorEdge Multipack Mirrored-Pair Configuration: Part 2

SPARC: Replacing a Sun StorEdge Multipack Storage Array

Disconnect the SCSI cables from the storage array. Disconnect the SCSI cable on the SCSI OUT connector, and then disconnect the cable on the SCSI IN connector, as shown in Figure 3–14.
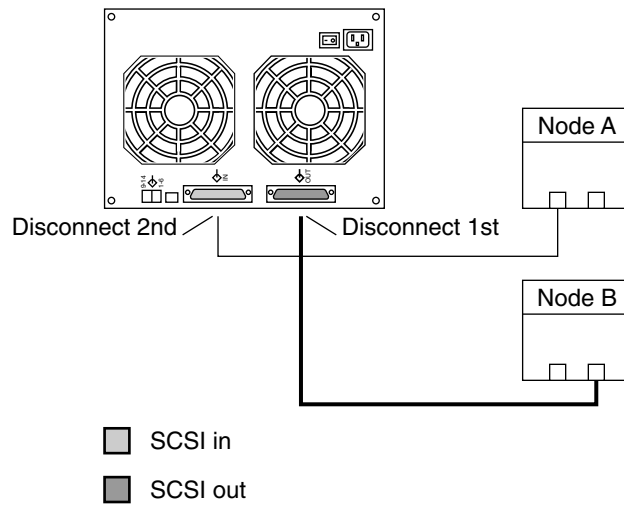
FIGURE 3–14 SPARC: Replacing a Sun StorEdge Multipack Storage Array

SPARC: Removing a Sun StorEdge Multipack Storage Array

Disconnect the SCSI cables from the storage array. Disconnect the SCSI cables in the order that is shown in Figure 3–15.

Disconnect 2nd   **Disconnect 1st**

☐ SCSI in

☐ SCSI out
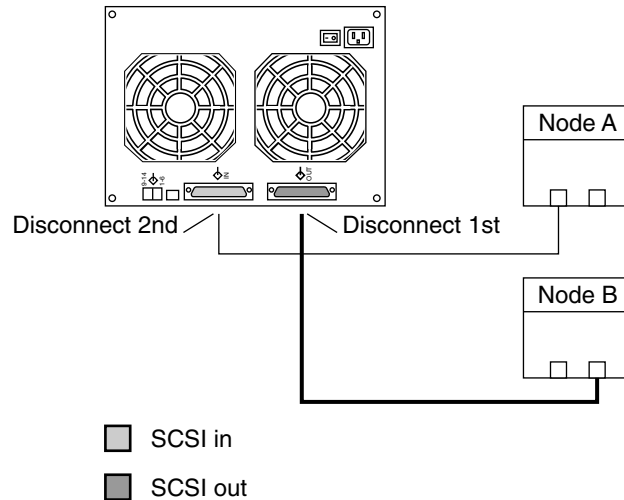
FIGURE 3–15 SPARC: Removing a Sun StorEdge Multipack Storage Array

SPARC: Sun StorEdge Netra D130/S1 Cabling Diagrams

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

If your host has only one SCSI port, see "SPARC: Single-SCSI-Port Hosts" on page 85. If your host has two SCSI ports, see "SPARC: Dual-SCSI-Port Hosts" on page 88.

