Sun microsystems

Sun Cluster 3.0-3.1 With StorEdge A1000 Array, Netra st A1000 Array, or StorEdge A3500 System Manual

For Solaris OS (SPARC Platform Edition)

Adobe PostScript™

050503@11223

Contents

Preface

The *Sun Cluster 3.0-3.1 With StorEdge A1000 Array, Netra st A1000 Array, or StorEdge A3500 System Manual* provides procedures specific to SCSI RAID storage devices that use Sun™ StorEdge™ RAID Manager software in a Sun™ Cluster environment.

Note – 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:

- Chapter 1 contains procedures about installing, configuring, and maintaining SCSI RAID storage devices that use RAID Manager software.

- Appendix A contains cabling diagrams for the various 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 StorEdge A1000 Array, Netra st A1000 Array, or StorEdge A3500 System Manual

Revision Date	New Information
	No revisions.

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 Cluster software and data services
- Storage Array and Storage System patches
- Netra™ st A1000 array
- Sun™ StorEdge™ A1000 array
- Sun StorEdge A3500 system
- Sun StorEdge Traffic Manager software
- Sun StorEdge RAID Manager software

TABLE P–2 Storage Array and Storage System Patches

Title
EarlyNotifier 20029, A1000/A3x00/A3500FC Software/Firmware Configuration Matrix
Available on http://sunsolve.sun.com

TABLE P–3 Sun StorEdge Traffic Manager and Sun StorEdge RAID Manager

Title	Part Number
Sun StorEdge RAID Manager Release Notes	805-7758
Available at http://docs.sun.com/db/doc/805-7758-13	
Sun StorEdge RAID Manager User's Guide	806-0478
Available at http://docs.sun.com/db/doc/806-0478-10	
Sun StorEdge Traffic Manager Installation and Configuration Guide	816-1420
Available at http://docs.sun.com/db/doc/816-1420-11	

TABLE P–4 Sun StorEdge A1000 and Netra st A1000 Storage Array

Title	Part Number
OpenBoot 2.x Command Reference Manual	806-2906
Available at http://docs.sun.com/db/doc/806-2906–10	
Netra st A1000/D1000 Installation and Maintenance Manual	805-7147
Available at http://docs.sun.com/db/doc/805-7147	
Sun StorEdge A1000 and D1000 Installation, Operations, and Service Manual	805-2624
Available at http://docs.sun.com/db/doc/805-2624-12	
Sun StorEdge D1000 Storage Guide	805-4013
Available at http://docs.sun.com/db/doc/805-4013-10	

TABLE P–5 Sun StorEdge A3500 Storage System

Title	Part Number
OpenBoot 2.x Command Reference Manual	806-2906
Available at http://docs.sun.com/db/doc/806-2906–10	
Sun StorEdge A3500/A3500FC Hardware Configuration Guide	805-4981
Available at http://docs.sun.com/db/doc/805-4981-13	

TABLE P–6 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*
Data service administration	*Sun Cluster Data Services Planning and Administration Guide for Solaris OS*
	Individual data service guides
Data service development	*Sun Cluster Data Services Developer's Guide for Solaris OS*
System administration	*Sun Cluster System Administration Guide for Solaris OS*
Error messages	*Sun Cluster Error Messages Guide for Solaris OS*
Command and function reference	*Sun Cluster Reference Manual for Solaris OS*
Release Notes	*Sun Cluster 3.1 8/05 Release Notes for Solaris OS*
	Sun Cluster 3.x Release Notes Supplement

Using UNIX Commands

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

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

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

Getting Help

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

- Your name and email address (if available)
- Your company name, address, and phone number
- The model number and serial number of your systems
- The release number of the operating environment (for example, Solaris 8)
- The release number of Sun Cluster (for example, Sun Cluster 3.0)

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

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

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

Accessing Sun Documentation Online

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

Ordering Sun Documentation

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

Typographic Conventions

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

TABLE P–7 Typographic Conventions

Typeface or Symbol	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.`
AaBbCc123	What you type, contrasted with onscreen computer output	`machine_name%` **su** `Password:`
AaBbCc123	Command-line 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*. Perform a *patch analysis*. Do *not* save the file. [Note that some emphasized items appear bold online.]

Shell Prompts in Command Examples

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

TABLE P–8 Shell Prompts

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

Product Training

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

Installing and Maintaining a SCSI RAID Storage Device

This chapter describes the procedures about how to install, configure, and maintain SCSI RAID storage devices that use Sun™ StorEdge™ RAID Manager software in a Sun Cluster environment.

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

- Netra™ st A1000 array
- Sun™ StorEdge™ A1000 array
- Sun™ StorEdge™ A3500 system

This chapter contains the following sections:

Restrictions and Requirements

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

- RAID Manager software 6.22.1 or later version is required with Sun Cluster 3.0 software.

- *(StorEdge A3500 Only)* Do not use LUNs as quorum devices.

Installing Storage Arrays

This section contains the instructions for installing storage arrays both in new clusters and existing clusters.

TABLE 1–1 Task Map: Installing Storage Arrays

Task	Information
Install an array in a new cluster, before the OS and Sun Cluster software are installed.	"How to Install a Storage Array in a New Cluster" on page 14
Add an array to an operational cluster.	"How to Add a Storage Array to an Existing Cluster" on page 18

▼ How to Install a Storage Array in a New Cluster

This procedure assumes you are installing one or more storage arrays at initial installation of a cluster.

This procedure uses an updated method for setting the scsi-initiator-id. The method that was published in earlier documentation is still applicable. However, if your cluster configuration uses a Sun StorEdge PCI Dual Ultra3 SCSI host adapter to connect to any other shared storage, you need to update your nvramrc script and set the scsi-initiator-id by following 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.
- Your nodes are powered off or are at the Open Boot PROM.
- Your arrays are powered off.
- Your interconnect hardware is connected to the nodes in your cluster.
- No software is installed.

Steps   1. **Install the host adapters in the nodes that connect to the storage arrays.**

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

2. **Cable the storage arrays.**

For cabling diagrams, see Appendix A.

3. **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
```

4. **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
...
```

5. **Install the Solaris Operating System, then apply any required Solaris patches.**

   For the most current list of patches, see http://sunsolve.sun.com.

6. **Read the following two conditions carefully to determine whether you must reboot the nodes.**

   - **If you are using a version of RAID Manager later than 6.22, proceed to Step 7.**

   - **If you are using a version of the Solaris Operating System earlier than Solaris 8 Update 4, proceed to Step 7.**

   - **If you are using RAID Manager 6.22** *and* **the Solaris 8 Update 4 or later operating environment, reboot both nodes.**

     ```
     # reboot
     ```

7. **Install the RAID Manager software.**

   For the procedure about how to install the RAID Manager software, see the *Sun StorEdge RAID Manager User's Guide*.

   For the required version of the RAID Manager software that Sun Cluster software supports, see "Restrictions and Requirements" on page 13.

