



Sun Cluster Software Installation Guide for Solaris OS

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

Part No: 817-6543-10
September 2004, Revision A

Copyright 2004 Sun Microsystems, Inc. 4150 Network Circle, Santa Clara, CA 95054 U.S.A. All rights reserved.

This product or document is protected by copyright and distributed under licenses restricting its use, copying, distribution, and decompilation. No part of this product or document may be reproduced in any form by any means without prior written authorization of Sun and its licensors, if any. Third-party software, including font technology, is copyrighted and licensed from Sun suppliers.

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, docs.sun.com, AnswerBook, AnswerBook2, Java, JumpStart, Solstice DiskSuite, Sun Fire, SunPlex, Sun StorEdge, 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. Netscape is a trademark or registered trademark of Netscape Communications Corporation in the United States and other countries. Netscape Navigator is a trademark or registered trademark of Netscape Communications Corporation in the United States and other countries. 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.

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.

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 2004 Sun Microsystems, Inc. 4150 Network Circle, Santa Clara, CA 95054 U.S.A. Tous droits réservés.

Ce produit ou document est protégé par un copyright et distribué avec des licences qui en restreignent l'utilisation, la copie, la distribution, et la décompilation. Aucune partie de ce produit ou document ne peut être reproduite sous aucune forme, par quelque moyen que ce soit, sans l'autorisation préalable et écrite de Sun et de ses bailleurs de licence, s'il y en a. Le logiciel détenu par des tiers, et qui comprend la technologie relative aux polices de caractères, est protégé par un copyright et licencié par des fournisseurs de Sun.

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

Sun, Sun Microsystems, le logo Sun, docs.sun.com, AnswerBook, AnswerBook2, Java, JumpStart, Solstice DiskSuite, Sun Fire, SunPlex, Sun StorEdge, 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. Netscape est une marque de Netscape Communications Corporation aux Etats-Unis et dans d'autres pays. Netscape Navigator est une marque de Netscape Communications Corporation aux Etats-Unis et dans d'autres pays. 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.

CETTE PUBLICATION EST FOURNIE "EN L'ETAT" ET AUCUNE GARANTIE, EXPRESSE OU IMPLICITE, N'EST ACCORDEE, Y COMPRIS DES GARANTIES CONCERNANT LA VALEUR MARCHANDE, L'APTITUDE DE LA PUBLICATION A REpondre A UNE UTILISATION PARTICULIERE, OU LE FAIT QU'ELLE NE SOIT PAS CONTREFAISANTE DE PRODUIT DE TIERS. CE DENI DE GARANTIE NE S'APPLIQUERAIT PAS, DANS LA MESURE OU IL SERAIT TENU JURIDIQUEMENT NUL ET NON AVENU.



040812@9495



Contents

Preface 9

1	Planning the Sun Cluster Configuration	15
	Where to Find Sun Cluster Installation Tasks	15
	Planning the Solaris OS	16
	Guidelines for Selecting Your Solaris Installation Method	17
	Solaris OS Feature Restrictions	17
	Solaris Software Group Considerations	17
	System Disk Partitions	18
	Planning the Sun Cluster Environment	21
	Licensing	21
	Software Patches	22
	IP Addresses	22
	Console-Access Devices	23
	Logical Addresses	23
	Public Networks	23
	Guidelines for NFS	24
	Service Restrictions	24
	Sun Cluster Configurable Components	25
	Planning the Global Devices and Cluster File Systems	30
	Guidelines for Highly Available Global Devices and Cluster File Systems	31
	Cluster File Systems	31
	Disk Device Groups	31
	Mount Information for Cluster File Systems	32
	Planning Volume Management	33
	Guidelines for Volume-Manager Software	34

Guidelines for Solstice DiskSuite or Solaris Volume Manager Software	35
SPARC: Guidelines for VERITAS Volume Manager Software	37
File-System Logging	38
Mirroring Guidelines	39
2 Installing and Configuring Sun Cluster Software	43
Task Map: Installing the Software	44
Installing the Software	46
▼ How to Prepare for Cluster Software Installation	46
▼ How to Install Cluster Control Panel Software on an Administrative Console	48
▼ How to Install Solaris Software	52
▼ How to Install Sun Cluster Software Packages	56
▼ How to Configure Sun Cluster Software on All Nodes (<code>scinstall</code>)	57
Using SunPlex Installer to Install Sun Cluster Software	63
▼ How to Install SunPlex Manager Software	65
▼ How to Install and Configure Sun Cluster Software (SunPlex Installer)	70
▼ How to Install Solaris and Sun Cluster Software (JumpStart)	76
How to Install Sun Cluster Software on a Single-Node Cluster	87
▼ How to Configure Sun Cluster Software on Additional Cluster Nodes (<code>scinstall</code>)	90
▼ SPARC: How to Install VERITAS File System Software	97
▼ How to Configure the Name-Service Switch	98
▼ How to Set Up the Root Environment	99
▼ How to Install Data-Service Software Packages (<code>installer</code>)	100
▼ How to Install Data-Service Software Packages (<code>scinstall</code>)	102
▼ How to Perform Postinstallation Setup and Configure Quorum Devices	104
▼ How to Verify the Quorum Configuration and Installation Mode	107
Task Map: Configuring the Cluster	107
Configuring the Cluster	108
▼ How to Create Cluster File Systems	108
▼ How to Configure Internet Protocol (IP) Network Multipathing Groups	113
▼ How to Change Private Hostnames	115
▼ How to Configure Network Time Protocol (NTP)	116
SPARC: Task Map: Installing the Sun Cluster Module for Sun Management Center	117
SPARC: Installing the Sun Cluster Module for Sun Management Center	118
SPARC: Installation Requirements for Sun Cluster Monitoring	118