SPARC: Single-SCSI-Port Hosts

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

```
          Ethernet Switch              Ethernet Switch

                  | LAN                        | LAN

          Node 1                       Node 2
          PCI                          PCI
          Interface                    Interface

                      / LAN              \ LAN
          Netra E1 Expander            Netra E1 Expander
                  SCSI                         SCSI

          SCSI          SCSI      SCSI      SCSI
          Out            In        In       Out
          Netra D130/StorEdge S1 Storage   Netra D130/StorEdge S1 Storage


              _ _ _ _   Ethernet

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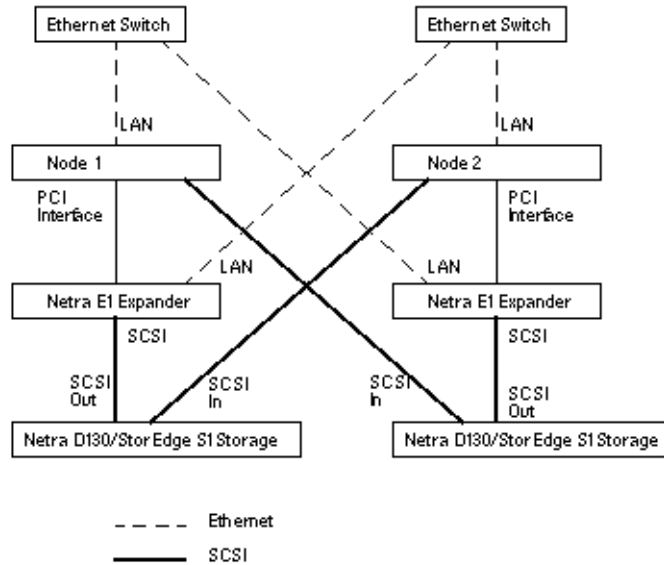
FIGURE 3–16 SPARC: Sun StorEdge Netra D130/S1 Mirrored-Pair Configuration With Single SCSI Port Hosts: An Example

