



Sun Cluster Quick Start Guide for Solaris OS

SPARC Platform Edition



Sun Microsystems, Inc.
4150 Network Circle
Santa Clara, CA 95054
U.S.A.

Part No: 819-0912-10
December 2006, Revision A

Copyright 2006 Sun Microsystems, Inc. 4150 Network Circle, Santa Clara, CA 95054 U.S.A. All rights reserved.

Sun Microsystems, Inc. has intellectual property rights relating to technology embodied in the product that is described in this document. In particular, and without limitation, these intellectual property rights may include one or more U.S. patents or pending patent applications in the U.S. and in other countries.

U.S. Government Rights – Commercial software. Government users are subject to the Sun Microsystems, Inc. standard license agreement and applicable provisions of the FAR and its supplements.

This distribution may include materials developed by third parties.

Parts of the product may be derived from Berkeley BSD systems, licensed from the University of California. UNIX is a registered trademark in the U.S. and other countries, exclusively licensed through X/Open Company, Ltd.

Sun, Sun Microsystems, the Sun logo, the Solaris logo, the Java Coffee Cup logo, docs.sun.com, OpenBoot, Sun StorEdge, Sun Fire, Sun Ultra, SunSolve, Java, and Solaris are trademarks or registered trademarks of Sun Microsystems, Inc. in the U.S. and other countries. All SPARC trademarks are used under license and are trademarks or registered trademarks of SPARC International, Inc. in the U.S. and other countries. Products bearing SPARC trademarks are based upon an architecture developed by Sun Microsystems, Inc. ORACLE is a registered trademark of Oracle Corporation. The Adobe PostScript logo is a trademark of Adobe Systems, Incorporated.

The OPEN LOOK and Sun™ Graphical User Interface was developed by Sun Microsystems, Inc. for its users and licensees. Sun acknowledges the pioneering efforts of Xerox in researching and developing the concept of visual or graphical user interfaces for the computer industry. Sun holds a non-exclusive license from Xerox to the Xerox Graphical User Interface, which license also covers Sun's licensees who implement OPEN LOOK GUIs and otherwise comply with Sun's written license agreements.

Products covered by and information contained in this publication are controlled by U.S. Export Control laws and may be subject to the export or import laws in other countries. Nuclear, missile, chemical or biological weapons or nuclear maritime end uses or end users, whether direct or indirect, are strictly prohibited. Export or reexport to countries subject to U.S. embargo or to entities identified on U.S. export exclusion lists, including, but not limited to, the denied persons and specially designated nationals lists is strictly prohibited.

DOCUMENTATION IS PROVIDED "AS IS" AND ALL EXPRESS OR IMPLIED CONDITIONS, REPRESENTATIONS AND WARRANTIES, INCLUDING ANY IMPLIED WARRANTY OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE OR NON-INFRINGEMENT, ARE DISCLAIMED, EXCEPT TO THE EXTENT THAT SUCH DISCLAIMERS ARE HELD TO BE LEGALLY INVALID.

Copyright 2006 Sun Microsystems, Inc. 4150 Network Circle, Santa Clara, CA 95054 U.S.A. Tous droits réservés.

Sun Microsystems, Inc. détient les droits de propriété intellectuelle relatifs à la technologie incorporée dans le produit qui est décrit dans ce document. En particulier, et ce sans limitation, ces droits de propriété intellectuelle peuvent inclure un ou plusieurs brevets américains ou des applications de brevet en attente aux Etats-Unis et dans d'autres pays.

Cette distribution peut comprendre des composants développés par des tierces personnes.

Certaines composants de ce produit peuvent être dérivées du logiciel Berkeley BSD, licenciés par l'Université de Californie. UNIX est une marque déposée aux Etats-Unis et dans d'autres pays; elle est licenciée exclusivement par X/Open Company, Ltd.

Sun, Sun Microsystems, le logo Sun, le logo Solaris, le logo Java Coffee Cup, docs.sun.com, OpenBoot, Sun StorEdge, Sun Fire, Sun Ultra, SunSolve, Java et Solaris sont des marques de fabrique ou des marques déposées de Sun Microsystems, Inc. aux Etats-Unis et dans d'autres pays. Toutes les marques SPARC sont utilisées sous licence et sont des marques de fabrique ou des marques déposées de SPARC International, Inc. aux Etats-Unis et dans d'autres pays. Les produits portant les marques SPARC sont basés sur une architecture développée par Sun Microsystems, Inc. ORACLE est une marque déposée registre de Oracle Corporation. Le logo Adobe PostScript est une marque déposée de Adobe Systems, Incorporated.

L'interface d'utilisation graphique OPEN LOOK et Sun a été développée par Sun Microsystems, Inc. pour ses utilisateurs et licenciés. Sun reconnaît les efforts de pionniers de Xerox pour la recherche et le développement du concept des interfaces d'utilisation visuelle ou graphique pour l'industrie de l'informatique. Sun détient une licence non exclusive de Xerox sur l'interface d'utilisation graphique Xerox, cette licence couvrant également les licenciés de Sun qui mettent en place l'interface d'utilisation graphique OPEN LOOK et qui, en outre, se conforment aux licences écrites de Sun.

Les produits qui font l'objet de cette publication et les informations qu'il contient sont régis par la législation américaine en matière de contrôle des exportations et peuvent être soumis au droit d'autres pays dans le domaine des exportations et importations. Les utilisations finales, ou utilisateurs finaux, pour des armes nucléaires, des missiles, des armes chimiques ou biologiques ou pour le nucléaire maritime, directement ou indirectement, sont strictement interdites. Les exportations ou réexportations vers des pays sous embargo des Etats-Unis, ou vers des entités figurant sur les listes d'exclusion d'exportation américaines, y compris, mais de manière non exclusive, la liste de personnes qui font objet d'un ordre de ne pas participer, d'une façon directe ou indirecte, aux exportations des produits ou des services qui sont régis par la législation américaine en matière de contrôle des exportations et la liste de ressortissants spécifiquement désignés, sont rigoureusement interdites.

LA DOCUMENTATION EST FOURNIE "EN L'ETAT" ET TOUTES AUTRES CONDITIONS, DECLARATIONS ET GARANTIES EXPRESSES OU TACITES SONT FORMELLEMENT EXCLUES, DANS LA MESURE AUTORISEE PAR LA LOI APPLICABLE, Y COMPRIS NOTAMMENT TOUTE GARANTIE IMPLICITE RELATIVE A LA QUALITE MARCHANDE, A L'APTITUDE A UNE UTILISATION PARTICULIERE OU A L'ABSENCE DE CONTREFAÇON.

Contents

Preface	5
1 Example of Installing and Configuring a Sun Cluster Configuration	9
Configuration Specifications and Assumptions	9
Hardware Configuration	10
Software Configuration	10
Public Network Addresses	11
Procedure Assumptions	12
Task Map: Creating a Sun Cluster Quick Start Configuration	12
Installing the Hardware	13
▼ How to Connect the Administrative Console	14
▼ How to Connect the Cluster Nodes	14
▼ How to Connect the Sun StorEdge 3510 FC RAID Array	15
▼ How to Configure the Storage Array	16
Installing the Software	17
▼ How to Install the Administrative Console	18
▼ How to Install the Solaris Operating System	19
▼ How to Set Up the User Environment	21
▼ How to Configure the Operating System	22
▼ How to Create State Database Replicas	25
▼ How to Mirror the Root (/) File System	25
▼ How to Install Sun Cluster Software	26
▼ How to Set Up the Oracle System Groups and User	27
Configuring the Cluster	28
▼ How to Establish the Cluster	28
Configuring Volume Management	29
▼ How to Create Disk Sets	29
▼ How to Add LUNs to Disk Sets	30
▼ How to Create and Activate an md . tab File	31

Creating File Systems	32
▼ How to Create File Systems	32
Installing and Configuring Application Software	34
▼ How to Configure Apache HTTP Server Software	34
▼ How to Install Oracle 10gR2 Software	35
▼ How to Create an Oracle Database	37
▼ How to Set Up Oracle Database Permissions	38
Configuring the Data Services	39
▼ How to Start Sun Cluster Manager	39
▼ How to Configure the Scalable Sun Cluster HA for Apache Data Service	39
▼ How to Configure the Sun Cluster HA for NFS Data Service	40
▼ How to Configure the Sun Cluster HA for Oracle Data Service	41
Index	43

Preface

The *Sun Cluster Quick Start Guide for Solaris OS* contains an example set of procedures for installing and configuring a specific Sun Cluster 3.2 configuration on a SPARC® based system. The configuration includes the Sun Cluster HA for Apache, Sun Cluster HA for NFS, and Sun Cluster HA for Oracle data services.

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

The instructions in this book assume knowledge of the Solaris™ Operating System (Solaris OS) and expertise with the volume-manager software that is used with Sun Cluster software.

Using UNIX Commands

This document contains information about commands that are used to install and configure 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 OS
- Other software documentation that you received with your system
- Solaris OS man pages

Typographic Conventions

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

TABLE P-1 Typographic Conventions

Typeface	Meaning	Example
AaBbCc123	The names of commands, files, and directories, and onscreen computer output	Edit your <code>.login</code> file. Use <code>ls -a</code> to list all files. <code>machine_name%</code> you have mail.
AaBbCc123	What you type, contrasted with onscreen computer output	<code>machine_name%</code> su Password:
<i>aabbcc123</i>	Placeholder: replace with a real name or value	The command to remove a file is <i>rm filename</i> .
<i>AaBbCc123</i>	Book titles, new terms, and terms to be emphasized	Read Chapter 6 in the <i>User's Guide</i> . A <i>cache</i> is a copy that is stored locally. Do <i>not</i> save the file. Note: Some emphasized items appear bold online.