▼ SPARC: How to Install the Sun Cluster Module for Sun Management Center	119
▼ SPARC: How to Start Sun Management Center	120
▼ SPARC: How to Add a Cluster Node as a Sun Management Center Agent Host Object	121
▼ SPARC: How to Load the Sun Cluster Module	121
Uninstalling the Software	123
▼ How to Uninstall Sun Cluster Software to Correct Installation Problems	123
▼ How to Uninstall the SUNWscrdt Package	124
▼ How to Unload the RSMRDT Driver Manually	125
3 Installing and Configuring Solstice DiskSuite or Solaris Volume Manager Software	127
Task Map: Installing and Configuring Solstice DiskSuite or Solaris Volume Manager Software	128
Installing and Configuring Solstice DiskSuite or Solaris Volume Manager Software	129
Solstice DiskSuite or Solaris Volume Manager Configuration Example	130
▼ How to Install Solstice DiskSuite Software	131
▼ How to Set the Number of Metadevice or Volume Names and Disk Sets	133
▼ How to Create State Database Replicas	135
Mirroring the Root Disk	136
▼ How to Mirror the Root (/) File System	136
▼ How to Mirror the Global Namespace	140
▼ How to Mirror File Systems Other Than Root (/) That Cannot Be Unmounted	143
▼ How to Mirror File Systems That Can Be Unmounted	147
Creating Disk Sets in a Cluster	150
▼ How to Create a Disk Set	151
Adding Drives to a Disk Set	153
▼ How to Repartition Drives in a Disk Set	155
▼ How to Create an md.tab File	156
▼ How to Activate Metadevices or Volumes	157
Configuring Dual-String Mediators	159
Requirements for Dual-String Mediators	159
▼ How to Add Mediator Hosts	160
▼ How to Check the Status of Mediator Data	160
▼ How to Fix Bad Mediator Data	161

4	SPARC: Installing and Configuring VERITAS Volume Manager	163
	SPARC: Installing and Configuring VxVM Software	163
	SPARC: Task Map: Installing and Configuring VxVM Software	164
	SPARC: Setting Up a Root Disk Group Overview	164
	▼ SPARC: How to Install VERITAS Volume Manager Software and Encapsulate the Root Disk	165
	▼ SPARC: How to Mirror the Encapsulated Root Disk	168
	▼ SPARC: How to Install VERITAS Volume Manager Software Only	170
	▼ SPARC: How to Create a Root Disk Group on a Nonroot Disk	172
	▼ SPARC: How to Create and Register a Disk Group	173
	▼ SPARC: How to Assign a New Minor Number to a Disk Device Group	175
	▼ SPARC: How to Verify the Disk Group Configuration	176
	▼ SPARC: How to Unencapsulate the Root Disk	177
5	Upgrading Sun Cluster Software	181
	Overview of Upgrading a Sun Cluster Configuration	181
	Upgrade Requirements and Support Guidelines	182
	Choosing a Sun Cluster Upgrade Method	182
	Upgrading to Sun Cluster 3.1 9/04 Software (Nonrolling)	183
	Task Map: Upgrading to Sun Cluster 3.1 9/04 Software (Nonrolling)	184
	▼ How to Prepare the Cluster for a Nonrolling Upgrade	184
	▼ How to Perform a Nonrolling Upgrade of the Solaris OS	189
	▼ How to Perform a Nonrolling Upgrade of Sun Cluster 3.1 9/04 Software	192
	▼ How to Finish a Nonrolling Upgrade to Sun Cluster 3.1 9/04 Software	198
	Upgrading to Sun Cluster 3.1 9/04 Software (Rolling)	201
	Task Map: Upgrading to Sun Cluster 3.1 9/04 Software (Rolling)	201
	▼ How to Prepare a Cluster Node for a Rolling Upgrade	202
	▼ How to Perform a Rolling Upgrade of a Solaris Maintenance Update	206
	▼ How to Perform a Rolling Upgrade of Sun Cluster 3.1 9/04 Software	207
	▼ How to Finish a Rolling Upgrade to Sun Cluster 3.1 9/04 Software	212
	Recovering From Storage Configuration Changes During Upgrade	214
	▼ How to Handle Storage Reconfiguration During an Upgrade	215
	▼ How to Resolve Mistaken Storage Changes During an Upgrade	216
	SPARC: Sun Management Center Software Upgrade	217
	▼ SPARC: How to Upgrade Sun Cluster-Module Software for Sun Management Center	217
	▼ SPARC: How to Upgrade Sun Management Center Software	218

6	Configuring Data Replication With Sun StorEdge Availability Suite 3.1 Software	221
	Introduction to Data Replication	222
	What Is Disaster Tolerance?	222
	Data Replication Methods Used by Sun StorEdge Availability Suite Software	222
	Guidelines for Configuring Data Replication	225
	Configuring Replication Resource Groups	226
	Configuring Application Resource Groups	226
	Guidelines for Managing a Failover or Switchover	229
	Example Configuration	230
	Connecting and Installing the Clusters	230
	Example of How to Configure Device Groups and Resource Groups	232
	Example of How to Enable Data Replication	241
	Example of How to Perform Data Replication	243
	Example of How to Verify That Replication Is Configured Correctly	245
	Example of How to Cope With a Failover or Switchover	247
A	Sun Cluster Installation and Configuration Worksheets	251
	Installation and Configuration Worksheets	252
	Local File System Layout Worksheet	254
	Public Networks Worksheet	256
	Local Devices Worksheets	258
	Disk Device Group Configurations Worksheet	260
	Volume-Manager Configurations Worksheet	262
	Metadevices Worksheet (Solstice DiskSuite or Solaris Volume Manager)	264
	Index	267

Preface

The *Sun Cluster Software Installation Guide for Solaris OS* contains guidelines for planning a Sun™ Cluster configuration, and provides procedures for installing, configuring, and upgrading the Sun Cluster software on both SPARC® based systems and x86 based systems. This book also provides a detailed example of how to use Sun StorEdge™ Availability Suite 3.1 software to configure data replication between clusters.