To cable the storage arrays correctly, adhere to the bus-length restrictions that are listed in the documentation that shipped with your host adapter. In a Sun Cluster environment, this cable-length measurement includes the cables to both nodes. This cable-length measurement also includes the bus length that is internal to each storage array, node, and host adapter.
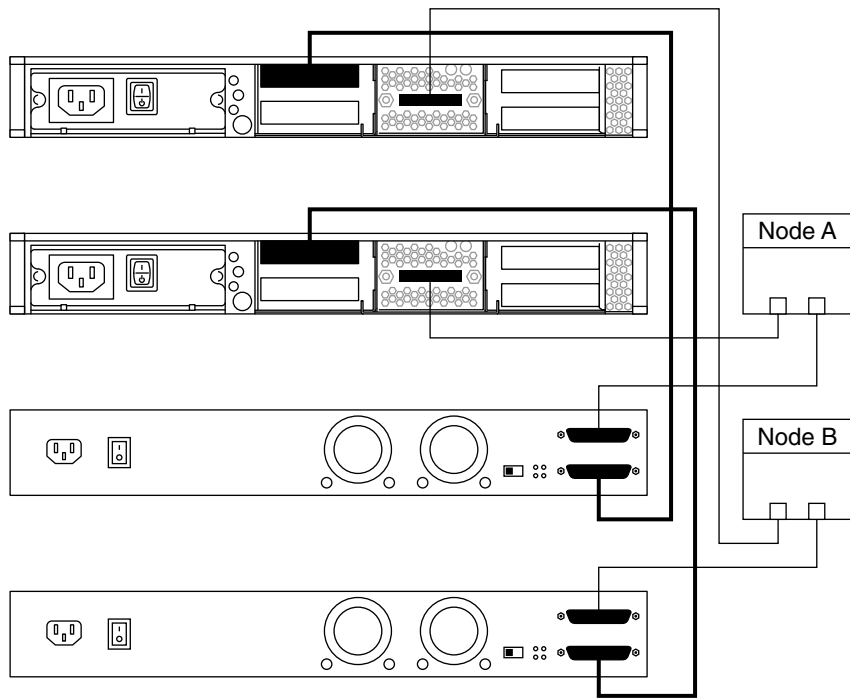
FIGURE 3–17 SPARC: Cabling a Sun StorEdge Netra D130/S1 Mirrored-Pair Configuration With Single-SCSI-Port Hosts
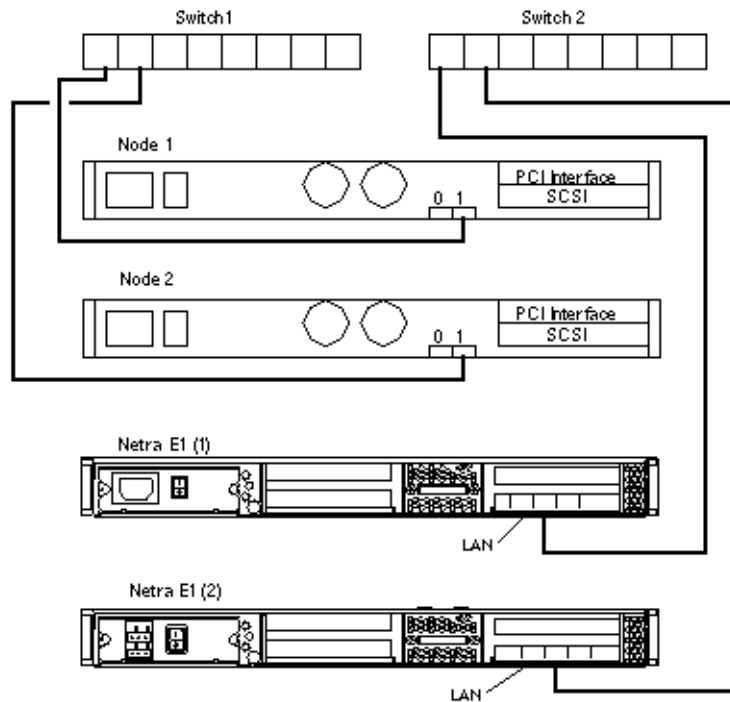
FIGURE 3–18 SPARC: Cabling a Sun StorEdge Netra D130/S1 Mirrored-Pair Configuration With Single-SCSI-Port Hosts and Netra E1 PCI Expanders

SPARC: Dual-SCSI-Port Hosts

To cable the storage arrays correctly, adhere to the bus-length restrictions that are listed in the documentation that shipped with your host adapter. In a Sun Cluster environment, this cable-length measurement includes the cables to both nodes. This cable-length measurement also includes the bus length that is internal to each storage array, node, and host adapter.
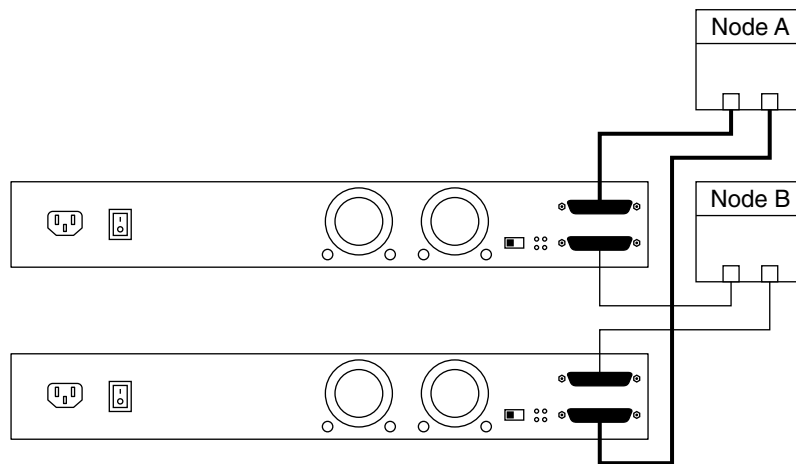