Shell Prompts in Command Examples

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

TABLE P-2 Shell Prompts

Shell	Prompt
C shell	<code>machine_name%</code>
C shell for superuser	<code>machine_name#</code>
Bourne shell and Korn shell	<code>\$</code>
Bourne shell and Korn shell for superuser	<code>#</code>

Related Documentation

Information about related Sun Cluster topics is available in the documentation that is listed in the following table. All Sun Cluster documentation is available at <http://docs.sun.com>.

Topic	Documentation
Overview	<i>Sun Cluster Overview for Solaris OS</i>
Concepts	<i>Sun Cluster Concepts Guide for Solaris OS</i>
Hardware installation and administration	<i>Sun Cluster 3.1 - 3.2 Hardware Administration Manual for Solaris OS</i> Individual hardware administration guides
Software installation	<i>Sun Cluster Software Installation Guide for Solaris OS</i>
Data service installation and administration	<i>Sun Cluster Data Services Planning and Administration Guide for Solaris OS</i> Individual data service guides
Data service development	<i>Sun Cluster Data Services Developer's Guide for Solaris OS</i>
System administration	<i>Sun Cluster System Administration Guide for Solaris OS</i>
Error messages	<i>Sun Cluster Error Messages Guide for Solaris OS</i>
Command and function references	<i>Sun Cluster Reference Manual for Solaris OS</i>

For a complete list of Sun Cluster documentation, see the release notes for your release of Sun Cluster software at <http://docs.sun.com>.

Related Third-Party Web Site References

Sun is not responsible for the availability of third-party web sites mentioned in this document. Sun does not endorse and is not responsible or liable for any content, advertising, products, or other materials that are available on or through such sites or resources. Sun will not be responsible or liable for any actual or alleged damage or loss caused or alleged to be caused by or in connection with use of or reliance on any such content, goods, or services that are available on or through such sites or resources.

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/>)

Getting Help

If you have problems installing or using Sun Cluster software, contact your service provider and supply 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 Solaris OS (for example, Solaris 10)
- The release number of Sun Cluster (for example, Sun Cluster 3.2)

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

Command	Function
<code>prtconf -v</code>	Displays the size of the system memory and reports information about peripheral devices
<code>psrinfo -v</code>	Displays information about processors
<code>showrev -p</code>	Reports which patches are installed
<code>SPARC: prtdiag -v</code>	Displays system diagnostic information
<code>/usr/cluster/bin/clnode show-rev</code>	Displays Sun Cluster release and package version information

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

Example of Installing and Configuring a Sun Cluster Configuration

The *Sun Cluster Quick Start Guide for Solaris OS* provides an example of how to install and configure a specific Sun Cluster configuration. These guidelines and procedures are SPARC® specific but can be extrapolated for x86 based configurations. These example procedures can also be used as a guideline to configure other hardware and software configuration combinations, when used in conjunction with the Sun™ Cluster hardware, software, and data service manuals.

This book contains the following guidelines and procedures:

- “Configuration Specifications and Assumptions” on page 9
- “Task Map: Creating a Sun Cluster Quick Start Configuration” on page 12
- “Installing the Hardware” on page 13
- “Installing the Software” on page 17
- “Configuring the Cluster” on page 28
- “Configuring Volume Management” on page 29
- “Creating File Systems” on page 32
- “Installing and Configuring Application Software” on page 34
- “Configuring the Data Services” on page 39

Perform these procedures in the order that they are presented in this manual.

Configuration Specifications and Assumptions

This section provides descriptions of the specific cluster configuration that is used in this manual.

- “Hardware Configuration” on page 10
- “Software Configuration” on page 10
- “Public Network Addresses” on page 11
- “Procedure Assumptions” on page 12

Hardware Configuration

The procedures in the *Sun Cluster Quick Start Guide for Solaris OS* assume that the cluster consists of the following hardware and that the server installation is already performed.

TABLE 1-1 Hardware Specifications

Hardware Product	Components per Machine	Installation Instructions
Two Sun Fire™ V440 servers	At least 2 Gbytes of memory Two internal disks Two onboard ports, configured for the private interconnect Two Sun Quad GigaSwift Ethernet (QGE) cards, for connection to the public network and to the management network Two Fibre Channel-Arbitrated Loops (FC-AL) cards, for connection to the storage	<i>Sun Fire V440 Server Installation Guide</i>
One Sun StorEdge™ 3510 FC RAID array with dual controllers	Twelve 73-Gbyte physical drives	<i>Sun StorEdge 3000 Family Installation, Operation, and Service Manual, Sun StorEdge 3510 FC Array</i>
One Sun Ultra™ 20 workstation	One QGE card, for connection to the public network	<i>Sun Ultra 20 Workstation Getting Started Guide</i> (819-2148)

Software Configuration

The procedures in the *Sun Cluster Quick Start Guide for Solaris OS* assume that you have the following versions of software to install.

TABLE 1-2 Software Specifications

Product	Included Products	Product Subcomponents
Solaris 10 11/06 software for SPARC platforms	Apache HTTP Server version 1.3 software, secure using <code>mod_ssl</code> NFS version 3 software Solaris Volume Manager software Solaris multipathing functionality	

TABLE 1-2 Software Specifications (Continued)

Product	Included Products	Product Subcomponents
Sun Java™ Availability Suite software	Sun Cluster 3.2 core software	Cluster Control Panel (cconsole)
		Sun Cluster Manager
	Sun Cluster agent software	Sun Cluster HA for Apache
		Sun Cluster HA for NFS
		Sun Cluster HA for Oracle
Oracle 10gR2		

The procedures in this manual configure the following data services:

- Sun Cluster HA for Apache on a cluster file system
- Sun Cluster HA for NFS on a highly available local file system
- Sun Cluster HA for Oracle on a highly available local file system

Public Network Addresses

The procedures in the *Sun Cluster Quick Start Guide for Solaris OS* assume that public-network IP addresses are created for the following components.

Note – The IP addresses in the following table are for example only and are not valid for use on the public network.

The following addresses are used for communication with the public-network subnet 192 . 168 . 10.

TABLE 1-3 Public Network Example IP Addresses

Component	IP Address	Name
Cluster nodes	192 . 168 . 10 . 1	phys - sun
	192 . 168 . 10 . 2	phys - moon
Sun Cluster HA for Apache logical hostname	192 . 168 . 10 . 3	apache - lh
Sun Cluster HA for NFS logical hostname	192 . 168 . 10 . 4	nfs - lh
Sun Cluster HA for Oracle logical hostname	192 . 168 . 10 . 5	oracle - lh
Administrative console	192 . 168 . 10 . 6	admincon

The following addresses are used for communication with the management-network subnet, 192 . 168 . 11.

TABLE 1-4 Management Network Example IP Addresses

Component	IP Address	Name
Cluster nodes	192.168.11.1	phys-sun-11
	192.168.11.2	phys-moon-11
Sun StorEdge 3510 FC RAID array	192.168.11.3	se3510fc
Administrative console	192.168.11.4	admincon-11

Procedure Assumptions

The procedures in this manual were developed with the following assumptions:

- **Name service** - The cluster uses a name service.
- **Shell environment** - All commands and environment setup in this manual are for the C shell environment. If you use a different shell, replace any C shell-specific information or instructions with the appropriate information for your preferred shell environment.
- **User login** - Unless otherwise specified, perform all procedures as superuser.
- **cconsole** - These procedures assume that you are using the `cconsole` utility. Use this utility to access individual nodes and to issue commands to all nodes simultaneously by using the `cconsole` master window.
- **Command prompts** - The `phys-schost-N#` prompt indicates that you type the command in the `cconsole` master console window. This action issues the command on both cluster nodes simultaneously.

The `phys-sun#`, `phys-moon#`, and `admincon#` prompts indicate that you type the command only in the console window for the specified machine.

Task Map: Creating a Sun Cluster Quick Start Configuration

The following task map lists the tasks that you perform to create a Sun Cluster configuration for the hardware and software components that are specified in this manual. Complete the tasks in the order that they are presented in this table.

TABLE 1-5 Task Map: Creating a Sun Cluster Quick Start Configuration

Task	Instructions
1. Connect the administrative console, cluster nodes, and storage array. Configure the storage array.	“Installing the Hardware” on page 13

TABLE 1-5 Task Map: Creating a Sun Cluster Quick Start Configuration *(Continued)*

Task	Instructions
2. Install the Solaris OS and Cluster Control Panel software on the administrative console. Install the Solaris OS and Sun Cluster software and patches on the cluster nodes. Configure the Solaris OS and IPMP groups. Create state database replicas. Mirror the root file system. Set up the Oracle system groups and user.	“Installing the Software” on page 17
3. Establish the cluster and verify the configuration.	“Configuring the Cluster” on page 28
4. Configure Solaris Volume Manager and create disk sets.	“Configuring Volume Management” on page 29
5. Create the cluster file system and the highly available local file systems.	“Creating File Systems” on page 32
6. Configure the Apache HTTP Server software. Install and configure Oracle software.	“Installing and Configuring Application Software” on page 34
7. Use Sun Cluster Manager to configure Sun Cluster HA for Apache, Sun Cluster HA for NFS, and Sun Cluster HA for Oracle.	“Configuring the Data Services” on page 39

Installing the Hardware

Perform the following procedures to connect the cluster hardware components. See your hardware documentation for additional information and instructions.