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

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.

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

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 on 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 changes that are used in this book.

TABLE P-1 Typographic Conventions

Typeface or Symbol	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>	Command-line placeholder: replace with a real name or value	The command to remove a file is <code>rm filename</code> .
<i>AaBbCc123</i>	Book titles, new terms, and terms to be emphasized	Read Chapter 6 in the <i>User's Guide</i> . These are called <i>class</i> options. Do <i>not</i> save the file. (Emphasis sometimes appears in 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-2 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	#

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.x 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>

Topic	Documentation
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>.

Accessing Sun Documentation Online

The docs.sun.comSM 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>.

Getting Help

If you have problems installing or using Sun Cluster, 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 8)
- The release number of Sun Cluster (for example, Sun Cluster 3.1 9/04)

Use the following commands to gather information on 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
SPARC: <code>prtdiag -v</code>	Displays system diagnostic information
<code>/usr/cluster/bin/scinstall -pv</code>	Displays Sun Cluster release and package version information

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

Planning the Sun Cluster Configuration

This chapter provides planning information and guidelines for installing a Sun Cluster configuration.

The following overview information is in this chapter:

- “Where to Find Sun Cluster Installation Tasks” on page 15
- “Planning the Solaris OS” on page 16
- “Planning the Sun Cluster Environment” on page 21
- “Planning the Global Devices and Cluster File Systems” on page 30
- “Planning Volume Management” on page 33

Where to Find Sun Cluster Installation Tasks

The following table shows where to find instructions for various installation tasks for Sun Cluster software installation and the order in which you should perform the tasks.

TABLE 1-1 Sun Cluster Software Installation Task Information

Task	Instructions
Set up cluster hardware.	<ul style="list-style-type: none">■ <i>Sun Cluster 3.x Hardware Administration Manual</i>■ Documentation that shipped with your server and storage devices
Plan cluster software installation.	<ul style="list-style-type: none">■ Chapter 1■ “Installation and Configuration Worksheets” on page 252

TABLE 1-1 Sun Cluster Software Installation Task Information (Continued)

Task	Instructions
Install a new cluster or add nodes to an existing cluster. Optionally, install and configure Sun StorEdge QFS software.	<p>“Installing the Software” on page 46</p> <p><i>Sun StorEdge QFS and Sun StorEdge SAM-FS Software Installation and Configuration Guide</i></p>
Install and configure Solstice DiskSuite™ or Solaris Volume Manager software.	<ul style="list-style-type: none"> ■ “Installing and Configuring Solstice DiskSuite or Solaris Volume Manager Software” on page 129 ■ Solstice DiskSuite or Solaris Volume Manager documentation
SPARC: Install and configure VERITAS Volume Manager (VxVM) software.	<ul style="list-style-type: none"> ■ “SPARC: Installing and Configuring VxVM Software” on page 163 ■ VxVM documentation
Configure cluster framework software and optionally install and configure the Sun Cluster module to Sun Management Center (which is available on SPARC based systems only).	<p>“Configuring the Cluster” on page 108</p>
Plan, install, and configure resource groups and data services.	<p><i>Sun Cluster Data Service Planning and Administration Guide for Solaris OS</i></p>
Develop custom data services.	<p><i>Sun Cluster Data Services Developer’s Guide for Solaris OS</i></p>
Upgrade to Sun Cluster 3.1 9/04 software.	<ul style="list-style-type: none"> ■ Chapter 5 ■ “Installing and Configuring Solstice DiskSuite or Solaris Volume Manager Software” on page 129 or “SPARC: Installing and Configuring VxVM Software” on page 163 ■ Volume manager documentation

Planning the Solaris OS

This section provides guidelines for planning Solaris software installation in a cluster configuration. For more information about Solaris software, see your Solaris installation documentation.

Guidelines for Selecting Your Solaris Installation Method

You can install Solaris software from a local CD-ROM or from a network installation server by using the JumpStart™ installation method. In addition, Sun Cluster software provides a custom method for installing both the Solaris OS and Sun Cluster software by using the JumpStart installation method. If you are installing several cluster nodes, consider a network installation.

See “How to Install Solaris and Sun Cluster Software (JumpStart)” on page 76 for details about the `scinstall` JumpStart installation method. See your Solaris installation documentation for details about standard Solaris installation methods.

Solaris OS Feature Restrictions

The following Solaris OS features are not supported in a Sun Cluster configuration:

- Solaris interface groups are not supported in a Sun Cluster configuration. The Solaris interface groups feature is disabled by default during Solaris software installation. Do not re-enable Solaris interface groups. See the `ifconfig(1M)` man page for more information about Solaris interface groups.
- Automatic power-saving shutdown is not supported in Sun Cluster configurations and should not be enabled. See the `pmconfig(1M)` and `power.conf(4)` man pages for more information.