FIGURE 3–19 SPARC: Cabling a Sun StorEdge Netra D130/S1 Mirrored-Pair Configuration With Dual-SCSI-Port Hosts

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

To cable the storage arrays correctly, adhere to the bus-length restrictions that are listed in the documentation that shipped with your host adapter. In a Sun Cluster environment, this cable-length measurement includes the cables to both nodes. This cable-length measurement also includes the bus length that is internal to each storage array, node, and host adapter.
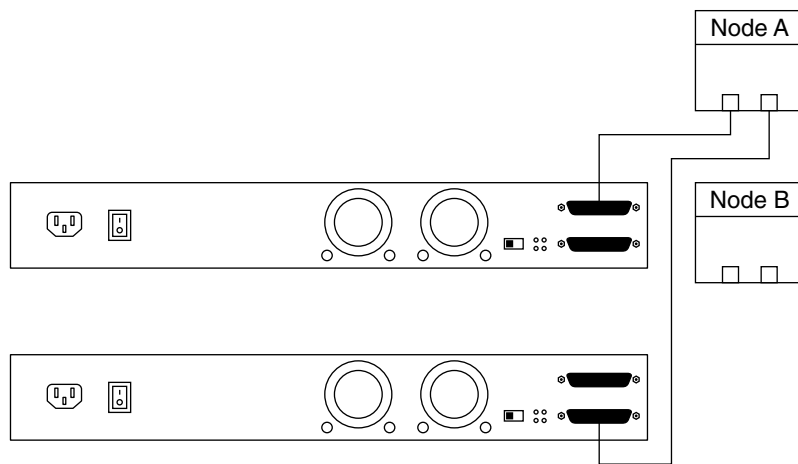
FIGURE 3–20 SPARC: Adding a Sun StorEdge Netra D130/S1 Storage Array Mirrored-Pair Configuration: Part 1
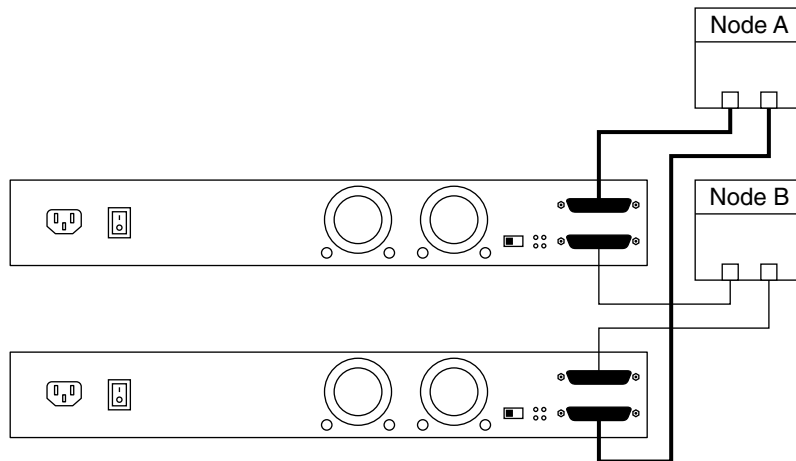
FIGURE 3–21 SPARC: Adding a Sun StorEdge Netra D130/S1 Mirrored-Pair Configuration: Part 2

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

Disconnect the SCSI cables from the storage arrays. Disconnect the cable on the SCSI OUT connector, and then disconnect the cable on the SCSI IN connector, as shown in Figure 3–22.