- “How to Connect the Administrative Console” on page 14
- “How to Connect the Cluster Nodes” on page 14
- “How to Connect the Sun StorEdge 3510 FC RAID Array” on page 15
- “How to Configure the Storage Array” on page 16

The following figure illustrates the cabling scheme for this configuration.

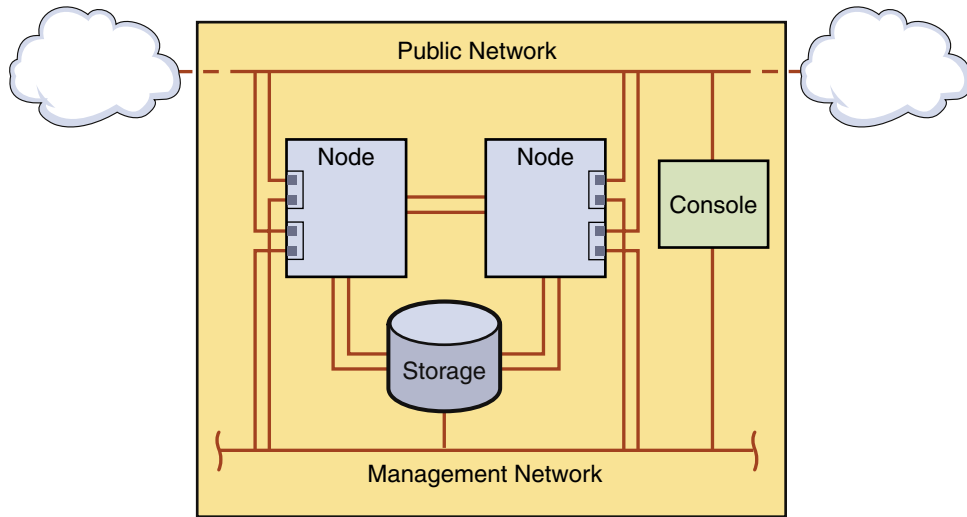


FIGURE 1-1 Cluster Topology and Cable Connections

▼ How to Connect the Administrative Console

For ease of installation, these example installation procedures apply to using an administrative console that is installed with Cluster Control Panel software. However, Sun Cluster software does not require that you use an administrative console. You can use other means to contact the cluster nodes, such as by using the `telnet` command to connect through the public network. Also, an administrative console does not have to be dedicated exclusively to use by a single cluster.

- 1 **Connect the administrative console to a management network that is connected to `phys - sun` and to `phys - moon`.**
- 2 **Connect the administrative console to the public network.**

▼ How to Connect the Cluster Nodes

- 1 **As the following figure shows, connect `ce0` and `ce9` on `phys - sun` to `ce0` and `ce9` on `phys - moon` by using switches.**

This connection forms the private interconnect.




FIGURE 1-2 Two-Node Cluster Interconnect

The use of switches in a two-node cluster permits ease of expansion if you decide to add more nodes to the cluster.

- 2 On each cluster node, connect from ce1 and ce5 to the public-network subnet.
- 3 On each cluster node, connect from ce2 and ce6 to the management network subnet.

▼ How to Connect the Sun StorEdge 3510 FC RAID Array

- 1 Connect the storage array to the management network.

Alternatively, connect the storage array by serial cable directly to the administrative console.

- 2 As the following figure shows, use fiber-optic cables to connect the storage array to the cluster nodes, two connections for each cluster node.

One node connects to a port on host channels 0 and 5. The other node connects to a port on host channels 1 and 4.

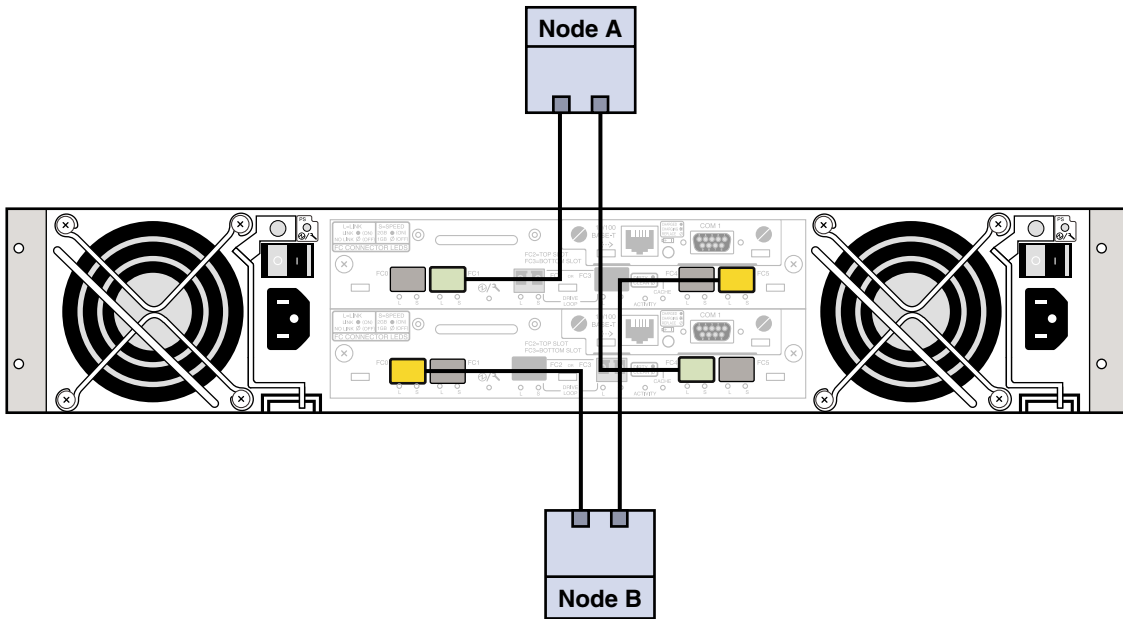


FIGURE 1-3 Sun StorEdge 3510 FC RAID Array Connection to Two Nodes

3 Power on the storage array and check LEDs.

Verify that all components are powered on and functional. Follow procedures in “First-Time Configuration for SCSI Arrays” in *Sun StorEdge 3000 Family Installation, Operation, and Service Manual, Sun StorEdge 3510 FC Array*.

▼ How to Configure the Storage Array

Follow procedures in the *Sun StorEdge 3000 Family RAID Firmware 4.1x User's Guide* to configure the storage array. Configure the array to the following specifications.

- 1 Create one global hot-spare drive from the unused physical drive.
- 2 Create two RAID-5 logical drives.
 - a. For redundancy, distribute the physical drives that you choose for each logical drive over separate channels.
 - b. Add six physical drives to one logical drive and assign the logical drive to the primary controller of the storage array, ports 0 and 5.
 - c. Add five physical drives to the other logical drive and assign the logical drive to the secondary controller, ports 1 and 4.

3 Partition the logical drives to achieve three partitions.

a. Allocate the entire six-drive logical drive to a single partition.

This partition will be for use by Sun Cluster HA for Oracle.

b. Create two partitions on the five-drive logical drive.

- Allocate 40% of space on the logical drive to one partition for use by Sun Cluster HA for NFS.
- Allocate 10% of space on the logical drive to the second partition for use by Sun Cluster HA for Apache.
- Leave 50% of space on the logical drive unallocated, for other use as needed.

4 Map each logical drive partition to a host logical unit number (LUN).

Partition Use	LUN
Oracle	LUN0
NFS	LUN1
Apache	LUN2

5 Note the World Wide Name (WWN) for each LUN.

You use this information when you create the disk sets later in this manual.

Installing the Software

Perform the following procedures to install the packages and patches for all software products and set up the user environment.

- [“How to Install the Administrative Console” on page 18](#)
- [“How to Install the Solaris Operating System” on page 19](#)
- [“How to Set Up the User Environment” on page 21](#)
- [“How to Configure the Operating System” on page 22](#)
- [“How to Create State Database Replicas” on page 25](#)
- [“How to Mirror the Root \(/\) File System” on page 25](#)
- [“How to Install Sun Cluster Software” on page 26](#)
- [“How to Set Up the Oracle System Groups and User” on page 27](#)

Note – You install the Oracle software later in this manual.

▼ How to Install the Administrative Console

Before You Begin Have the following available:

- The Java Availability Suite DVD-ROM that contains Sun Cluster 3.2 software.
- Access to the SunSolveSM web site to download patches.

1 Become superuser on the administrative console.

2 Configure the preinstalled Solaris 10 11/06 software, if you have not already done so.

For more information, see the *Sun Ultra 20 Workstation Getting Started Guide* (819–2148).

3 Download, install, and configure Sun Update Connection.

See <http://www.sun.com/service/sunupdate/gettingstarted.html>

(<http://www.sun.com/service/sunupdate/gettingstarted.html>) for details. Documentation for Sun Update Connection is available at <http://docs.sun.com/app/docs/coll/1320.2>

(<http://docs.sun.com/app/docs/coll/1320.2>).