8. **Install patches for the controller modules and RAID Manager software.**

   For the most current list of patches, see `http://sunsolve.sun.com`.

9. **Check the NVSRAM file revision for the storage arrays. If necessary, install the most recent revision.**

   For the NVSRAM file revision number, boot level, and procedure about how to upgrade the NVSRAM file, see the *Sun StorEdge RAID Manager Release Notes*.

10. **Check the controller module firmware revision for the storage arrays. If necessary, install the most recent revision.**

    For the firmware revision number and boot level, see the *Sun StorEdge RAID Manager Release Notes*. For the procedure about how to upgrade the firmware, see the *Sun StorEdge RAID Manager User's Guide*.

11. **Set the `Rdac` parameters in the `/etc/osa/rmparams` file on both nodes.**

    ```
    Rdac_RetryCount=1
    Rdac_NoAltOffline=TRUE
    ```

12. **Ensure that the controller module is set to `active/active` mode.**

    For more information about controller modes, see the *Sun StorEdge RAID Manager Installation and Support Guide* and the *Sun StorEdge RAID Manager User's Guide*.

13. **Set up the storage arrays with logical unit numbers (LUNs) and hot spares.**

    For the procedure about how to set up the storage array with LUNs and hot spares, see the *Sun StorEdge RAID Manager User's Guide*.

    Note – Use the `format` command to verify Solaris logical device names.

    ---

14. **Copy the `/etc/raid/rdac_address` file from the node on which you created the LUNs to the other node. If you copy this file to the other node, you ensure consistency across both nodes.**

15. **Ensure that the new logical name for the LUN that you created in Step 13 appears in the `/dev/rdsk` directory on both nodes.**

    ```
    # /etc/raid/bin/hot_add
    ```

See Also   To continue with Sun Cluster software and data services installation tasks, see your Sun Cluster software installation documentation and the Sun Cluster data services developer's documentation. For a list of Sun Cluster documentation, see "Related Documentation" on page 6.

▼ How to Add a Storage Array to an Existing Cluster

Use this procedure to add a storage device to an existing cluster. If you need to install a storage device in a new cluster, use the procedure in "How to Install a Storage Array in a New Cluster" on page 14.

You might want to perform this procedure in the following scenarios.

- You need to increase available storage.
- You need to upgrade to a higher-quality or 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, use this procedure.

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.

- This procedure uses an updated method for setting the scsi-initiator-id. For this storage array, the method that was published in earlier documentation is still applicable. However, if your cluster configuration uses a Sun StorEdge PCI Dual Ultra3 SCSI host adapter to connect to any other shared storage, you need to update your nvramrc script and set the scsi-initiator-id by using this procedure.

Steps 1. **Install the RAID Manager software.**

- For the required version of the RAID Manager software that Sun Cluster software supports, see "Restrictions and Requirements" on page 13.

- For the procedure about how to install RAID Manager software, see the *Sun StorEdge RAID Manager Installation and Support Guide*.

- For the most current list of software, firmware, and patches that your storage array or storage system requires, refer to the appropriate EarlyNotifier that is outlined in "Related Documentation" on page 6. This document is available online to Sun service providers and to customers with SunSolve service contracts at the SunSolve site: http://sunsolve.sun.com.

2. **Install the storage array or storage system patches.**

For the location of patches and installation instructions, see your Sun Cluster release notes documentation. For a list of Sun Cluster documentation, see "Related Documentation" on page 6.

3. **Set the Rdac parameters in the /etc/osa/rmparams file.**

```
Rdac_RetryCount=1
Rdac_NoAltOffline=TRUE
```

4. **Power on the storage array or storage system.**

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

5. **Are you installing new host adapters in your nodes?**

- **If no, skip to Step 7.**

- **If yes, shut down and power off Node A.**

For the procedure about how to shut down and power off a node, see your Sun Cluster system administration documentation. For a list of Sun Cluster documentation, see "Related Documentation" on page 6.

6. **Install the host adapters in Node A.**

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

7. **Cable the storage array or storage system to Node A.**

For cabling diagrams, see Appendix A.