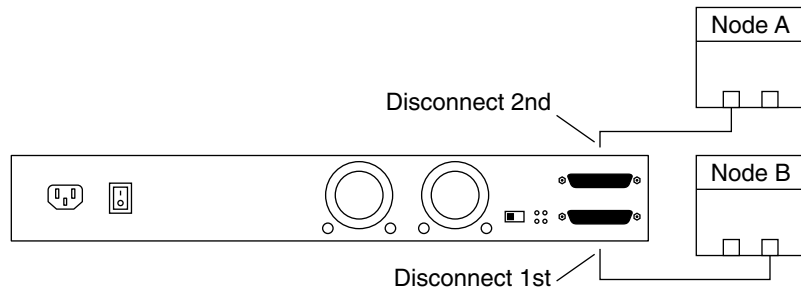
FIGURE 3–22 SPARC: Replacing a Sun StorEdge Netra D130/S1 Storage Array

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

Disconnect the SCSI cables from the storage arrays. Disconnect the SCSI cable in the order that is shown in Figure 3–23.
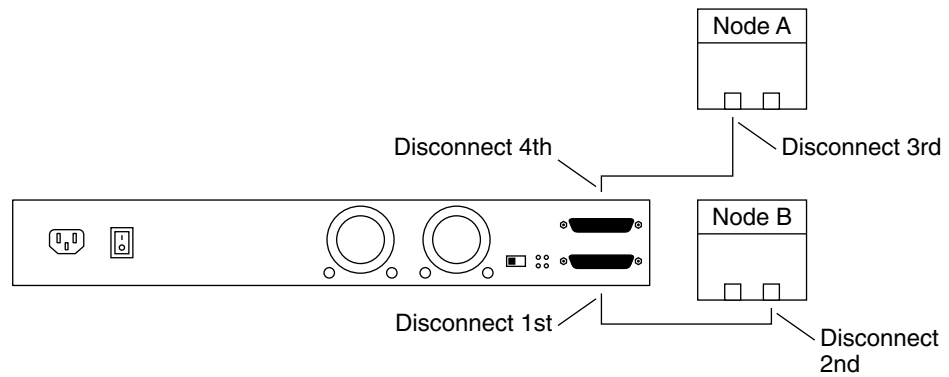
FIGURE 3–23 SPARC: Removing a Sun StorEdge Netra D130/S1 Storage Array

SPARC: Sun StorEdge D2 Cabling Diagrams

SPARC: Installing a Sun StorEdge D2 Storage Array

You can install a single storage array by using a split-bus configuration. You can install two storage arrays by using a single-bus configuration.
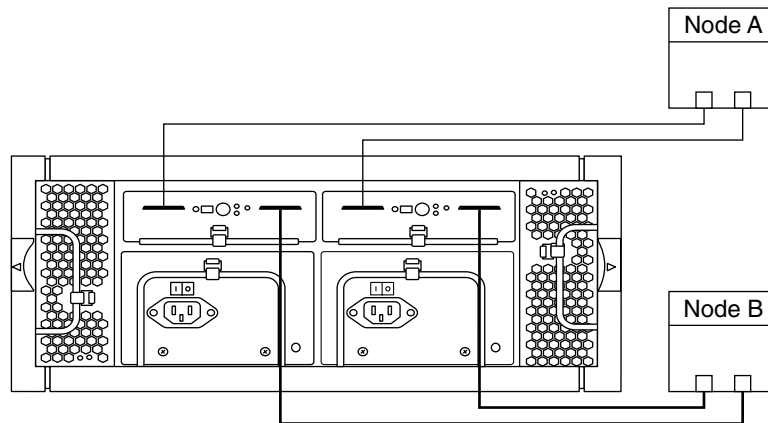
FIGURE 3–24 SPARC: Installing a Sun StorEdge D2 Split-Bus Configuration: An Example

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

As Figure 3–24 shows, the split-bus configuration requires that the storage array contain two Environmental Service Modules (ESM). Each ESM is connected to two nodes. Each ESM controls half the disks in the storage array and the data is mirrored by using a volume manager.