4 Download and apply any Solaris 10 patches by using Sun Update Connection.

5 Load the Java Availability Suite DVD-ROM into the DVD-ROM drive.

6 Change to the `Solaris_sparc/Product/sun_cluster/Solaris_10/Packages/` directory.

7 Install software packages for the Cluster Control Panel and man pages.

```
admincon# pkgadd -d . SUNWcccon SUNWscman
```

8 Change to a directory that does *not* reside on the DVD-ROM and eject the DVD-ROM.

```
host# cd /
host# eject cdrom
```

9 Create an `/etc/cluster` file that contains the cluster name and the two node names.

```
admincon# vi /etc/cluster
scccluster phys-sun phys-moon
```

10 Create an `/etc/serialports` file that contains both node names and the hostname and port number that each node uses to connect to the management network.

```
admincon# vi /etc/serialports
phys-sun phys-sun 46
phys-moon phys-moon 47
```

11 Add the Sun Cluster PATH and MANPATH to the `.cshrc` user initialization file.

- To the PATH entry, add `/opt/SUNWcluster/bin`.
- To the MANPATH entry, add `/opt/SUNWcluster/man` and `/usr/cluster/man`.

12 Initialize your modifications.

```
admincon# cd
admincon# source .cshrc
```

▼ How to Install the Solaris Operating System

This procedure describes how to install the Solaris 10 OS to meet Sun Cluster software installation requirements.

Note – If your system comes with the Solaris OS preinstalled but does not meet Sun Cluster software installation requirements, perform this procedure to reinstall Solaris software to meet installation requirements.

Before You Begin Have the following available:

- The Solaris 10 11/06 Software for SPARC platforms DVD-ROM.
- Access to the SunSolve web site to download patches.

1 Add all public hostnames and logical addresses for the cluster to the naming service.

Note – The IP addresses in this step are for example only and are not valid for use on the public network. Substitute your own IP addresses when you perform this step.

192.168.10.1	phys-sun
192.168.10.2	phys-moon
192.168.10.3	apache-lh
192.168.10.4	nfs-lh
192.168.10.5	oracle-lh
192.168.10.6	admincon
192.168.11.1	phys-sun-11
192.168.11.2	phys-moon-11
192.168.11.3	se3510fc
192.168.11.4	admincon-11

For more information about naming services, see *System Administration Guide: Naming and Directory Services (DNS, NIS, and LDAP)*.

2 From the administrative console, start the cconsole(1M) utility.

```
admincon# cconsole &
```

Use the `cconsole` utility to communicate with each individual cluster node or use the master window to send commands to both nodes simultaneously.

3 Insert the Solaris 10 11/06 DVD-ROM in the DVD-ROM drive of phys-sun.

4 Access the console window for phys - sun.**5 Boot phys - sun.**

- If the system is new, turn on the system.
- If the system is currently running, shut down the system.

```
phys-sun# init 0
```

The ok prompt is displayed.

6 Disable automatic reboot.

```
ok setenv auto-boot? false
```

Disabling automatic reboot prevents continuous boot cycling.

7 Create an alias for each disk.

The assignment of aliases to the disks enables you to access and boot from the second disk if you cannot boot from the default disk.

a. Display the disks and choose the boot disk.

```
ok show-disks
```

```
...
```

```
Enter selection, q to quit: X
```

b. Assign the alias name rootdisk to the disk that you chose.

```
ok nvalias rootdisk Control-Y
```

The Control-Y keystroke combination enters the disk name that you chose from the show-disks menu.

c. Save the disk alias.

```
ok nvstore
```

d. Repeat the preceding steps to identify and assign the alias name backup_root to the alternate boot disk.**e. Set the boot-device environment variable to the aliases for the default boot disk and backup boot disk.**

```
ok setenv boot-device rootdisk backup_root
```

For more information, see *OpenBoot 4.x Command Reference Manual*.

8 Start the Solaris installation program.

```
ok boot cdrom
```

9 Follow the prompts.

- Make the following installation choices:

Prompt	Value
Solaris Software Group	Entire Plus OEM Support
Partitions	Manual formatting
Root password	Same password on both nodes
Automatic reboot	No

- Set the following partition sizes and file-system names, if not already set:

Slice	Size	File System Name
0	remaining free space	/
1	2 Gbyte	swap
4	512 Mbyte	/globaldevices
5	2 Gbyte	/var
7	32 Mbyte	for Solaris Volume Manager use

10 Return to [Step 3](#) and repeat these steps on `phys - moon`.**11 On both nodes, download, install, and configure Sun Update Connection.**

See <http://www.sun.com/service/sunupdate/gettingstarted.html> (<http://www.sun.com/service/sunupdate/gettingstarted.html>) for details. Documentation for Sun Update Connection is available at <http://docs.sun.com/app/docs/coll/1320.2> (<http://docs.sun.com/app/docs/coll/1320.2>).

12 On both nodes, download and apply any Solaris 10 patches by using Sun Update Connection.

▼ How to Set Up the User Environment

Perform this procedure on both nodes. The steps in this procedure use the C shell environment. If you are using a different shell, perform the equivalent tasks for your preferred shell environment.

For more information, see “Customizing a User’s Work Environment” in *System Administration Guide: Basic Administration*.

1 Open the cconsole master console window, if it is not already open.

Use the master console window to perform the steps in this procedure on both nodes at the same time.

2 Display the settings for the umask and the environment variables.

```
phys-X# umask
phys-X# env | more
```

3 If not already set, set the umask to 22.

This entry sets the default permissions for newly created files.

```
umask 022
```

4 Ensure that the PATH includes the following paths.

- /usr/bin
- /usr/cluster/bin
- /usr/sbin
- /usr/ccs/bin
- /oracle/oracle/product/10.2.0/bin

5 (Optional) Add the following paths to the MANPATH.

- /usr/cluster/man
- /usr/apache/man

6 Set the ORACLE_BASE and ORACLE_SID environment variables.

```
ORACLE_BASE=/oracle
ORACLE_SID=orasrvr
```

7 Verify the setting changes that you made.

```
phys-X# umask
phys-X# env | more
```

▼ How to Configure the Operating System

This procedure describes how to modify certain system settings to support the Quick Start configuration.

1 On both nodes, enable Solaris multipathing functionality.

```
phys-X# /usr/sbin/stmsboot -e
```

-e Enables Solaris I/O multipathing

For more information, see the `stmsboot(1M)` man page.

2 On both nodes, update the `/etc/inet/ipnodes` file with all public hostnames and logical addresses for the cluster.

Except for the `loghost` entries, these entries are the same on both nodes.

Note – The IP addresses in this step are for example only and are not valid for use on the public network. Substitute your own IP addresses when you perform this step.

`phys-X# vi /etc/inet/ipnodes`

■ **On `phys-sun`, add the following entries:**

```
127.0.0.1      localhost
192.168.10.1   phys-sun  loghost
192.168.10.2   phys-moon
192.168.10.3   apache-lh
192.168.10.4   nfs-lh
192.168.10.5   oracle-lh
192.168.10.6   admincon

192.168.11.1   phys-sun-11
192.168.11.2   phys-moon-11
192.168.11.3   se3510fc-11
192.168.11.4   admincon-11
```

■ **On `phys-moon`, add the following entries:**

```
127.0.0.1      localhost
192.168.10.1   phys-sun
192.168.10.2   phys-moon  loghost
192.168.10.3   apache-lh
192.168.10.4   nfs-lh
192.168.10.5   oracle-lh
192.168.10.6   admincon

192.168.11.1   phys-sun-11
192.168.11.2   phys-moon-11
192.168.11.3   se3510fc-11
192.168.11.4   admincon-11
```

3 On both nodes, ensure that the following kernel parameters are set to at least the minimum values that Oracle requires.

a. Display the settings for the `default` project.

`phys-X# prctl -i project default`

- b. If no kernel parameters are set, or if any kernel parameters are not set to the minimum required value for Oracle as shown in the following table, set the parameter.**

```
phys-X# projmod -s -K "parameter=(priv,value,deny)" default
```

Oracle Kernel Parameter	Minimum Required Value
process.max-sem-nsems	256
project.max-sem-ids	100
project.max-shm-ids	100
project.max-shm-memory	4294967295

- c. Verify the new settings.**

```
phys-X# prctl -i project default
```

These settings are the minimum required values to support the Oracle software in a Sun Cluster Quick Start configuration. For more information about these parameters, see the *Oracle10g Installation Guide*.

- 4 On both nodes, add the following entries to the /etc/system file.**

```
phys-X# vi /etc/system
set ce:ce_taskq_disable=1
exclude:lofs
```

- The first entry supports ce adapters for the private interconnect.
- The second entry disables the loopback file system (LOFS), which must be disabled when Sun Cluster HA for NFS is configured on a highly available local file system. For more information and alternatives to disabling LOFS when Sun Cluster HA for NFS is configured, see the information about loopback file systems in “Solaris OS Feature Restrictions” in *Sun Cluster Software Installation Guide for Solaris OS*.

These changes take effect at the next system reboot.