Solaris Software Group Considerations

Sun Cluster 3.1 9/04 software requires at least the End User Solaris Software Group. However, other components of your cluster configuration might have their own Solaris software requirements as well. Consider the following information when you decide which Solaris software group you are installing.

- Check your server documentation for any Solaris software requirements. For example, Sun Enterprise 10000 servers require the Entire Solaris Software Group Plus OEM Support.
- If you intend to use SCI-PCI adapters, which are available for use in SPARC based clusters only, or the Remote Shared Memory Application Programming Interface (RSMAPI), ensure that you install the RSMAPI software packages (`SUNWrsm`, `SUNWrsmx`, `SUNWrsmo`, and `SUNWrsmoz`). The RSMAPI software packages are included in only some Solaris software groups. For example, the Developer Solaris Software Group includes the RSMAPI software packages but the End User Solaris Software Group does not.

If the software group that you install does not include the RSMAPI software packages, install the RSMAPI software packages manually before you install Sun Cluster software. Use the `pkgadd(1M)` command to manually install the software

packages. See the Solaris 8 Section (3RSM) man pages for information about using the RSM API.

- You might need to install other Solaris software packages that are not part of the End User Solaris Software Group. The Apache HTTP server packages are one example. Third-party software, such as ORACLE®, might also require additional Solaris software packages. See your third-party documentation for any Solaris software requirements.

Tip – To avoid the need to manually install Solaris software packages, install the Entire Solaris Software Group Plus OEM Support.

System Disk Partitions

Add this information to the appropriate “Local File System Layout Worksheet” on page 254.

When you install the Solaris OS, ensure that you create the required Sun Cluster partitions and that all partitions meet minimum space requirements.

- **swap** – The combined amount of swap space that is allocated for Solaris and Sun Cluster software must be no less than 750 Mbytes. For best results, add at least 512 Mbytes for Sun Cluster software to the amount that is required by the Solaris OS. In addition, allocate any additional swap amount that is required by applications that are to run on the cluster node.

Note – If you intend to create an additional swap file, do not create the swap file on a global device. Only use a local disk as a swap device for the node.

- `/globaldevices` – Create a 512-Mbyte file system that is to be used by the `scinstall(1M)` utility for global devices.
- **Volume manager** – Create a 20-Mbyte partition on a slice at the end of the disk (slice 7) for volume manager use. If your cluster uses VERITAS Volume Manager (VxVM) and you intend to encapsulate the root disk, you need to have two unused slices available for use by VxVM.

To meet these requirements, you must customize the partitioning if you are performing interactive installation of the Solaris OS.

See the following guidelines for additional partition planning information:

- “Guidelines for the Root (/) File System” on page 19
- “Guidelines for the `/globaldevices` File System” on page 19
- “Volume Manager Requirements” on page 20

Guidelines for the Root (/) File System

As with any other system running the Solaris OS, you can configure the root (/), /var, /usr, and /opt directories as separate file systems. Or, you can include all the directories in the root (/) file system. The following describes the software contents of the root (/), /var, /usr, and /opt directories in a Sun Cluster configuration. Consider this information when you plan your partitioning scheme.

- **root (/)** – The Sun Cluster software itself occupies less than 40 Mbytes of space in the root (/) file system. Solstice DiskSuite or Solaris Volume Manager software requires less than 5 Mbytes, and VxVM software requires less than 15 Mbytes. To configure ample additional space and inode capacity, add at least 100 Mbytes to the amount of space you would normally allocate for your root (/) file system. This space is used for the creation of both block special devices and character special devices used by the volume management software. You especially need to allocate this extra space if a large number of shared disks are in the cluster.
- **/var** – The Sun Cluster software occupies a negligible amount of space in the /var file system at installation time. However, you need to set aside ample space for log files. Also, more messages might be logged on a clustered node than would be found on a typical standalone server. Therefore, allow at least 100 Mbytes for the /var file system.
- **/usr** – Sun Cluster software occupies less than 25 Mbytes of space in the /usr file system. Solstice DiskSuite or Solaris Volume Manager and VxVM software each require less than 15 Mbytes.
- **/opt** – Sun Cluster framework software uses less than 2 Mbytes in the /opt file system. However, each Sun Cluster data service might use between 1 Mbyte and 5 Mbytes. Solstice DiskSuite or Solaris Volume Manager software does not use any space in the /opt file system. VxVM software can use over 40 Mbytes if all of its packages and tools are installed.

In addition, most database and applications software is installed in the /opt file system.

SPARC: If you use Sun Management Center software to monitor the cluster, you need an additional 25 Mbytes of space on each node to support the Sun Management Center agent and Sun Cluster module packages.

Guidelines for the /globaldevices File System

Sun Cluster software requires you to set aside a special file system on one of the local disks for use in managing global devices. This file system is later mounted as a cluster file system. Name this file system /globaldevices, which is the default name that is recognized by the `scinstall(1M)` command.

The `scinstall` command later renames the file system /global/.devices/node@nodeid, where *nodeid* represents the number that is assigned to a node when it becomes a cluster member. The original /globaldevices mount point is removed.

The `/globaldevices` file system must have ample space and ample inode capacity for creating both block special devices and character special devices. This guideline is especially important if a large number of disks are in the cluster. A file system size of 512 Mbytes should suffice for most cluster configurations.