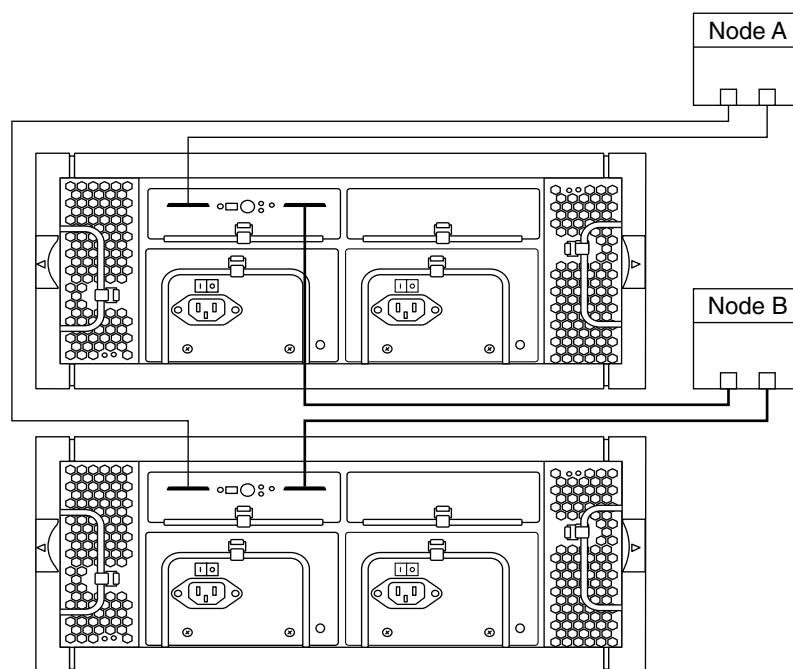
FIGURE 3–25 SPARC: Installing a Sun StorEdge D2 Single-Bus Configuration: An Example

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

SPARC: Adding a Sun StorEdge D2 Storage Array

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

- For split-bus configurations, see Figure 3–26 and Figure 3–28.
- For single-bus configurations, see Figure 3–27 and Figure 3–29.
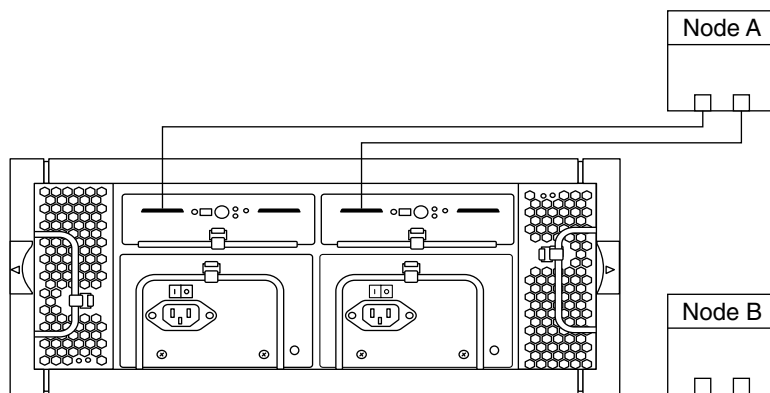
FIGURE 3–26 SPARC: Adding a Sun StorEdge D2 Split-Bus Configuration: Part 1
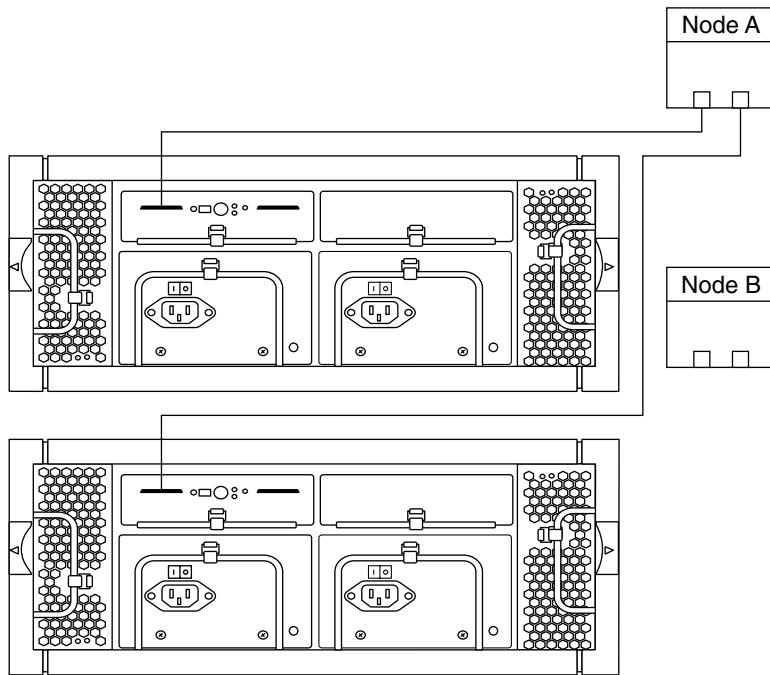
FIGURE 3–27 SPARC: Adding a Sun StorEdge D2 Single-Bus Configuration: Part 1
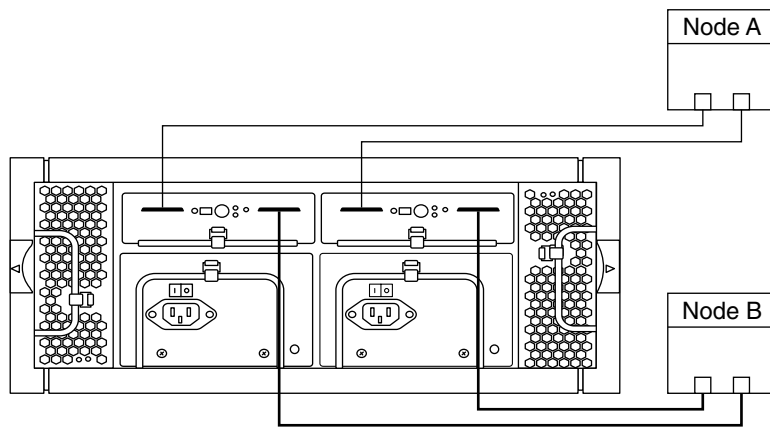
FIGURE 3–28 SPARC: Adding a Sun StorEdge D2 Split-Bus Configuration: Part 2
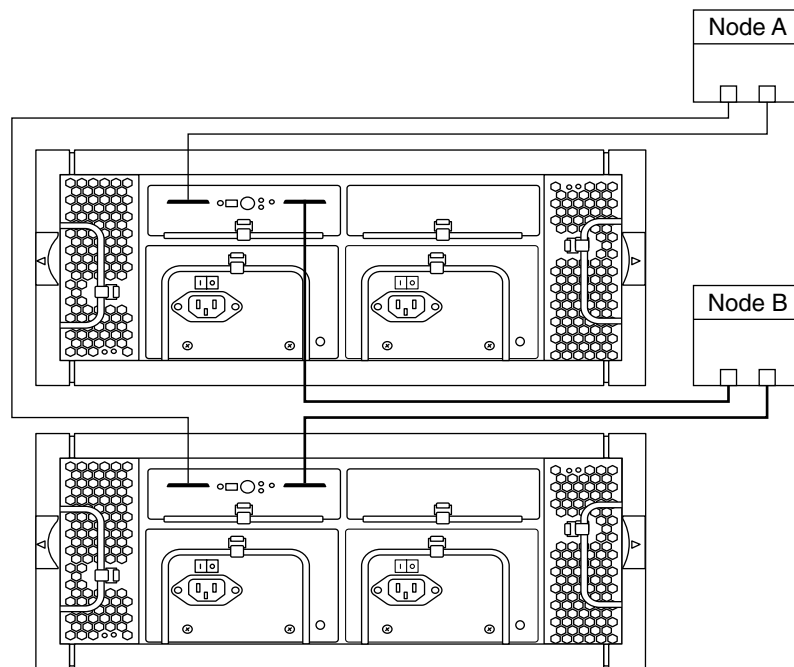
FIGURE 3–29 SPARC: Adding a Sun StorEdge D2 Single-Bus Configuration: Part 2

Index