- 5 On both nodes, set NFS version 3 as the default version.**

- a. Add the following entry to the /etc/default/nfs file.**

```
NFS_SERVER_VERSMAX=3
```

- b. Disable the NFS service.**

```
phys-X# svcadm disable network/nfs/server
```

- c. Re-enable the NFS service.**

```
phys-X# svcadm enable network/nfs/server
```


- 6 On both nodes, update the /devices and /dev entries.**

```
phys-X# devfsadm -C
```

- 7 On both nodes, confirm that the storage array is visible.**

```
phys-X# luxadm probe
```

▼ How to Create State Database Replicas

This procedure assumes that the specified disks are available for creation of database replicas. Substitute your own disk names in this procedure.

- 1 On both nodes, create state database replicas.**

Create three replicas on each of the two internal disks.

```
phys-X# metadb -af -c 3 c0t0d0s7
```

```
phys-X# metadb -a -c 3 c0t1d0s7
```

- 2 On both nodes, verify the replicas.**

```
phys-X# metadb
```

flags		first blk	block count	
a	u	16	8192	/dev/dsk/c0t0d0s7
a	u	8208	8192	/dev/dsk/c0t0d0s7
a	u	16400	8192	/dev/dsk/c0t0d0s7
a	u	16	8192	/dev/dsk/c0t1d0s7
a	u	8208	8192	/dev/dsk/c0t1d0s7
a	u	16400	8192	/dev/dsk/c0t1d0s7

▼ How to Mirror the Root (/) File System

Perform this procedure on one node at a time.

This procedure assumes that the cluster node contains the internal nonshared disks c0t0d0 and c0t1d0. Substitute your own internal disk names if necessary in the steps of this procedure.

- 1 On phys-sun, place the root slice c0t0d0s0 in a single-slice (one-way) concatenation.**

```
phys-sun# metainit -f d10 1 1 c0t0d0s0
```

- 2 Create a second concatenation with the other internal disk, c0t1d0s0.**

```
phys-sun# metainit d20 1 1 c0t1d0s0
```

- 3 Create a one-way mirror with one submirror.**

```
phys-sun# metainit d0 -m d10
```

4 Set up the system files for the root directory.

```
phys-sun# metaroot d0
```

The `metaroot` command edits the `/etc/vfstab` and `/etc/system` files so that the system can be booted with the root (`/`) file system on a metadevice or volume. For more information, see the `metaroot(1M)` man page.

5 Flush all file systems.

```
phys-sun# lockfs -fa
```

The `lockfs` command flushes all transactions from the log and writes the transactions to the master file system on all mounted UFS file systems. For more information, see the `lockfs(1M)` man page.

6 Reboot the node to remount the newly mirrored root (`/`) file system.

```
phys-sun# init 6
```

7 Attach the second submirror to the mirror.

```
phys-sun# metattach d0 d20
```

For more information, see the `metattach(1M)` man page.

8 Record the alternate boot path for possible future use.

If the primary boot device fails, you can then boot from this alternate boot device. For more information about alternate boot devices, see “Creating a RAID-1 Volume” in *Solaris Volume Manager Administration Guide*.

```
phys-sun# ls -l /dev/rdisk/c0t1d0s0
```

9 Repeat Step 1 through Step 8 on `phys-moon`.

▼ How to Install Sun Cluster Software

This procedure installs software packages for the Sun Cluster framework and for the Sun Cluster HA for Apache, Sun Cluster HA for NFS, and Sun Cluster HA for Oracle data services.

Before You Begin Have available the following:

- The Java Availability Suite DVD-ROM that contains Sun Cluster 3.2 software.
- Access to the SunSolve web site to download patches.

1 On `phys-sun`, load the Java Availability Suite DVD-ROM in the DVD-ROM drive.**2 Start the Java Enterprise System (ES) installer program.**

```
phys-sun# ./installer
```

For more information about using the Java ES installer program, see the *Sun Java Enterprise System 5 Installation Guide for UNIX*.

3 Follow the onscreen instructions to install the Sun Cluster framework packages.

Screen Name	Instructions
Software License Agreement	Accept the license agreement.
Language Support	Choose any languages that you want to install in addition to English.
Installation Type	Answer no when asked if you want to install the full set of Java ES software.
Component Selection	<p>Choose Sun Cluster and Sun Cluster Agents. Do <i>not</i> deselect Sun Cluster Manager. Confirm your selection when prompted.</p> <p>Follow the onscreen instructions to install the following data service packages:</p> <ul style="list-style-type: none"> ▪ Sun Cluster HA for Apache ▪ Sun Cluster HA for NFS ▪ Sun Cluster HA for Oracle
Shared Component Upgrades Required	Accept upgrade of the list of shared components.
Configuration Type	Choose Configure Later .

After the installation is finished, the `installer` program provides an installation summary. This summary enables you to view logs that the program created during the installation. These logs are located in the `/var/sadm/install/logs/` directory.

4 Change to a directory that does *not* reside on the DVD-ROM and eject the DVD-ROM.

```
host# cd /
host# eject cdrom
```

5 Return to [Step 1](#) and repeat all steps on `phys-moon`.

6 On both nodes, use Sun Update Connection to download and apply any needed patches.

▼ How to Set Up the Oracle System Groups and User

Perform the steps in this procedure on both nodes.

1 Open the `cconsole` master console window, if it is not already open.

Use the master console window to perform the steps in this procedure on both nodes at the same time.

2 Create the Oracle Inventory group, `oinstall`, and the database administrator group, `dba`.

```
phys-X# groupadd oinstall
phys-X# groupadd dba
```

3 Create the Oracle user account, oracle.

Specify the Oracle home directory, `/oracle/oracle/product/10.2.0`. Set `dba` as the primary group and set `oinstall` as the secondary group.

```
phys-X# useradd -g dba -G oinstall -d /oracle/oracle/product/10.2.0 oracle
```

4 Set the oracle password.

```
phys-X# passwd -r files oracle
```

Configuring the Cluster

Perform the following procedure to establish the cluster.

▼ How to Establish the Cluster

1 From phys-moon, start the interactive scinstall utility.

```
phys-moon# scinstall
```

The `scinstall` Main Menu is displayed.

2 Type the number that corresponds to the option for Create a new cluster or new cluster node and press the Return key.

The New Cluster and Cluster Node Menu is displayed.

3 Type the number that corresponds to the option for Create a new cluster and press the Return key.

The Typical or Custom Mode menu is displayed.

4 Type the number that corresponds to the option for Typical and press the Return key.**5 Follow the menu prompts to supply the following information:**

Note – The adapter names that are used in the following table are arbitrarily selected for this example only.

Component	Description	Answer
Cluster Name	What is the name of the cluster that you want to establish?	<code>sccluster</code>
Cluster Nodes	List the names of the other nodes.	<code>phys-sun</code>

Component	Description	Answer
Cluster Transport Adapters and Cables	What are the names of the two cluster transport adapters that attach the node to the private interconnect?	ce0, ce9
Quorum Configuration	Do you want to disable automatic quorum device selection?	No
Check	Do you want to interrupt installation for sccheck errors?	No

The `scinstall` utility configures the cluster and reboots both nodes. It also automatically creates a link-based multiple-adapter IPMP group for each set of public-network adapters in the cluster that use the same subnet. The cluster is established when both nodes have successfully booted into the cluster. Sun Cluster installation output is logged in a `/var/cluster/logs/install/scinstall.log.N` file.

6 From `phys-sun`, verify that the nodes and the quorum device are successfully configured.

If the cluster is successfully established, you will see output similar to the following.

```
phys-sun# clquorum list
d5
phys-sun
phys-moon
```

Configuring Volume Management

Perform the following procedures to configure volume management.

- “How to Create Disk Sets” on page 29
- “How to Add LUNs to Disk Sets” on page 30
- “How to Create and Activate an `md.tab` File” on page 31

▼ How to Create Disk Sets

1 From `phys-sun`, create one disk set for each data service that you will configure.

a. Make `phys-sun` the primary node for the Apache and NFS data services.

```
phys-sun# metaset -s nfsset -a -h phys-sun phys-moon
phys-sun# metaset -s apacheset -a -h phys-sun phys-moon
```

b. Make `phys-moon` the primary node for the Oracle data service.

```
phys-sun# metaset -s oraset -a -h phys-moon phys-sun
```

2 Verify that the configuration of the disk sets is correct and visible to both nodes.

```
phys-X# metaset
Set name = nfsset, Set number = 1
...
```

```

Set name = apacheset, Set number = 2
...
Set name = oraset, Set number = 3
...

```

▼ How to Add LUNs to Disk Sets

1 From phys-sun, list the DID mappings.

Output is similar to the following, where *WWN* stands for the unique World Wide Number (WWN) of the disk target.

```

phys-sun# clddevice show | grep Device
=== DID Device Instances ===
DID Device Name:                /dev/did/rdisk/d1
  Full Device Path:              phys-sun:/dev/rdisk/c0t0d0
DID Device Name:                /dev/did/rdisk/d2
  Full Device Path:              phys-sun:/dev/rdisk/c0t6d0
DID Device Name:                /dev/did/rdisk/d3
  Full Device Path:              phys-sun:/dev/rdisk/c1t WWNd0
  Full Device Path:              phys-moon:/dev/rdisk/c1t WWNd0
DID Device Name:                /dev/did/rdisk/d4
  Full Device Path:              phys-sun:/dev/rdisk/c1t WWNd0
  Full Device Path:              phys-moon:/dev/rdisk/c1t WWNd0
DID Device Name:                /dev/did/rdisk/d5
  Full Device Path:              phys-sun:/dev/rdisk/c0t WWNd0
  Full Device Path:              phys-moon:/dev/rdisk/c0t WWNd0
...