8. **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
```

9. **Are you installing new host adapters in Node B to connect Node B to the storage
array or storage system?**

- **If no, skip to Step 11.**

- **If yes, shut down and power off the node.**

    For the procedure about how to shut down and power off a node, see your Sun Cluster system administration documentation. For a list of Sun Cluster documentation, see "Related Documentation" on page 6.

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

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

11. **Cable the storage array or storage system to `Node B`.**

    For cabling diagrams, see "Adding a Sun StorEdge A3500 Storage System" on page 49.

12. **Did you power off `Node B` to install a host adapter?**

    - **If no, skip to Step 14.**

    - **If yes, power on `Node B` and the storage array or storage system.** *Do not enable the node to boot.* **If necessary, halt the system to continue with OpenBoot PROM (OBP) Monitor tasks.**

13. **Verify that `Node B` recognizes the new host adapters and disk drives.**

    If the node does not recognize the new hardware, check all hardware connections and repeat the installation steps you performed in Step 10.

    ```
    {0} ok show-disks
    ...
    b) /sbus@6,0/QLGC,isp@2,10000/sd...
    d) /sbus@2,0/QLGC,isp@2,10000/sd...{0} ok
    ```

14. **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
       ...
       ```

15. **Did you power off `Node B` to install a host adapter?**

    - **If no, skip to Step 19.**

- **If yes, perform a reconfiguration boot to create the new Solaris device files and links.**

16. **Check the controller module NVSRAM file revision. If necessary, install the most recent revision.**

    For the NVSRAM file revision number and boot level, see the *Sun StorEdge RAID Manager Release Notes*. For the procedure about how to upgrade the NVSRAM file, see the *Sun StorEdge RAID Manager User's Guide*.

17. **Verify the controller module firmware revision. If necessary, install the most recent firmware revision.**

    For the revision number and boot level of the controller module firmware, see the *Sun StorEdge RAID Manager Release Notes*. For the procedure about how to upgrade the controller firmware, see "How to Upgrade Controller Module Firmware" on page 40.

18. **One node at a time, boot each node into cluster mode.**

    ```
    # reboot
    ```

19. **On one node, verify that the device IDs have been assigned to the LUNs for all nodes. These nodes are attached to the storage array or storage system.**

    ```
    # scdidadm -L
    ```

20. *(StorEdge A3500 Only)* **Verify that the controller module is set to `active/active` mode.**

    For more information about controller modes, see the *Sun StorEdge RAID Manager Installation and Support Guide* and the *Sun StorEdge RAID Manager User's Guide*.

See Also   To create a LUN from disk drives that are unassigned, see "How to Create a LUN" on page 24.

To upgrade controller module firmware, see "How to Upgrade Controller Module Firmware" on page 40.

Configuring Storage Arrays

This section contains the procedures about how to configure a storage array or storage system *after* you install Sun Cluster software. Table 1–2 lists these procedures.

To configure a storage array or storage system *before* you install Sun Cluster software, use the same procedures you use in a noncluster environment. For the procedures about how to configure a storage system before you install Sun Cluster software, see the *Sun StorEdge RAID Manager User's Guide*.

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

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

To fix device IDs that report this error, run the `scdidadm -R` command for each affected device.

TABLE 1–2 Task Map: Configuring Disk Drives

Task	Information
Create a logical unit number (LUN).	"How to Create a LUN" on page 24
Remove a LUN.	"How to Delete a LUN" on page 26
Reset the LUN configuration.	"How to Reset the LUN Configuration" on page 28

The following is a list of administrative tasks that require no cluster-specific procedures. See the *Sun StorEdge RAID Manager User's Guide* for these procedures.

- Rebalancing running LUNs
- Creating a hot spare
- Deleting a hot spare
- Increasing the size of a drive group

▼ How to Create a LUN

Use this procedure to create a logical unit number (LUN) from unassigned disk drives or remaining capacity. For information about LUN administration, see the *Sun StorEdge RAID Manager Release Notes*.

This product supports the use of hardware RAID and host-based software RAID. For host-based software RAID, this product supports RAID levels 0+1 and 1+0.

Note – You must use hardware RAID for Oracle Parallel Server (OPS) data stored on the storage array. Do not place OPS data under volume management control. You must place all non-OPS data that is stored on the storage arrays under volume management control. Use either hardware RAID, host-based software RAID, or both types of RAID to manage your non-OPS data.

Hardware RAID uses the storage array's or storage system's hardware redundancy to ensure that independent hardware failures do not impact data availability. If you mirror across separate storage arrays, host-based software RAID ensures that independent hardware failures do not impact data availability when an entire storage array is offline. Although you can use hardware RAID and host-based software RAID concurrently, you need only one RAID solution to maintain a high degree of data availability.

Note – When you use host-based software RAID with hardware RAID, the hardware RAID levels you use affect hardware maintenance. If you use hardware RAID level 1, 3, or 5, you can perform most maintenance procedures without volume management disruptions. If you use hardware RAID level 0, some maintenance procedures require additional volume management administration because the availability of the LUNs is impacted.

Steps   1. **With all nodes booted and attached to the storage array or storage system, create the LUN on one node.**

     After the LUN formatting completes, a logical name for the new LUN appears in `/dev/rdsk` on all nodes. These nodes are attached to the storage array or storage system.

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

     ```
     ...
     corrupt label - wrong magic number
     ```

     For the procedure about how to create a LUN, refer to your storage device's documentation. Use the `format`(1M) command to verify Solaris logical device names.

   2. **Copy the `/etc/raid/rdac_address` file from the node on which you created the LUN to the other node. If you copy this file to the other node, you ensure consistency across both nodes.**

   3. **Ensure that the new logical name for the LUN that you created appears in the `/dev/rdsk` directory on both nodes.**

       ```
       # /etc/raid/bin/hot_add
       ```

4. **On one node, update the global device namespace.**

   ```
   # scgdevs
   ```

5. **Ensure that the device ID numbers for the LUNs are the same on both nodes. In the sample output that follows, the device ID numbers are different.**

   ```
   # scdidadm -L
   ...
   33        e07a:/dev/rdsk/c1t4d2           /dev/did/rdsk/d33
   33        e07c:/dev/rdsk/c0t4d2           /dev/did/rdsk/d33
   ```

6. **Are the device ID numbers that you received from running the `scdidadm` command in Step 5 the same for both nodes?**

   - **If the device ID numbers are the same, proceed to Step 7.**

   - **If the device ID numbers are different, perform the procedure in "How to Correct Mismatched Device ID Numbers" on page 31 before you continue with Step 7 of this procedure.**

7. *(A1000 Only)* **If you want a volume manager to manage the new LUN, incorporate the new LUN into a diskset or disk group.**

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

   Note – The StorEdge A3500 system does not support using LUNs as quorum devices.

▼ How to Delete a LUN

Use this procedure to delete one or more LUNs. You might need to delete a LUN to free up or reallocate resources, or to use the disks for other purposes. See the *Sun StorEdge RAID Manager Release Notes* for the latest information about LUN administration.

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

- All data on the LUN that you delete will be removed.
- You are not deleting LUN 0.

Steps 1. **From one node that is connected to the storage array or storage system, determine the paths to the LUN that you are deleting.**

   ```
   # format
   ```

For example:

```
phys-schost-1# format
Searching for disks...done
AVAILABLE DISK SELECTIONS:
0. c0t5d0 <SYMBIOS-StorEdgeA3500FCr-0301 cyl3 alt2 hd64 sec64>
/pseudo/rdnexus@0/rdriver@5,0
1. c0t5d1 <SYMBIOS-StorEdgeA3500FCr-0301 cyl2025 alt2 hd64 sec64>
/pseudo/rdnexus@0/rdriver@5,1
```

2. *(A1000 Only)* **Is the LUN a quorum device? This LUN is the LUN that you are removing.**

Note – Your storage array or storage system might not support LUNs as quorum devices. To determine if this restriction applies to your storage array or storage system, see "Restrictions and Requirements" on page 13.

```
# scstat -q
```

- **If no, proceed to Step 3.**

- **If yes, relocate that quorum device to another suitable storage array.**

  For procedures about how to add and remove quorum devices, see your Sun Cluster system administration documentation. For a list of Sun Cluster documentation, see "Related Documentation" on page 6.

3. **Does a volume manager manage the LUN that you are deleting?**

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

- **If yes, remove the LUN from any diskset or disk group. For more information, see your Solstice DiskSuite/Solaris Volume Manager or VERITAS Volume Manager documentation.**

  LUNs that were managed by VERITAS Volume Manager must be removed from VERITAS Volume Manager control before you can delete the LUNs. To remove the LUNs, after you delete the LUN from any disk group, use the following commands.

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

4. **Delete the LUN.**

For the procedure about how to delete a LUN, refer to your storage device's documentation.

5. **Remove the paths to the LUNs you are deleting.**

```
# rm /dev/rdsk/cNtXdY*
# rm /dev/dsk/cNtXdY*
```

6. **Complete the removal of the paths by issuing the following RAID Manager commands.**

   ```
   # rm /dev/osa/dev/dsk/cNtXdY*
   # rm /dev/osa/dev/rdsk/cNtXdY*
   ```

7. *(StorEdge A3500 Only)* **Determine the *alternate* paths to the LUNs you are deleting.**

   The RAID Manager software creates two paths to the LUN in the /dev/osa/dev/rdsk directory. Substitute the cNtXdY number from the other controller module in the storage array to determine the alternate path.

   For example:

   ```
   # lad
   c0t5d0 1T93600714 LUNS: 0 1
   c1t4d0 1T93500595 LUNS: 2
   ```

   Therefore, the alternate paths are as follows:

   ```
   /dev/osa/dev/dsk/c1t4d1*
   /dev/osa/dev/rdsk/c1t4d1*
   ```

8. *(StorEdge A3500 Only)* **Remove the *alternate* paths to the LUNs you are deleting.**

   ```
   # rm /dev/osa/dev/dsk/cNtXdY*
   # rm /dev/osa/dev/rdsk/cNtXdY*
   ```

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

   ```
   # scdidadm -C
   ```

10. **Move all resource groups and device groups off the node.**

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

11. **Shut down the node.**

    For the procedure about how to shut down and power off a node, see your Sun Cluster system administration documentation. For a list of Sun Cluster documentation, see "Related Documentation" on page 6.

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

13. **Repeat Step 4 through Step 12 on the other node that is attached to the storage array or storage system.**

▼ How to Reset the LUN Configuration

Use this procedure to completely remove and reset the LUN configuration.

⚠ Caution – If you reset a LUN configuration, a new device ID number is assigned to LUN 0. This change occurs because the software assigns a new world wide name (WWN) to the new LUN.

Steps 1. **From one node that is connected to the storage array or storage system, determine the paths to the LUNs you are resetting.**