Volume Manager Requirements

If you use Solstice DiskSuite or Solaris Volume Manager software, you must set aside a slice on the root disk for use in creating the state database replica. Specifically, set aside a slice for this purpose on each local disk. But, if you only have one local disk on a node, you might need to create three state database replicas in the same slice for Solstice DiskSuite or Solaris Volume Manager software to function properly. See your Solstice DiskSuite or Solaris Volume Manager documentation for more information.

SPARC: If you use VERITAS Volume Manager (VxVM) and you intend to encapsulate the root disk, you need to have two unused slices that are available for use by VxVM. Additionally, you need to have some additional unassigned free space at either the beginning or the end of the disk. See your VxVM documentation for more information about root disk encapsulation.

Example—Sample File-System Allocations

Table 1–2 shows a partitioning scheme for a cluster node that has less than 750 Mbytes of physical memory. This scheme is to be installed with the End User Solaris Software Group, Sun Cluster software, and the Sun Cluster HA for NFS data service. The last slice on the disk, slice 7, is allocated with a small amount of space for volume-manager use.

This layout allows for the use of either Solstice DiskSuite or Solaris Volume Manager software or VxVM software. If you use Solstice DiskSuite or Solaris Volume Manager software, you use slice 7 for the state database replica. If you use VxVM, you later free slice 7 by assigning the slice a zero length. This layout provides the necessary two free slices, 4 and 7, as well as provides for unused space at the end of the disk.

TABLE 1–2 Example File-System Allocation

Slice	Contents	Allocation (in Mbytes)	Description
0	/	6.75GB	Remaining free space on the disk after allocating space to slices 1 through 7. Used for the Solaris OS, Sun Cluster software, data-services software, volume-manager software, Sun Management Center agent and Sun Cluster module agent packages, root file systems, and database and application software.

TABLE 1-2 Example File-System Allocation (Continued)

Slice	Contents	Allocation (in Mbytes)	Description
1	swap	1GB	512 Mbytes for the Solaris OS. 512 Mbytes for Sun Cluster software.
2	overlap	8.43GB	The entire disk.
3	/globaldevices	512MB	The Sun Cluster software later assigns this slice a different mount point and mounts the slice as a cluster file system.
4	unused	-	Available as a free slice for encapsulating the root disk under VxVM.
5	unused	-	-
6	unused	-	-
7	volume manager	20MB	Used by Solstice DiskSuite or Solaris Volume Manager software for the state database replica, or used by VxVM for installation after you free the slice.

Planning the Sun Cluster Environment

This section provides guidelines for planning and preparing the following components for Sun Cluster software installation and configuration:

- “Licensing” on page 21
- “Software Patches” on page 22
- “IP Addresses” on page 22
- “Console-Access Devices” on page 23
- “Logical Addresses” on page 23
- “Public Networks” on page 23
- “Guidelines for NFS” on page 24
- “Service Restrictions” on page 24
- “Sun Cluster Configurable Components” on page 25

For detailed information about Sun Cluster components, see the *Sun Cluster Overview for Solaris OS* and the *Sun Cluster Concepts Guide for Solaris OS*.

Licensing

Ensure that you have available all necessary license certificates before you begin software installation. Sun Cluster software does not require a license certificate, but each node installed with Sun Cluster software must be covered under your Sun Cluster software license agreement.

For licensing requirements for volume-manager software and applications software, see the installation documentation for those products.

Software Patches

After installing each software product, you must also install any required patches.

- For information about current required patches, see “Patches and Required Firmware Levels” in *Sun Cluster Release Notes for Solaris OS* or consult your Sun service provider.
- For general guidelines and procedures for applying patches, see “Patching Sun Cluster Software and Firmware” in *Sun Cluster System Administration Guide for Solaris OS*.

IP Addresses

You must set up a number of IP addresses for various Sun Cluster components, depending on your cluster configuration. Each node in the cluster configuration must have at least one public-network connection to the same set of public subnets.

The following table lists the components that need IP addresses assigned. Add these IP addresses to any naming services that are used. Also add these IP addresses to the local `/etc/inet/hosts` file on each cluster node after you install Solaris software.

- For more information about IP addresses, see *System Administration Guide, Volume 3* (Solaris 8) or *System Administration Guide: IP Services* (Solaris 9).
- For more information about test IP addresses to support IP Network Multipathing, see *IP Network Multipathing Administration Guide*.

TABLE 1-3 Sun Cluster Components That Use IP Addresses

Component	Number of IP Addresses Needed
Administrative console	1 per subnet
IP Network Multipathing groups	<ul style="list-style-type: none"> ■ Single-adapter groups – 1 ■ Multiple-adapter groups – 1 primary IP address plus 1 test IP address for each adapter in the group
Cluster nodes	1 per node, per subnet
Domain console network interface (Sun Fire™ 15000)	1 per domain
Console-access device	1

TABLE 1-3 Sun Cluster Components That Use IP Addresses (Continued)

Component	Number of IP Addresses Needed
Logical addresses	1 per logical host resource, per subnet

Console-Access Devices

You must have console access to all cluster nodes. If you install Cluster Control Panel software on your administrative console, you must provide the hostname of the console-access device that is used to communicate with the cluster nodes.

- A terminal concentrator is used to communicate between the administrative console and the cluster node consoles.
- A Sun Enterprise 10000 server uses a System Service Processor (SSP) instead of a terminal concentrator.
- A Sun Fire™ server uses a system controller instead of a terminal concentrator.

For more information about console access, see the *Sun Cluster Concepts Guide for Solaris OS*.