```

2 Map LUN0, LUN1, and LUN2 to their DID device names.

Compare the information that you saved when you created the LUNs with the output of the `clddevice` command. For each LUN, locate the `/dev/rdisk/cNt WWNdY` name that is associated with the LUN. Then find that same disk name in the `clddevice` output to determine the DID device name.

These procedures assume the following mappings for the purposes of this example. Substitute your own disk names and DID names when you perform the remainder of these procedures.

Data Service	LUN Name	Raw Disk Device Name	DID Name
Sun Cluster HA for Oracle	LUN0	/dev/did/rdisk/c1t WWNd0	dsk/d3
Sun Cluster HA for NFS	LUN1	/dev/did/rdisk/c1t WWNd0	dsk/d4
Sun Cluster HA for Apache	LUN2	/dev/did/rdisk/c0t WWNd0	dsk/d5

3 Take ownership of the Oracle disk set oraset.

```
phys-sun# cldevicegroup switch -n phys-sun oraset
```

4 Add LUN0 to the Oracle disk set.

Use the full DID path name.

```
phys-sun# metaset -s oraset -a /dev/did/rdisk/d3
```

5 Verify that the configuration of the disk set is correct.

```
phys-sun# metaset -s oraset
```

6 Repeat the process to add LUN1 to the NFS disk set nfsset.

```
phys-sun# cldevicegroup switch -n phys-sun nfsset
```

```
phys-sun# metaset -s nfsset -a /dev/did/rdisk/d4
```

```
phys-sun# metaset -s nfsset
```

7 Repeat the process to add LUN2 to the Apache disk set apacheset.

```
phys-sun# cldevicegroup switch -n phys-sun apacheset
```

```
phys-sun# metaset -s apacheset -a /dev/did/rdisk/d5
```

```
phys-sun# metaset -s apacheset
```

▼ How to Create and Activate an md.tab File

1 On both nodes, create an /etc/lvm/md.tab file with the following entries.

These entries define the volumes for each disk set. The one-way mirrors provide flexibility to add a mirror later without unmounting the file system. You can create the file on one node and copy it to the other node, or you can create it on both nodes at the same time by using the `cconsole(1M)` utility.

```
apacheset/d0 -m apacheset/d10
    apacheset/d10 1 1 /dev/did/rdisk/d3s0
```

```
nfsset/d1 -m nfsset/d11
    nfsset/d11 1 1 /dev/did/rdisk/d4s0
```

```
oraset/d2 -m oraset/d12
    oraset/d12 1 1 /dev/did/rdisk/d5s0
```

```
oraset/d0 -p oraset/d2 3G
oraset/d1 -p oraset/d2 3G
```

2 From phys-sun, take ownership of each of the disk sets and activate their volumes.

```
phys-sun# cldevicegroup switch -n phys-sun apacheset
```

```
phys-sun# metainit -s apacheset -a
```

```
phys-sun# cldevicegroup switch -n phys-sun nfsset
phys-sun# metainit -s nfsset -a
```

```
phys-moon# cldevicegroup switch -n phys-sun oraset
phys-moon# metainit -s oraset -a
```

3 Check the status of the volumes for each disk set.

```
phys-sun# metastat
...
Status: Okay
...
```

Creating File Systems

Perform the following procedure to create a cluster file system and local file systems to support the data services.

▼ How to Create File Systems

This procedure creates a cluster file system for use by Sun Cluster HA for Apache and local file systems for use by Sun Cluster HA for NFS and Sun Cluster HA for Oracle. Later in this manual, the local file systems are configured as highly available local file systems by using HAStoragePlus.

1 From phys-sun, create the UFS file systems.

```
phys-sun# newfs /dev/md/apacheset/rdisk/d0
phys-sun# newfs /dev/md/nfsset/rdisk/d1
phys-sun# newfs /dev/md/oraset/rdisk/d0
phys-sun# newfs /dev/md/oraset/rdisk/d1
```

2 On each node, create a mount-point directory for each file system.

```
phys-X# mkdir -p /global/apache
phys-X# mkdir -p /local/nfs
phys-X# mkdir -p /oracle/oracle/product/10.2.0
phys-X# mkdir -p /oradata/10gR2
```

3 For the Oracle home directory and database directory, set the owner, group, and mode.

a. Set the owner as oracle and the group as dba.

```
phys-X# chown -R oracle:dba /oracle/oracle/product/10.2.0
phys-X# chown -R oracle:dba /oradata/10gR2
```

b. Make the Oracle directories writable only by the owner and the group.

```
phys-X# chmod -R 775 /oracle/oracle/product/10.2.0
phys-X# chmod -R 775 /oradata/10gR2
```

4 On each node, add an entry to the `/etc/vfstab` file for each mount point.

Note – Only the cluster file system for Apache uses the `global` mount option. Do not specify the `global` mount option for the local file systems for NFS and Oracle.

```
phys-X# vi /etc/vfstab
#device          device          mount FS      fsck   mount  mount
#to mount        to fsck         point type   pass   at boot options
#
/dev/md/apacheset/dsk/d0 /dev/md/apacheset/rdisk/d0 /global/apache ufs 2 yes global,logging
/dev/md/nfsset/dsk/d1 /dev/md/nfsset/rdisk/d1 /local/nfs ufs 2 no logging
/dev/md/oraset/dsk/d0 /dev/md/oraset/rdisk/d0 /oracle/oracle/product/10.2.0 ufs 2 no logging
/dev/md/oraset/dsk/d1 /dev/md/oraset/rdisk/d1 /oradata/10gR2 ufs 2 no logging,forcedirectio
```

5 From `phys-sun`, verify that the mount points exist.

```
phys-sun# cluster check
```

If no errors occur, nothing is returned.

6 From `phys-sun`, mount the file systems.

```
phys-sun# mount /global/apache
```

```
phys-sun# mount /local/nfs
```

```
phys-sun# mount /oracle/oracle/product/10.2.0
```

```
phys-sun# mount /oradata/10gR2
```

7 On each node, verify that the file systems are mounted.

Note – Only the cluster file system for Apache is displayed on both nodes.

```
phys-sun# mount
```

```
...
```

```
/global/apache on /dev/md/apacheset/dsk/d0 read/write/setuid/global/logging
on Sun Oct 3 08:56:16 2005
```

```
/local/nfs on /dev/md/nfsset/dsk/d1 read/write/setuid/logging
on Sun Oct 3 08:56:16 2005
```

```
/oracle/oracle/product/10.2.0 on /dev/md/oraset/dsk/d0 read/write/setuid/logging
on Sun Oct 3 08:56:16 2005
```

```
/oradata/10gR2 on /dev/md/oraset/dsk/d1 read/write/setuid/logging/forcedirectio
on Sun Oct 3 08:56:16 2005
```

```
phys-moon# mount
```

```
...
```

```
/global/apache on /dev/md/apacheset/dsk/d0 read/write/setuid/global/logging
on Sun Oct 3 08:56:16 2005
```

Installing and Configuring Application Software

Perform the following procedures to configure Apache software, install Oracle software, and configure the Oracle database.

- “How to Configure Apache HTTP Server Software” on page 34
- “How to Install Oracle 10gR2 Software” on page 35
- “How to Create an Oracle Database” on page 37
- “How to Set Up Oracle Database Permissions” on page 38

▼ How to Configure Apache HTTP Server Software

This procedure configures secure Apache HTTP Server version 1.3 software by using `mod_ssl`. For additional information, see the installed Apache online documentation at <file:///usr/apache/htdocs/manual/index.html>, the Apache HTTP Server web site at <http://httpd.apache.org/docs/1.3/>, and the Apache `mod_ssl` web site at <http://www.modssl.org/docs/>.

1 Use the `cconsole` master window to access both nodes.

You can perform the next steps on both nodes at the same time.

2 Modify the `/etc/apache/httpd.conf` configuration file.

a. If necessary, copy the `/etc/apache/httpd.conf-example` template as `/etc/apache/httpd.conf`.

b. Set the following directives:

Apache Directive	Value
<code>ServerType</code>	<code>Standalone</code>
<code>ServerName</code>	<code>apache-lh</code>
<code>DocumentRoot</code>	<code>/var/apache/htdocs</code>

3 Install all certificates and keys.

4 In the `/usr/apache/bin` directory, create the file `keypass`.

Set file permissions for owner access only.

```
phys-X# cd /usr/apache/bin
phys-X# touch keypass
phys-X# chmod 700 keypass
```

5 Edit the `keypass` file so that it prints the pass phrase for the encrypted key that corresponds to a host and a port.

This file will be called with `server:port algorithm` as arguments. Ensure that the file can print the pass phrase for each of your encrypted keys when called with the correct parameters.

Later, when you attempt to start the web server manually, it must not prompt you for a pass phrase. For example, suppose that a secure web server is listening on ports 8080 and 8888, with private keys for both ports that are encrypted by using RSA. The `keypass` file could be the following:

```
# !/bin/ksh
host='echo $1 | cut -d: -f1'
port='echo $1 | cut -d: -f2'
algorithm=$2

if [ "$host" = "apache-lh.example.com" -a "$algorithm" = "RSA" ]; then
    case "$port" in
        8080) echo passphrase-for-8080;;
        8888) echo passphrase-for-8888;;
        esac
    fi
```

6 Update the paths in the Apache start/stop script file, `/usr/apache/bin/apachectl`, if they differ from your Apache directory structure.