```
# format
```
For example:

```
phys-schost-1# format
Searching for disks...done
AVAILABLE DISK SELECTIONS:
0. c0t5d0 <SYMBIOS-StorEdgeA3500FCr-0301 cyl3 alt2 hd64 sec64>
/pseudo/rdnexus@0/rdriver@5,0
1. c0t5d1 <SYMBIOS-StorEdgeA3500FCr-0301 cyl2025 alt2 hd64 sec64>
/pseudo/rdnexus@0/rdriver@5,1
```

2. *(A1000 Only)* **Is the LUN a quorum device? This LUN is the LUN that you are resetting.**

```
# scstat -q
```

- If no, proceed to Step 3.
- If yes, relocate that quorum device to another suitable storage array.

  For procedures about how to add and remove quorum devices, see your Sun Cluster system administration documentation. For a list of Sun Cluster documentation, see "Related Documentation" on page 6.

3. **Does a volume manager manage the LUNs on the controller module you are resetting?**

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

- **If yes, remove the LUN from any diskset or disk group. For more information, see your Solstice DiskSuite/Solaris Volume Manager or VERITAS Volume Manager documentation.**

  You must completely remove LUNs that were managed by VERITAS Volume Manager from VERITAS Volume Manager control before you can delete the LUNs.

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

4. **On one node, reset the LUN configuration.**

   For the procedure about how to reset the LUN configuration, see the *Sun StorEdge RAID Manager User's Guide*.

5. *(StorEdge A3500 Only)* **Set the controller module back to `active/active` mode.**

For more information about controller modes, see the *Sun StorEdge RAID Manager Installation and Support Guide* and the *Sun StorEdge RAID Manager User's Guide*.

6. **Use the `format` command to label the new LUN 0.**

7. **Remove the paths to the old LUNs you reset.**

   ```
   # rm /dev/rdsk/cNtXdY*
   # rm /dev/dsk/cNtXdY*

   # rm /dev/osa/dev/dsk/cNtXdY*
   # rm /dev/osa/dev/rdsk/cNtXdY*
   ```

8. *(StorEdge A3500 Only)* **Determine the** *alternate* **paths to the old LUNs you reset. Use the `lad` command.**

   The RAID Manager software creates two paths to the LUN in the `/dev/osa/dev/rdsk` directory. Substitute the `cNtXdY` number from the other controller module in the storage array to determine the alternate path.

   Example:

   ```
   # lad
   c0t5d0 1T93600714 LUNS: 0 1
   c1t4d0 1T93500595 LUNS: 2
   ```

   Therefore, the alternate paths are as follows:

   ```
   /dev/osa/dev/dsk/c1t4d1*
   /dev/osa/dev/rdsk/c1t4d1*
   ```

9. *(StorEdge A3500 Only)* **Remove the** *alternate* **paths to the old LUNs you reset.**

   ```
   # rm /dev/osa/dev/dsk/cNtXdY*
   # rm /dev/osa/dev/rdsk/cNtXdY*
   ```

10. **On both nodes, update device namespaces.**

    ```
    # devfsadm -C
    ```

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

    ```
    # scdidadm -C
    ```

12. **Move all resource groups and device groups off the node.**

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

13. **Shut down the node.**

    For the procedure about how to shut down and power off a node, see your Sun Cluster system administration documentation. For a list of Sun Cluster documentation, see "Related Documentation" on page 6.

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

    If an error message like the following appears, ignore it. Continue with the next step.