Logical Addresses

Each data-service resource group that uses a logical address must have a hostname specified for each public network from which the logical address can be accessed.

- For more information, see the *Sun Cluster Data Service Planning and Administration Guide for Solaris OS*.
- For additional information about data services and resources, also see the *Sun Cluster Overview for Solaris OS* and the *Sun Cluster Concepts Guide for Solaris OS*.

Public Networks

Public networks communicate outside the cluster. Consider the following points when you plan your public-network configuration:

- Public networks and the private network (cluster interconnect) must use separate adapters.
- You must have at least one public network that is connected to all cluster nodes.
- You can have as many additional public-network connections as your hardware configuration allows.
- Sun Cluster software supports IPv4 and IPv6 addresses on the public network, both for failover and scalable data services. However, Sun Cluster software does not support IPv6 communications over the private interconnects.

- The `local-mac-address?` variable must use the default value `true` for Ethernet adapters. Sun Cluster software does not support a `local-mac-address?` value of `false` for Ethernet adapters. This requirement is a change from Sun Cluster 3.0, which did require a `local-mac-address?` value of `false`.
- During Sun Cluster installation, the `scinstall` utility automatically configures a single-adapter IP Network Multipathing group for each public-network adapter. To modify these backup groups after installation, follow the procedures in “Deploying Network Multipathing” in *IP Network Multipathing Administration Guide* (Solaris 8) or “Administering Network Multipathing (Task)” in *System Administration Guide: IP Services* (Solaris 9).

See “IP Network Multipathing Groups” on page 29 for guidelines on planning public-network-adapter backup groups. For more information about public-network interfaces, see *Sun Cluster Concepts Guide for Solaris OS*.

Guidelines for NFS

Consider the following points when you plan the use of Network File System (NFS) in a Sun Cluster configuration.

- No Sun Cluster node can be an NFS client of a Sun Cluster HA for NFS-exported file system being mastered on a node in the same cluster. Such cross-mounting of Sun Cluster HA for NFS is prohibited. Use the cluster file system to share files among cluster nodes.
- Applications that run locally on the cluster must not lock files on a file system exported through NFS. Otherwise, local blocking (for example, `fcntl(3UCB)` or `fcntl(2)`) might interfere with the ability to restart the lock manager (`lockd`). During restart, a blocked local process might be granted a lock which might be intended to be reclaimed by a remote client. This would cause unpredictable behavior.
- Sun Cluster software does not support Secure NFS or the use of Kerberos with NFS. In particular, Sun Cluster software does not support the `secure` and `kerberos` options to the `share_nfs(1M)` subsystem.

However, Sun Cluster software does support the use of secure ports for NFS. You enable secure ports for NFS by adding the entry `set nfssrv:nfs_portmon=1` to the `/etc/system` file on cluster nodes.

Service Restrictions

Observe the following service restrictions for Sun Cluster configurations:

- Do not configure cluster nodes as routers (gateways). If the system goes down, the clients cannot find an alternate router and cannot recover.
- Do not configure cluster nodes as NIS or NIS+ servers. There is no data service available for NIS or NIS+. However, cluster nodes can be NIS or NIS+ clients.

- Do not use a Sun Cluster configuration to provide a highly available boot or installation service on client systems.
- Do not use a Sun Cluster configuration to provide an `rarpd` service.
- If you install an RPC service on the cluster, the service must not use any of the following program numbers:
 - 100141
 - 100142
 - 100248

These numbers are reserved for the Sun Cluster daemons `rgmd_receptionist`, `fed`, and `pmfd`, respectively.

If the RPC service that you install also uses one of these program numbers, you must change that RPC service to use a different program number.

- Sun Cluster software does not support the running of high-priority process scheduling classes on cluster nodes. Do not run either of the following types of processes on cluster nodes:
 - Processes that run in the time-sharing scheduling class with a high priority
 - Processes that run in the real-time scheduling class

Sun Cluster software relies on kernel threads that do not run in the real-time scheduling class. Other time-sharing processes that run at higher-than-normal priority or real-time processes can prevent the Sun Cluster kernel threads from acquiring needed CPU cycles.

Sun Cluster Configurable Components

This section provides guidelines for the following Sun Cluster components that you configure:

- “Cluster Name” on page 26
- “Node Names” on page 26
- “Private Network” on page 26
- “Private Hostnames” on page 27
- “Cluster Interconnect” on page 28
- “IP Network Multipathing Groups” on page 29
- “Quorum Devices” on page 30

Add this information to the appropriate configuration worksheet.

TABLE 1-4 Worksheets for Sun Cluster Configuration

Configuration Worksheet	Location
Table 2-2 (to use defaults) or Table 2-3 (to customize)	"How to Configure Sun Cluster Software on All Nodes (scinstall)" on page 57
Table 2-6	"How to Install and Configure Sun Cluster Software (SunPlex Installer)" on page 70
Table 2-7	"How to Install Solaris and Sun Cluster Software (JumpStart)" on page 76
Table 2-8	"How to Configure Sun Cluster Software on Additional Cluster Nodes (scinstall)" on page 90

Cluster Name

Specify a name for the cluster during Sun Cluster configuration. The cluster name should be unique throughout the enterprise.

Node Names

The node name is the name that you assign to a machine when you install the Solaris OS. During Sun Cluster configuration, you specify the names of all nodes that you are installing as a cluster. In single-node cluster installations, the default node name is the same as the cluster name.

Private Network

Note – You do not need to configure a private network for a single-node cluster.