7 Verify your configuration changes.

a. Check the `/etc/apache/httpd.conf` file for correct syntax.

```
phys-X# /usr/apache/bin/apachectl configtest
```

b. Ensure that any logical hostnames or shared addresses that Apache uses are configured and online.

c. On `phys-sun`, start the Apache server.

```
phys-sun# /usr/apache/bin/apachectl startssl
```

- Ensure that the web server does not ask you for a pass phrase.
- If Apache does not start properly, correct the problem.

d. On `phys-sun`, stop the Apache server.

```
phys-sun# /usr/apache/bin/apachectl stopssl
```

▼ How to Install Oracle 10gR2 Software

Before You Begin Have available the following:

- The Oracle 10gR2 installation DVD.

- Any Oracle patches.
- Oracle installation documentation.

1 On phys-sun, become user oracle.

```
phys-sun# su - oracle
```

2 Change to the /tmp directory.

```
phys-sun# cd /tmp
```

3 Insert the Oracle product disc.

If the volume management daemon `volld(1M)` is running and is configured to manage DVD-ROMs, the daemon automatically mounts the Oracle 10gR2 DVD-ROM on the `/cdrom/cdrom0` directory.

4 Start the Oracle Universal Installer.

```
phys-sun# /cdrom/cdrom0/Disk1/runInstaller
```

For more information about using the Oracle Universal Installer, see the *Oracle Database Client Installation Guide for Solaris Operating System (SPARC 64-Bit)*.

5 Follow the prompts to install Oracle software.

Specify the following values:

Oracle Component	Value
Source file location	/cdrom/cdrom0/Disk1/products.jar
Destination file location (<i>the value of \$ORACLE_HOME</i>)	/oracle/oracle/product/10.2.0
UNIX group name	dba
Available products	Oracle 10g Enterprise Edition or Standard Edition
Database configuration type	General Purpose
Installation type	Typical
Global database name	orasrvr
Oracle System Identifier (SID)	orasrvr
Database file location	/oradata/10gR2
Database character set	default

For more information, see the *Oracle Database Client Installation Guide for Solaris Operating System (SPARC 64-Bit)*.

6 Change to a directory that does *not* reside on the DVD and eject the DVD.

```
phys-sun# eject cdrom
```

- 7 **Apply any Oracle patches.**
- 8 **Verify that the owner, group, and mode of the `/oracle/oracle/product/10.2.0/bin/oracle` file are correct.**

```
phys-sun# ls -l /oracle/oracle/product/10.2.0/bin/oracle
-rwsr-s--x 1 oracle dba 3195 Apr 27 2005 oracle
```
- 9 **Verify that the listener binaries exist in the `/oracle/oracle/product/10.2.0/bin/` directory.**
Oracle listener binaries include the `lsnrctl` command and the `tnsping` command.
- 10 **Exit from the user `oracle`.**
The superuser prompt is again displayed.
- 11 **Prevent the Oracle `cssd` daemon from being started.**
Remove the following entry from the `/etc/inittab` file. This action prevents unnecessary error messages from being displayed.

```
h1:23:respawn:/etc/init.d/init.cssd run >/dev/null 2>&| > </dev/null
```
- 12 **Repeat this procedure on `phys-moon`.**

▼ How to Create an Oracle Database

Before You Begin Have available your Oracle installation documentation. Refer to those procedures to perform the following tasks.

- 1 **On `phys-sun`, prepare the database configuration files.**
 - Place all of the database-related files (data files, redo log files, and control files) on the `/oradata/10gR2` directory.
 - Within the `init$ORACLE_SID.ora` file or the `config$ORACLE_SID.ora` file, modify the assignments for `control_files` and `background_dump_dest` to specify the location of the control files.
- 2 **Start the creation of the database by using a utility from the following list:**
 - The Oracle Database Configuration Assistant (DBCA)
 - The Oracle `sqlplus(1M)` command

During creation, ensure that all of the database-related files are placed in the `/oradata/10gR2` directory.
- 3 **Verify that the file names of your control files match the file names in your configuration files.**

4 Create the `v$sysstat` view.

Run the catalog scripts that create the `v$sysstat` view. The Sun Cluster HA for Oracle fault monitor uses this view.

▼ How to Set Up Oracle Database Permissions

Perform this procedure on both nodes.

1 Enable access for the Oracle user and password to be used for fault monitoring.

Use the Oracle authentication method to grant to the `oracle` user authority on the `v_$sysstat` view and the `v_$archive_dest` view.

```
phys-X# sqlplus "/ as sysdba"
```

```
sql> grant connect, resource to user identified by passwd;  
sql> alter user oracle default tablespace system quota 1m on system;  
sql> grant select on v_$sysstat to oracle;  
sql> grant select on v_$archive_dest to oracle;  
sql> grant create session to oracle;  
sql> grant create table to oracle;
```

```
sql> exit;  
#
```

2 Configure NET8 for the Sun Cluster software.

a. Set the following entries in the default

`/oracle/oracle/product/10.2.0/network/admin/listener.ora` file.

```
HOST = oracle-lh  
POST = 1521
```

b. Set the same entries in the default

`/oracle/oracle/product/10.2.0/network/admin/tnsnames.ora` file.

Note – The values that you set in the `listener.ora` file and in the `tnsnames.ora` file must be the same.

Configuring the Data Services

Perform the following procedures to use Sun Cluster Manager to configure the data services.

- “How to Start Sun Cluster Manager” on page 39
- “How to Configure the Scalable Sun Cluster HA for Apache Data Service” on page 39
- “How to Configure the Sun Cluster HA for NFS Data Service” on page 40
- “How to Configure the Sun Cluster HA for Oracle Data Service” on page 41

▼ How to Start Sun Cluster Manager

Alternatively, you can run the `clsetup` utility to use the equivalent text-based interface.

- 1 From the administrative console, start a browser.
- 2 Connect to the Sun Java Web Console port on `phys-sun`.
`https://phys-sun:6789`
- 3 From the Sun Java Web Console screen, choose the Sun Cluster Manager link.
- 4 From the Sun Cluster Manager screen, choose **Tasks** from the sidebar.

▼ How to Configure the Scalable Sun Cluster HA for Apache Data Service

- 1 From the Sun Cluster Manager **Tasks** screen, under **Configure Data Services for Applications**, choose **Apache Web Server**.

The configuration wizard is displayed.

- 2 Follow the prompts to configure a scalable Sun Cluster HA for Apache data service.

Specify the following information. Otherwise, accept the default.

Component	Value
Apache configuration mode	Scalable Mode
Nodes or zones	<code>phys-sun, phys-moon</code>
Apache configuration file	<code>/etc/apache/httpd.conf</code>
Apache document root directory	Click Next to copy <code>/var/apache/htdocs</code> to a highly available file system

Component	Value
Cluster file-system mount point	/global/apache
Network resource	apache-lh

When all information is supplied, the wizard creates the data service and displays the commands that were used. The wizard performs validation checks on all Apache properties.

▼ How to Configure the Sun Cluster HA for NFS Data Service

- 1 From the Sun Cluster Manager Tasks screen, under Configure Data Services for Applications, choose NFS.

The configuration wizard is displayed.

- 2 Follow the prompts to configure a Sun Cluster HA for NFS data service.

Specify the following information. Otherwise, accept the default.

Component	Value
Node list	phys-sun, phys-moon
Logical hostname	nfs-lh
File-system mount point	/local/nfsset
Path prefix	/local/nfsset
Share options	
Access permissions	rw
nosuid	Off
Security	Default
Path	/local/nfsset

When all information is supplied, the wizard creates the data service and displays the commands that were used.

▼ How to Configure the Sun Cluster HA for Oracle Data Service

- 1 From the Sun Cluster Manager Tasks screen, under Configure Data Services for Applications, choose Oracle.

The configuration wizard is displayed.

- 2 Follow the prompts to configure the Sun Cluster HA for Oracle data service.

Specify the following information. Otherwise, accept the default.

Component	Value
Node list	phys-moon, phys-sun
Oracle components to configure	Server and Listener
Oracle home directory	/oracle/oracle/product/10.2.0
Oracle system identifier (SID)	orasrvr
Sun Cluster resource properties	
Alert_log_file	/oracle/oracle/product/10.2.0/alert_log
Connect_string	oracle/oracle-password
Server:Debug_level	1
Listener_name	LISTENER
Listener:Debug_level	1
Logical hostname	oracle-lh

When all information is supplied, the wizard creates the data service and displays the commands that were used. The wizard performs validation checks on all Oracle properties.

- 3 Log out of Sun Cluster Manager.