```
device id for '/dev/rdsk/c0t5d0' does not match physical disk's
id.
```

15. **After the node reboots and joins the cluster, repeat Step 7 through Step 14 on the other node. This node is attached to the storage array or storage system.**

The device ID number for the original LUN 0 is removed. A new device ID is assigned to LUN 0.

▼ How to Correct Mismatched Device ID Numbers

Use this section to correct mismatched device ID numbers that might appear during the creation of LUNs. You correct the mismatch by deleting Solaris and Sun Cluster paths to the LUNs that have device ID numbers that are different. After rebooting, the paths are corrected.

Note – Use this procedure only if you are directed to do so from "How to Create a LUN" on page 24.

Steps 1. **From one node that is connected to the storage array or storage system, determine the paths to the LUNs. These LUNs have different device ID numbers.**

   # **format**

2. **Remove the paths to the LUNs that have different device ID numbers.**

   ```
   # rm /dev/rdsk/cNtXdY*
   # rm /dev/dsk/cNtXdY*
   # rm /dev/osa/dev/dsk/cNtXdY*
   # rm /dev/osa/dev/rdsk/cNtXdY*
   ```

3. *(StorEdge A3500 Only)* **Use the lad command to determine the *alternate* paths to the LUNs that have different device ID numbers.**

   The RAID Manager software creates two paths to the LUN in the /dev/osa/dev/rdsk directory. Substitute the cNtXdY number from the other controller module in the storage array to determine the alternate path.

   For example:

   ```
   # lad
   c0t5d0 1T93600714 LUNS: 0 1
   c1t4d0 1T93500595 LUNS: 2
   ```

   Therefore, the alternate paths are as follows:

   ```
   /dev/osa/dev/dsk/c1t4d1*
   /dev/osa/dev/rdsk/c1t4d1*
   ```

4. *(StorEdge A3500 Only)* **Remove the *alternate* paths to the LUNs that have different device ID numbers.**

   ```
   # rm /dev/osa/dev/dsk/cNtXdY*
   # rm /dev/osa/dev/rdsk/cNtXdY*
   ```

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

   ```
   # scdidadm -C
   ```

6. **Move all resource groups and device groups off the node.**

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

7. **Shut down the node.**

   For the procedure about how to shut down and power off a node, see your Sun Cluster system administration documentation. For a list of Sun Cluster documentation, see "Related Documentation" on page 6.

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

9. **Repeat Step 1 through Step 8 on the other node. This node is attached to the storage array or storage system.**

10. **Return to "How to Create a LUN" on page 24.**

Maintaining Storage Arrays

The maintenance procedures in "FRUs That *Do Not* Require Sun Cluster Maintenance Procedures" on page 33 are performed the same as in a noncluster environment. Table 1–3 lists the procedures that require cluster-specific steps.

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

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

To fix device IDs that report this error, run the scdidadm -R command for each affected device.

TABLE 1–3 Task Map: Maintaining a Storage Array or Storage System

Task	Information
Remove a storage array or storage system	"How to Remove a Storage Array" on page 35
Replace a storage array or storage system Replacing a storage array or storage system, requires first removing the storage array or storage system, then adding a new storage array or storage system to the configuration.	"How to Add a Storage Array to an Existing Cluster" on page 18 "How to Remove a Storage Array" on page 35
Replace a failed controller module or restore an offline controller module	"How to Replace a Failed Controller or Restore an Offline Controller" on page 38
Upgrade controller module firmware and NVSRAM file	"How to Upgrade Controller Module Firmware" on page 40
Add a disk drive	"How to Add a Disk Drive" on page 41
Replace a disk drive	"How to Replace a Disk Drive" on page 42
Remove a disk drive	"How to Remove a Disk Drive" on page 43
Upgrade disk drive firmware	"How to Upgrade Disk Drive Firmware" on page 44
Replace a host adapter	"How to Replace a Host Adapter" on page 45

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:

- Netra st A1000 array
- Sun StorEdge A1000 array
- Sun StorEdge A3500 system

Sun StorEdge A1000 Array and Netra st A1000 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 *Netra st A1000/D1000 Installation and Maintenance Manual* for these procedures.

- Replacing a power supply
- Replacing a battery unit
- Replacing a storage array cooling canister
- Replacing a power cable on the storage array

Replacing a storage array-to-host SCSI cable requires no cluster-specific procedures. See the *Sun StorEdge RAID Manager User's Guide* and the *Sun StorEdge RAID Manager Release Notes* for these procedures.

Sun StorEdge A3500 System FRUs

With the exception of one item, the following is a list of administrative tasks that require no cluster-specific procedures. Shut down the cluster, and then see the *Sun StorEdge A3500/A3500FC Controller Module Guide*, the *Sun StorEdge A1000 and D1000 Installation, Operations, and Service Manual*, and the *Sun StorEdge Expansion Cabinet Installation and Service Manual* for the following procedures. See the Sun Cluster system administration documentation for procedures about how to shut down a cluster. For a list of Sun Cluster documentation, see "Related Documentation" on page 6.

- Replacing a power cord that connects to the cabinet power distribution unit (see the *Sun StorEdge Expansion Cabinet Installation and Service Manual*).

- Replacing a power cord to a storage array (see the *Sun StorEdge A1000 and D1000 Installation, Operations, and Service Manual*).

- Replacing the power-supply housing.

- Replacing the battery unit.

- Replacing a DC-power or battery harness.

- Replacing a controller module assembly.

- Replacing a controller module power cord.

- Replacing a controller module card cage.

The following is a list of administrative tasks that require no cluster-specific procedures. See the *Sun StorEdge A3500/A3500FC Controller Module Guide*, the *Sun StorEdge RAID Manager User's Guide*, the *Sun StorEdge RAID Manager Release Notes*, the *Sun StorEdge FC-100 Hub Installation and Service Manual*, and the documentation that shipped with your FC hub or FC switch for the following procedures.

- Replacing a storage array-to-host SCSI cable.

- Replacing a controller module fan canister.

- Replacing the power supply fan canister.
- Replacing a SCSI cable from the controller module to the storage array.
- Replacing a storage array-to-host or storage array–to-hub fiber-optic cable.
- Replacing an FC hub (see the *Sun StorEdge FC-100 Hub Installation and Service Manual*).
- Replacing an FC hub gigabit interface converter (GB(C) or Small Form-Factor Pluggable (SFP) that connects cables to the host or hub.
- Replacing a GBIC or an SFP on a node.
- Replacing an FC switch.

▼ How to Remove a Storage Array

Caution – This procedure removes all data that is on the storage array or storage system you are removing.

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

- Your cluster is operational.
- You no longer need the data that is stored on the storage array or storage system you are removing.

Steps 1. **Migrate any Oracle Real Application Clusters tables, data services, or volumes off the storage array or storage system.**

2. **Is one of the LUNs in the storage array a quorum device? This storage array is the storage array that you are removing.**

   Note – Your storage array or storage system might not support LUNs as quorum devices. To determine if this restriction applies to your storage array or storage system, see "Restrictions and Requirements" on page 13.

   ```
   # scstat -q
   ```

   - **If no, proceed to Step 3.**

   - **If yes, relocate that quorum device to another suitable storage array.**
     For procedures about how to add and remove quorum devices, see your Sun Cluster system administration documentation. For a list of Sun Cluster documentation, see "Related Documentation" on page 6.

3. **Halt all activity to the controller module.**

   For instructions, see your storage device documentation and your operating system documentation.

4. **Does a volume manager manage any of the LUNs on the controller module you are removing?**

   - **If no, proceed to Step 10.**

   - **If yes, remove the LUN from any diskset or disk group. For more information, see your Solstice DiskSuite/Solaris Volume Manager or VERITAS Volume Manager documentation.**

     You must completely remove LUNs that were managed by VERITAS Volume Manager from VERITAS Volume Manager control before you can delete the LUNs.

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

5. **Delete the LUN.**

   For the procedure about how to delete a LUN, see your storage device's documentation.

6. **Remove the paths to the LUNs you deleted in Step 5.**

   ```
   # rm /dev/rdsk/cNtXdY*
   # rm /dev/dsk/cNtXdY*