Sun Cluster software uses the private network for internal communication between nodes. A Sun Cluster configuration requires at least two connections to the cluster interconnect on the private network. You specify the private-network address and netmask when you configure Sun Cluster software on the first node of the cluster. You can either accept the default private-network address (172.16.0.0) and netmask (255.255.0.0) or type different choices if the default network address is already in use elsewhere in the same enterprise.

Note – After the installation utility (`scinstall`, SunPlex Installer, or JumpStart) has finished processing and the cluster is established, you cannot change the private-network address and netmask. You must uninstall and reinstall the cluster software to use a different private-network address or netmask.

If you specify a private-network address other than the default, the address must meet the following requirements:

- The address must use zeroes for the last two octets of the address, as in the default address 172.16.0.0. Sun Cluster software requires the last 16 bits of the address space for its own use.
- The address must be included in the block of addresses that RFC 1918 reserves for use in private networks. You can contact the InterNIC to obtain copies of RFCs or view RFCs online at <http://www.rfcs.org>.
- You can use the same private network address in more than one cluster. Private IP network addresses are not accessible from outside the cluster.
- Sun Cluster software does not support IPv6 addresses for the private interconnect.

Although the `scinstall` utility lets you specify an alternate netmask, best practice is to accept the default netmask, 255.255.0.0. There is no benefit if you specify a netmask that represents a larger network. And the `scinstall` utility does not accept a netmask that represents a smaller network.

See “Planning Your TCP/IP Network” in *System Administration Guide, Volume 3* (Solaris 8) or “Planning Your TCP/IP Network (Task)” in *System Administration Guide: IP Services* (Solaris 9) for more information about private networks.

Private Hostnames

The private hostname is the name that is used for internode communication over the private-network interface. Private hostnames are automatically created during Sun Cluster configuration. These private hostnames follow the naming convention `clusternodenodeid-priv`, where `nodeid` is the numeral of the internal node ID. During Sun Cluster configuration, the node ID number is automatically assigned to each node when the node becomes a cluster member. After the cluster is configured, you can rename private hostnames by using the `scsetup(1M)` utility.

Cluster Interconnect

Note – You do not need to configure a cluster interconnect for a single-node cluster. However, if you anticipate eventually adding nodes to a single-node cluster configuration, you might want to configure the cluster interconnect for future use.

The cluster interconnects provide the hardware pathways for private-network communication between cluster nodes. Each interconnect consists of a cable that is connected in one of the following ways:

- Between two transport adapters
- Between a transport adapter and a transport junction
- Between two transport junctions

During Sun Cluster configuration, you specify the following information for two cluster interconnects:

- **Transport adapters** – For the transport adapters, such as ports on network interfaces, specify the transport adapter names and transport type. If your configuration is a two-node cluster, you also specify whether your interconnect is direct connected (adapter to adapter) or uses a transport junction. If your two-node cluster is direct connected, you can still specify a transport junction for the interconnect.

Tip – If you specify a transport junction, you can more easily add another node to the cluster in the future.

Consider the following guidelines and restrictions:

- **SBus SCI adapters** – The SBus Scalable Coherent Interface (SCI) is not supported as a cluster interconnect. However, the SCI-PCI interface is supported.
- **Logical network interfaces** – Logical network interfaces are reserved for use by Sun Cluster software. See the `sconf_trans_adap_*(1M)` family of man pages for information about a specific transport adapter.
- **Transport junctions** – If you use transport junctions, such as a network switch, specify a transport junction name for each interconnect. You can use the default name `switchN`, where `N` is a number that is automatically assigned during configuration, or create another name. The exception is the Sun Firelink adapter, which requires the junction name `sw-rsmN`. The `scinstall` utility automatically uses this junction name after you specify a Sun Firelink adapter (`wrsmN`).

Also specify the junction port name or accept the default name. The default port name is the same as the internal node ID number of the node that hosts the adapter end of the cable. However, you cannot use the default port name for certain adapter types, such as SCI-PCI.

Note – Clusters with three or more nodes *must* use transport junctions. Direct connection between cluster nodes is supported only for two-node clusters.

You can configure additional private-network connections after the cluster is established by using the `scsetup(1M)` utility.

For more information about the cluster interconnect, see “Cluster Interconnect” in *Sun Cluster Overview for Solaris OS* and *Sun Cluster Concepts Guide for Solaris OS*.

IP Network Multipathing Groups

Add this planning information to the “[Public Networks Worksheet](#)” on page 256.

Internet Protocol (IP) Network Multipathing groups, which replace Network Adapter Failover (NAFO) groups, provide public-network adapter monitoring and failover, and are the foundation for a network-address resource. A multipathing group provides high availability when the multipathing group is configured with two or more adapters. If one adapter fails, all of the addresses on the failed adapter fail over to another adapter in the multipathing group. In this way, the multipathing-group adapters maintain public-network connectivity to the subnet to which the adapters in the multipathing group connect.

Consider the following points when you plan your multipathing groups.

- Each public network adapter must belong to a multipathing group.
- For multipathing groups that contain two or more adapters, you must configure a test IP address for each adapter in the group. If a multipathing group contains only one adapter, you do not need to configure a test IP address.
- Test IP addresses for all adapters in the same multipathing group must belong to a single IP subnet.
- Test IP addresses must not be used by normal applications because the test IP addresses are not highly available.
- In the `/etc/default/mpathd` file, the value of `TRACK_INTERFACES_ONLY_WITH_GROUPS` must be `yes`.
- The name of a multipathing group has no requirements or restrictions.