Next Steps Installation and configuration of your Sun Cluster Quick Start configuration is complete. Information about administering your cluster is available in the following documentation:

Topic	Documentation
Hardware	<p data-bbox="569 215 1191 241"><i>Sun Cluster 3.1 - 3.2 Hardware Administration Manual for Solaris OS</i></p> <p data-bbox="569 258 1292 314"><i>Sun Cluster 3.1 - 3.2 With Sun StorEdge 3510 or 3511 FC RAID Array Manual for Solaris OS</i></p>
Cluster Software	<p data-bbox="569 336 1071 362"><i>Sun Cluster System Administration Guide for Solaris OS</i></p>
Data Services	<p data-bbox="569 381 1253 407"><i>Sun Cluster Data Services Planning and Administration Guide for Solaris OS</i></p> <p data-bbox="569 425 1078 451"><i>Sun Cluster Data Service for Apache Guide for Solaris OS</i></p> <p data-bbox="569 468 1049 494"><i>Sun Cluster Data Service for NFS Guide for Solaris OS</i></p> <p data-bbox="569 512 1071 538"><i>Sun Cluster Data Service for Oracle Guide for Solaris OS</i></p>

Index

A

- adding
 - logical addresses to the ipnodes file, 23
 - LUNs to disk sets, 30-31
 - public-network IP addresses
 - ipnodes file, 23
 - naming service, 19
- administrative console
 - adding
 - to MANPATH, 18
 - to PATH, 18
 - connecting to the networks, 14
 - installing
 - patches, 18
 - software, 18-19
 - specifications, 10
- alternate boot path, displaying, 26
- Apache software
 - apachectl script, 35
 - configuration file, 34
 - configuring, 34-35
 - directives, 34
 - documentation, 34
 - keypass file, 34
 - mod_ssl web site, 34
 - specifications, 10
 - starting, 35
 - stopping, 35
 - verifying configuration changes, 35
- automatic reboot, disabling, 20

B

- boot-device environment variable, setting, 20
- boot devices, alternate boot path, 26

C

- cable connections
 - to the administrative console, 14
 - to the cluster nodes, 14-15
 - to the Sun StorEdge 3510 FC RAID array, 15-16
- cconsole command
 - See also* Cluster Control Panel software
 - starting, 19
- CCP, *See* Cluster Control Panel software
- ce:ce_taskq_disable environment variable, setting, 24
- cldevicegroup command, taking ownership of disk sets, 31
- Cluster Control Panel software, installing, 18
- cluster file systems
 - creating, 32-33
 - mounting, 33
 - verifying mounts, 33
- cluster nodes
 - adding
 - to MANPATH, 22
 - to PATH, 22
 - applying patches, 21
 - configuring, 28-29
 - connecting
 - to the management network, 15
 - to the public network, 15
 - to the storage, 15
 - specifications, 10

configuring

See also creating

Apache software, 34-35

Solaris software, 22-25

storage, 16-17

Sun Cluster HA for Apache, 39-40

Sun Cluster HA for NFS, 40

Sun Cluster HA for Oracle, 41-42

Sun Cluster software, 28-29

creating

See also configuring

Apache keypass file, 34

disk aliases, 20

disk sets, 29-30

/etc/lvm/md.tab file, 31-32

file systems, 32-33

Oracle

database, 37-38

dba system group, 27

home directory, 28

oinstall system group, 27

oracle user account, 28

password, 28

system groups, 27

v\$sysstat view, 38

state database replicas, 25

verifying cluster creation, 29

cssd daemon, disabling, 37

D**data service software****configuring**

Sun Cluster HA for Apache, 39-40

Sun Cluster HA for NFS, 40

Sun Cluster HA for Oracle, 41-42

installing, 26-27

dba system group, creating, 27

DID mappings, listing, 30

disabling

automatic reboot, 20

Oracle cssd daemon, 37

disk aliases

creating, 20

setting the boot-device environment variable, 20

disk sets

adding LUNs, 30-31

creating, 29-30

taking ownership, 31

verifying, 29

E

enabling, Solaris multipathing, 22

environment variables

boot-device, 20

ce:ce_taskq_disable, 24

exclude:lofs, 24

ORACLE_BASE, 22

ORACLE_SID, 22

error messages, cluster, 8

/etc/apache/httpd.conf file, modifying, 34

/etc/inet/ipnodes file, adding public-network IP addresses, 23

/etc/lvm/md.tab file, creating, 31-32

/etc/system file, setting environment variables, 24

/etc/vfstab file, adding mount points, 33

exclude:lofs, setting, 24

F**file systems**

creating, 32-33

mounting, 33

verifying mounts, 33

H

help, 8

httpd.conf file, modifying, 34

I**installer program**

log files, 27

starting, 26

installing
 Oracle software, 35-37
 patches
 administrative console, 18
 cluster nodes, 21
 Solaris software
 administrative console, 18
 cluster nodes, 19-21
 Sun Cluster software
 Cluster Control Panel, 18
 data services, 26-27
 framework, 26-27
 IP addresses
 adding
 to the ipnodes file, 23
 to the naming service, 19
 example
 management-network addresses, 11
 public-network addresses, 11-12
 ipnodes file, adding public-network IP addresses, 23

J

Java ES installer program
 log files, 27
 starting, 26

K

kernel parameters, Oracle settings, 23
 keypass file, creating, 34

L

listing, DID mappings, 30
 LOFS, *See* loopback file system (LOFS)
 log files
 Java ES installer program, 27
 scinstall utility, 29
 logical addresses, adding to the ipnodes file, 23
 logical drives
 mapping to LUNs, 17
 partitioning, 17

logical unit numbers (LUNs)
 adding to disk sets, 30-31
 mapping to logical drives, 17
 loopback file system (LOFS), disabling, 24
 LUNs, *See* logical unit numbers (LUNs)

M

management network
 connecting
 to the administrative console, 14
 to the cluster nodes, 15
 to the storage, 15
 example addresses, 11
 MANPATH
 adding to the administrative console, 18
 adding to the cluster nodes, 22
 md.tab file, creating, 31-32
 messages files, cluster, 8
 modifying, Apache httpd.conf file, 34
 mount points
 adding to the /etc/vfstab file, 33
 verifying, 33
 mounting file systems, 33

N

naming service, adding public-network IP addresses, 19
 NFS software
 setting the version, 24
 specifications, 10

O

oinstall system group, creating, 27
 ORACLE_BASE, setting, 22
 ORACLE_SID, setting, 22
 Oracle software
 creating
 database, 37-38
 dba system group, 27
 home directory, 28
 oinstall system group, 27

Oracle software, creating (*Continued*)

- oracle user account, 28
- password, 28
- v\$sysstat view, 38

disabling the cssd daemon, 37

installing, 35-37

setting

- database permissions, 38
- directory permissions, 32
- kernel parameters, 23

specifications, 11

P

partitioning

- logical drives, 17
- root disk, 21

patches

- administrative console, 18
- cluster nodes, 21

PATH

- adding to the administrative console, 18
- adding to the cluster nodes, 22

permissions

- Oracle database, 38
- Oracle directories, 32

public network

adding IP addresses

- ipnodes file, 23
- naming service, 19

connecting

- to the administrative console, 14
- to the cluster nodes, 15

example addresses, 11-12

Q

quorum devices, verifying, 29

R

replicas, creating, 25

root (/) file systems, mirroring, 25-26

root disk, partitions, 21

S

scinstall utility

- log files, 29
- starting, 28

servers, *See* cluster nodes

setting

- Apache directives, 34
- NFS version, 24
- ORACLE_BASE environment variable, 22
- Oracle kernel parameters, 23
- ORACLE_SID environment variable, 22
- umask, 22

Solaris multipathing, enabling, 22

Solaris software

- configuring, 22-25
- installing, 19-21
- partitions, 21

Solaris Volume Manager

disk sets

- adding LUNs, 30-31
- creating, 29-30
- taking ownership, 31
- verifying, 29
- mirroring root (/) file systems, 25-26

specifications

hardware, 10

- administrative console, 10
- Sun Fire V440 servers, 10
- Sun StorEdge 3510 FC RAID array, 10

software, 10-11

- Apache software, 10
- NFS software, 10
- Oracle software, 11

starting

- Apache server, 35
- cconsole software, 19
- Java ES installer program, 26
- scinstall utility, 28
- Sun Cluster Manager, 39

state database replicas, creating, 25

stmsboot command, 22

stopping, Apache server, 35

- storage
 - configuring, 16-17
 - connecting
 - to the cluster nodes, 15
 - to the management network, 15
 - specifications, 10
- Sun Cluster HA for Apache, configuring, 39-40
- Sun Cluster HA for NFS, configuring, 40
- Sun Cluster HA for Oracle, configuring, 41-42
- Sun Cluster Manager, starting, 39
- Sun Cluster software
 - configuring, 28-29
 - installing
 - Cluster Control Panel software, 18
 - data service software, 26-27
 - framework software, 26-27
- Sun Fire V440 servers, specifications, 10
- Sun StorEdge 3510 FC RAID array
 - connecting, 15-16
 - specifications, 10
- Sun Ultra 20 workstation, specifications, 10
- Sun Update Connection software, 21
- `/var/adm/messages` file, 8
- `/var/cluster/logs/install/scinstall.log.N` file, 29
- `/var/sadm/install/logs/` directory, 27
- verifying
 - Apache configuration changes, 35
 - disk sets, 29
 - file system mounts, 33
 - mount points, 33
 - quorum devices, 29
- `vfstab` file, adding mount points, 33

T

- technical support, 8

U

- `umask`, setting, 22
- Update Connection software, 21
- `/usr/cluster/bin/cconsole` command
 - See also* Cluster Control Panel software
 - starting, 19
- `/usr/cluster/bin/cldevicegroup` command, taking
 - ownership of disk sets, 31
- `/usr/cluster/bin/scinstall` utility
 - log files, 29
 - starting, 28

V

- `v$sysstat` view, creating, 38