   # rm /dev/osa/dev/dsk/cNtXdY*
   # rm /dev/osa/dev/rdsk/cNtXdY*
   ```

7. **On all nodes, remove references to the storage array.**

   ```
   # scdidadm -C
   ```

8. *(StorEdge A3500 Only)* **Use the `lad` command to determine the** *alternate* **paths to the LUN you are deleting.**

   The RAID Manager software creates two paths to the LUN in the /dev/osa/dev/rdsk directory. Substitute the cNtXdY number from the other controller module in the storage array to determine the alternate path.

   For example:

   ```
   # lad
   c0t5d0 1T93600714 LUNS: 0 1
   c1t4d0 1T93500595 LUNS: 2
   ```

   Therefore, the alternate paths are as follows:

   ```
   /dev/osa/dev/dsk/c1t4d1*
   /dev/osa/dev/rdsk/c1t4d1*
   ```

9. *(StorEdge A3500 Only)* **Remove the** *alternate* **paths to the LUNs you deleted in Step 5.**

   ```
   # rm /dev/osa/dev/dsk/cNtXdY*
   # rm /dev/osa/dev/rdsk/cNtXdY*
   ```

10. **Disconnect all cables from the storage array and storage system. Remove the hardware from your cluster.**

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 no other parallel SCSI devices are connected to the nodes, you can delete the contents of the nvramrc script. At the OpenBoot PROM, set setenv use-nvramrc? to false.

    ---

12. **Remove any unused host adapter from nodes that were attached to the storage array or storage system.**

    a. **Shut down and power off Node A, from which you are removing a host adapter.**

       For the procedure about how to shut down and power off a node, see your Sun Cluster system administration documentation. For a list of Sun Cluster documentation, see "Related Documentation" on page 6.

    b. **Remove the host adapter from Node A.**

       For the procedure about how to remove a host adapter, see the documentation that shipped with your node hardware.

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

    d. **Repeat Step a through Step c for Node B that was attached to the storage array or storage system.**

13. **Return resource groups to their primary nodes.**

    ```
    # scswitch -Z
    ```

14. **Are you removing the *last* storage array or storage system from your cluster?**

    - **If no, you are finished with this procedure.**

    - **If yes, proceed to Step 15.**

15. **Remove RAID Manager patches, then remove RAID Manager software packages.**

Caution – If you improperly remove RAID Manager packages, the next reboot of the node fails. Before you remove RAID Manager software packages, see the *Sun StorEdge RAID Manager Release Notes*.

For the procedure about how to remove software packages, see the documentation that shipped with your storage array or storage system.

▼ How to Replace a Failed Controller or Restore an Offline Controller

This procedure assumes that your cluster is operational. For conceptual information about SCSI reservations and failure fencing, see your Sun Cluster concepts documentation. For a list of Sun Cluster documentation, see "Related Documentation" on page 6.

Steps 1. *(StorEdge A1000 Only)* **Is one of the LUNs in the storage array a quorum device?**

Note – Your storage array or storage system might not support LUNs as quorum devices. To determine if this restriction applies to your storage array or storage system, see "Restrictions and Requirements" on page 13.

```
# 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 your Sun Cluster system administration documentation. For a list of Sun Cluster documentation, see "Related Documentation" on page 6.

2. *(StorEdge A3500 Only)* **On both nodes, to prevent LUNs from automatic assignment to the controller that is being brought online, set the System_LunReDistribution parameter in the /etc/raid/rmparams file to false.**

Caution – You must set the System_LunReDistribution parameter in the /etc/raid/rmparams file to false so that no LUNs are assigned to the controller being brought online. After you verify in Step 6 that the controller has the correct SCSI reservation state, you can balance LUNs between both controllers.

For the procedure about how to modify the rmparams file, see the *Sun StorEdge RAID Manager Installation and Support Guide*.

3. **Restart the RAID Manager daemon.**

   ```
   # /etc/init.d/amdemon stop
   # /etc/init.d/amdemon start
   ```

4. **Do you have a failed controller?**

   - **If your controller module is offline, but does not have a failed controller, proceed to Step 5.**

   - **If you have a failed controller, replace the failed controller with a new controller.** *Do not bring the controller online.*

     For the procedure about how to replace controllers, see the *Sun StorEdge A3500/A3500FC Controller Module Guide* and the *Sun StorEdge RAID Manager Installation and Support Guide* for additional considerations.

5. **On one node, use the RAID Manager GUI's Recovery application to restore the controller online.**

   Note – You must use the RAID Manager GUI's Recovery application to bring the controller online. Do not use the Redundant Disk Array Controller Utility (rdacutil) because this utility ignores the value of the System_LunReDistribution parameter in the /etc/raid/rmparams file.

   For information about the Recovery application, see the *Sun StorEdge RAID Manager User's Guide*. If you have problems with bringing the controller online, see the *Sun StorEdge RAID Manager Installation and Support Guide*.

6. **On one node that is connected to the storage array or storage system, verify that the controller has the correct SCSI reservation state.**

   Run the scdidadm(1M) repair option (-R) on LUN 0 of the controller you want to bring online.

   ```
   # scdidadm -R /dev/dsk/cNtXdY
   ```

7. *(StorEdge A3500 Only)* **Set the controller to `active/active` mode. Assign LUNs to the controller.**

   For more information about controller modes, see the *Sun StorEdge RAID Manager Installation and Support Guide* and the *Sun StorEdge RAID Manager User's Guide*.

8. *(StorEdge A3500 Only)* **Reset the `System_LunReDistribution` parameter in the `/etc/raid/rmparams` file to `true`.**

   For the procedure about how to change the rmparams file, see the *Sun StorEdge RAID Manager Installation and Support Guide*.

9. *(StorEdge A3500 Only)* **Restart the RAID Manager daemon.**

```
# /etc/init.d/amdemon stop
# /etc/init.d/amdemon start
```

▼ How to Upgrade Controller Module Firmware

Use either the *online* or the *offline* method to upgrade your NVSRAM firmware. The method that you choose depends on your firmware.

Before You Begin   This procedure assumes that your cluster is operational

Steps   1. **Are you upgrading the NVSRAM firmware file?**

- If you are not upgrading the NVSRAM file, you can use the online method.

  Upgrade the firmware by using the online method, as described in the *Sun StorEdge RAID Manager User's Guide*. No special steps are required for a cluster environment.

- If you are upgrading the NVSRAM file, you must use an offline method. Use one of the following procedures.

  - If the data on your controller module is mirrored on another controller module, use the procedure in Step 2.

  - If the data on your controller module is *not* mirrored on another controller module, use the procedure in Step 3.

2. **Use this step if you are upgrading the NVSRAM and other firmware files on a controller module. This controller module must have mirrored data.**

   a. **Halt all activity to the controller module.**

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

   b. **Update the firmware files by using the offline method, as described in the** *Sun StorEdge RAID Manager User's Guide***.**

   c. **Restore all activity to the controller module.**

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

   This step completes the firmware upgrade.

3. **Use this step if you are upgrading the NVSRAM and other firmware files on a controller module. This controller module must** *not* **have mirrored data.**

   a. **Shut down the entire cluster.**

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

b. Boot one node that is attached to the controller module into noncluster mode.

For the procedure about how to boot a node in noncluster mode, see your Sun Cluster system administration documentation. For a list of Sun Cluster documentation, see "Related Documentation" on page 6.

c. Update the firmware files by using the offline method, as described in the *Sun StorEdge RAID Manager User's Guide***.**

d. Boot both nodes into cluster mode.

For more information about how to boot nodes, see your Sun Cluster system administration documentation. For a list of Sun Cluster documentation, see "Related Documentation" on page 6.

This step completes the firmware upgrade.

▼ 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 – If the disk drive that you are adding was previously owned by another controller module, reformat the disk drive to wipe clean the old DacStore information before adding the disk drive to this storage array. If any old DacStore information remains, it can cause aberrant behavior including the appearance of ghost disks or LUNs in the RAID Manager interfaces.

Before You Begin

This procedure relies on the following prerequisites and assumptions.

- Your cluster is operational.
- 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. For a list of Sun Cluster documentation, see "Related Documentation" on page 6.

Steps **1. Verify that the new disk drive is formatted.**

For information about how to move drives between storage arrays, see the *Sun StorEdge RAID Manager Release Notes*.

2. Install the new disk drive to the storage array.

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

3. **Enable the disk drive to spin up approximately 30 seconds.**

4. **Run Health Check to ensure that the new disk drive is not defective.**

   For instructions about how to run Recovery Guru and Health Check, see the *Sun StorEdge RAID Manager User's Guide*.

5. **Fail the new drive, then revive the drive to update DacStore on the drive.**

   For the procedure about how to fail and revive drives, see the *Sun StorEdge RAID Manager User's Guide*.

6. **Repeat Step 1 through Step 5 for each disk drive you are adding.**

See Also   To create LUNs for the new drives, see "How to Create a LUN" on page 24 for more information.

▼ How to Replace a Disk Drive

Removing a disk drive enables you to reduce or reallocate your existing storage pool. You might want to perform this procedure if a disk has failed or is behaving in an unreliable manner.

For conceptual information about quorum, quorum devices, global devices, and device IDs, see your Sun Cluster concepts documentation. For a list of Sun Cluster documentation, see "Related Documentation" on page 6.

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. For a list of Sun Cluster documentation, see "Related Documentation" on page 6.

Steps   1. **Does replacing the disk drive affect any LUN's availability?**

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

   - **If yes, remove the LUNs from volume management control. For more information, see your Solstice DiskSuite/Solaris Volume Manager or VERITAS Volume Manager documentation.**

2. **Replace the disk drive in the storage array.**

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

3. **Run Health Check to ensure that the new disk drive is not defective.**

    For the procedure about how to run Recovery Guru and Health Check, see the *Sun StorEdge RAID Manager User's Guide*.

4. **Does the failed drive belong to a drive group?**

    - **If the drive does *not* belong to a device group, proceed to Step 5.**

    - **If the drive is part of a device group, reconstruction is started automatically. If reconstruction does not start automatically for any reason, then select Reconstruct from the Manual Recovery application. Do not select Revive. When reconstruction is complete, skip to Step 6.**

5. **Fail the new drive, then revive the drive to update DacStore on the drive.**

    For the procedure about how to fail and revive drives, see the *Sun StorEdge RAID Manager User's Guide*.

6. **If you removed LUNs from volume management control in Step 1, return the LUNs to volume management control.**

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

▼ How to Remove a Disk Drive

Removing a disk drive enables you to reduce 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 your Sun Cluster concepts documentation.

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

- Your cluster is operational.
- 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 35.

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

    If you need to replace your storage array's chassis, see "FRUs That *Do Not* Require Sun Cluster Maintenance Procedures" on page 33.

- 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. For a list of Sun Cluster documentation, see "Related

Steps   **1. Is the logical unit number (LUN) that is associated with the disk drive a quorum device? This disk drive is the disk drive that you are removing.**

Note – Your storage array or storage system might not support LUNs as quorum devices. To determine if this restriction applies to your storage array or storage system, see "Restrictions and Requirements" on page 13.

```
# 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 your Sun Cluster system administration documentation.

2. Remove the LUN that is associated with the disk drive you are removing.

For the procedure about how to remove a LUN, see "How to Delete a LUN" on page 26.

3. Remove the disk drive 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.

Caution – After you remove the disk drive, install a dummy drive to maintain proper cooling.

How to Upgrade Disk Drive Firmware

Caution – You must be a Sun service provider to perform disk drive firmware updates. If you need to upgrade drive firmware, contact your Sun service provider.

▼ How to Replace a Host Adapter

Note – Several steps in this procedure require you to halt I/O activity. To halt I/O activity, take the controller module offline by using the RAID Manager GUI's manual recovery procedure in the *Sun StorEdge RAID Manager User's Guide*.

Before You Begin

This procedure relies on the following prerequisites and assumptions.

- Your cluster is operational.
- The node on which the host adapter resides is attached to a *SCSI-based* storage array or storage system.
- This procedure defines `Node A` as the node with the host adapter on SCSI bus A. This host adapter is the host adapter that you are replacing. `Node B` is the node that remains in service.

Steps

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

   `# scstat`

   Record this information because you will use it in Step 23 of this procedure to return resource groups and device groups to this node.

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

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

3. **Without powering off the node, shut down `Node A`.**

   For the procedure about how to shut down and power off a node, see your Sun Cluster system administration documentation. For a list of Sun Cluster documentation, see "Related Documentation" on page 6.

4. **From `Node B`, halt I/O activity to SCSI bus A.**

   For instructions, see the *Sun StorEdge RAID Manager User's Guide*.

5. **From the controller module end of the SCSI cable, disconnect the SCSI bus A cable. This cable connects the controller module to `Node A`. Afterward, replace this cable with a differential SCSI terminator.**

6. **Restart I/O activity on SCSI bus A.**

   For instructions, see the *Sun StorEdge RAID Manager User's Guide*.

7. **Does servicing the failed host adapter affect SCSI bus B?**

   - **If no, skip to Step 11.**

   - **If yes, proceed to Step 8.**

8. **From `Node B`, halt I/O activity to the controller module on SCSI bus B.**

For instructions, see the *Sun StorEdge RAID Manager User's Guide.*

9. **From the controller module end of the SCSI cable, disconnect the SCSI bus B cable. This cable connects the controller module to `Node A`. Afterward, replace this cable with a differential SCSI terminator.**

10. **Restart I/O activity on SCSI bus B.**

    For instructions, see the *Sun StorEdge RAID Manager User's Guide.*

11. **Power off `Node A`.**

12. **Replace `Node A`'s host adapter.**

    For the procedure about how to replace a host adapter, see the documentation that shipped with your node hardware.

13. **Power on `Node A`.** *Do not enable the node to boot*. **If necessary, halt the system.**

14. **From `Node B`, halt I/O activity to the controller module on SCSI bus A.**

    For instructions, see the *Sun StorEdge RAID Manager User's Guide.*

15. **Remove the differential SCSI terminator from SCSI bus A. Afterward, reinstall the SCSI cable to connect the controller module to `Node A`.**

16. **Restart I/O activity on SCSI bus A.**

    For instructions, see the *Sun StorEdge RAID Manager User's Guide.*

17. **Did you install a differential SCSI terminator to SCSI bus B in Step 9?**

    ■ **If no, skip to Step 20.**

    ■ **If yes, halt I/O activity on SCSI bus B, then continue with Step 18.**

18. **Remove the differential SCSI terminator from SCSI bus B. Afterward, reinstall the SCSI cable to connect the controller module to `Node A`.**

19. **Restart I/O activity on SCSI bus B.**

    For instructions, see the *Sun StorEdge RAID Manager User's Guide.*

20. **Bring the controller module online.**

    For instructions, see the *Sun StorEdge RAID Manager User's Guide.*

21. **Rebalance all logical unit numbers (LUNs).**

    For instructions, see the *Sun StorEdge RAID Manager User's Guide.*

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

23. **(Optional) Return resource groups and device groups to `Node A`.**

Cabling Diagrams

- Sun StorEdge A1000 and Netra st A1000 arrays
- Sun StorEdge A3500 system

Sun StorEdge A1000 and Netra st A1000 Cabling Diagrams

Installing a Sun StorEdge A1000 or Netra st A1000 Storage Array

The storage arrays must be configured in pairs for the Sun Cluster environment.

To cable the storage arrays, connect the differential SCSI cable between the node and the storage array. Ensure that the *entire* SCSI bus length in each SCSI chain 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 host adapter. Figure A–1 and Figure A–2 illustrate a storage array that is cabled in a Sun Cluster environment.

47
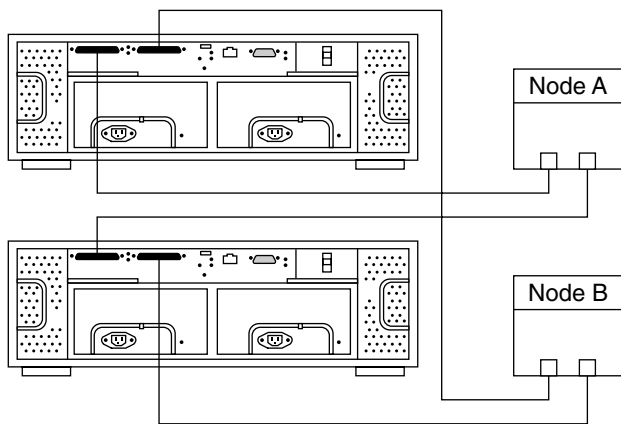
FIGURE A–1 Installing a Storage Array With One Pair

Figure A–2 illustrates how to cable the storage array in a daisy-chain configuration. If you place two storage arrays on a SCSI chain, ensure that each storage array's ID dial is set to a unique number.
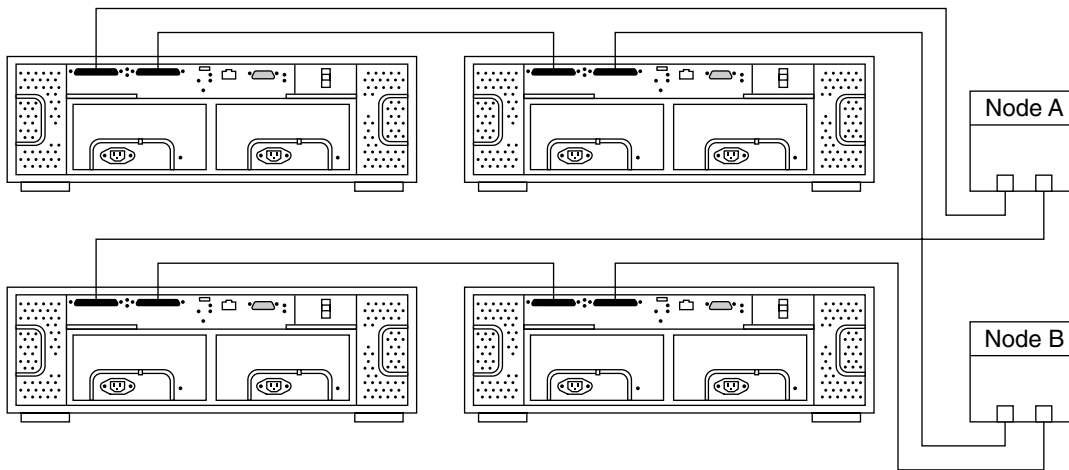
FIGURE A–2 Storage Array With Two Pairs and Daisy Chaining

Sun StorEdge A3500 Cabling Diagrams

Installing a Sun StorEdge A3500 Storage System
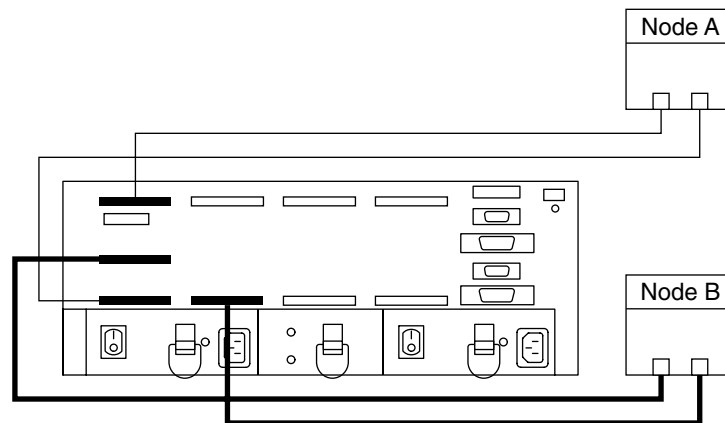
FIGURE A–3 Installing a StorEdge A3500 Storage System: An Example

For more sample configurations, see the *Sun StorEdge A3500/A3500FC Hardware Configuration Guide*.

For the procedure about how to install the cables, see the *Sun StorEdge A3500/A3500FC Controller Module Guide*.

Adding a Sun StorEdge A3500 Storage System

Connect the differential SCSI cable between the node and the controller module, as shown in Figure A–4. Ensure that the *entire* SCSI bus length 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 host adapter.
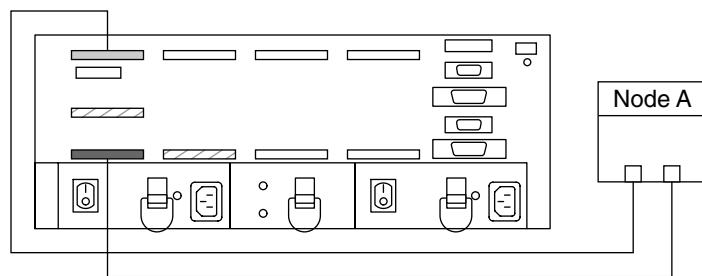
Differential SCSI terminators

Controller A

Controller B

FIGURE A–4 Cabling a StorEdge A3500 Storage System: An Example, Part 1
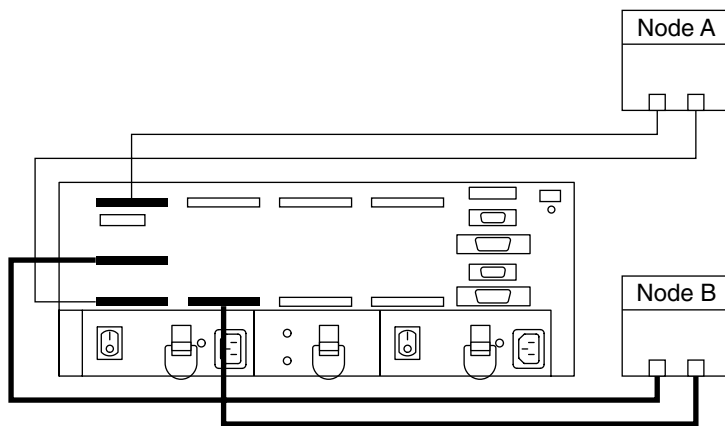
FIGURE A–5 Cabling a StorEdge A3500 Storage System: An Example, Part 2

Index