Most procedures, guidelines, and restrictions that are identified in the Solaris documentation for IP Network Multipathing are the same for both cluster and noncluster environments. Therefore, see the appropriate Solaris document for additional information about IP Network Multipathing:

- For the Solaris 8 OS, see “Deploying Network Multipathing” in *IP Network Multipathing Administration Guide*.
- For the Solaris 9 OS, see “Administering Network Multipathing (Task)” in *System Administration Guide: IP Services*.

Also see “IP Network Multipathing Groups” in *Sun Cluster Overview for Solaris OS* and *Sun Cluster Concepts Guide for Solaris OS*.

Quorum Devices

Sun Cluster configurations use quorum devices to maintain data and resource integrity. If the cluster temporarily loses connection to a node, the quorum device prevents amnesia or split-brain problems when the cluster node attempts to rejoin the cluster. You configure quorum devices by using the `scsetup(1M)` utility.

Note – You do not need to configure quorum devices for a single-node cluster.

Consider the following points when you plan quorum devices.

- **Minimum** – A two-node cluster must have at least one shared disk assigned as a quorum device. For other topologies, quorum devices are optional.
- **Odd-number rule** – If more than one quorum device is configured in a two-node cluster, or in a pair of nodes directly connected to the quorum device, configure an odd number of quorum devices. This configuration ensures that the quorum devices have completely independent failure pathways.
- **Connection** – You must connect a quorum device to at least two nodes.

For more information about quorum devices, see “Quorum and Quorum Devices” in *Sun Cluster Concepts Guide for Solaris OS* and “Quorum Devices” in *Sun Cluster Overview for Solaris OS*.

Planning the Global Devices and Cluster File Systems

This section provides the following guidelines for planning global devices and for planning cluster file systems:

- “Guidelines for Highly Available Global Devices and Cluster File Systems” on page 31
- “Disk Device Groups” on page 31
- “Mount Information for Cluster File Systems” on page 32

For more information about global devices and about cluster files systems, see *Sun Cluster Overview for Solaris OS* and *Sun Cluster Concepts Guide for Solaris OS*.

Guidelines for Highly Available Global Devices and Cluster File Systems

Sun Cluster software does not require any specific disk layout or file system size. Consider the following points when you plan your layout for global devices and for cluster file systems.

- **Mirroring** – You must mirror all global devices for the global device to be considered highly available. You do not need to use software mirroring if the storage device provides hardware RAID as well as redundant paths to disks.
- **Disks** – When you mirror, lay out file systems so that the file systems are mirrored across disk arrays.
- **Availability** – You must physically connect a global device to more than one node in the cluster for the global device to be considered highly available. A global device with multiple physical connections can tolerate a single-node failure. A global device with only one physical connection is supported, but the global device becomes inaccessible from other nodes if the node with the connection is down.
- **Swap devices** - Do not create a swap file on a global device.

Cluster File Systems

Consider the following points when you plan cluster file systems.

- **Quotas** - Quotas are not supported on cluster file systems.
- **Loopback file system (LOFS)** - Sun Cluster software does not support the use of the loopback file system (LOFS) on cluster nodes.
- **Communication end-points** - The cluster file system does not support any of the file-system features of Solaris software by which one would put a communication end-point in the file-system name space.
 - Although you can create a UNIX domain socket whose name is a path name into the cluster file system, the socket would not survive a node failover.
 - Any FIFOs or named pipes that you create on a cluster file system would not be globally accessible.

Therefore, do not attempt to use the `fattach` command from any node other than the local node.

Disk Device Groups

Add this planning information to the [“Disk Device Group Configurations Worksheet”](#) on page 260.

You must configure all volume-manager disk groups as Sun Cluster disk device groups. This configuration enables a secondary node to host multihost disks if the primary node fails. Consider the following points when you plan disk device groups.

- **Failover** – You can configure multihost disks and properly configured volume-manager devices as failover devices. Proper configuration of a volume-manager device includes multihost disks and correct setup of the volume manager itself. This configuration ensures that multiple nodes can host the exported device. You cannot configure tape drives, CD-ROMs, or single-ported devices as failover devices.
- **Mirroring** – You must mirror the disks to protect the data from disk failure. See “Mirroring Guidelines” on page 39 for additional guidelines. See “Installing and Configuring Solstice DiskSuite or Solaris Volume Manager Software” on page 129 or “SPARC: Installing and Configuring VxVM Software” on page 163 and your volume-manager documentation for instructions on mirroring.

For more information about disk device groups, see “Devices” in *Sun Cluster Overview for Solaris OS* and *Sun Cluster Concepts Guide for Solaris OS*.

Mount Information for Cluster File Systems

Consider the following points when you plan mount points for cluster file systems.

- **Mount-point location** – Create mount points for cluster file systems in the `/global` directory, unless you are prohibited by other software products. By using the `/global` directory, you can more easily distinguish cluster file systems, which are globally available, from local file systems.
- **SPARC: VxFS mount requirement** – If you use VERITAS File System (VxFS), globally mount and unmount a VxFS file system from the primary node. The primary node is the node that masters the disk on which the VxFS file system resides. This method ensures that the mount or unmount operation succeeds. A VxFS file-system mount or unmount operation that is performed from a secondary node might fail.
- The following VxFS features are not supported in a Sun Cluster 3.1 cluster file system. They are, however, supported in a local file system.
 - Quick I/O
 - Snapshots
 - Storage checkpoints
 - VxFS-specific mount options:
 - `convosync` (Convert `O_SYNC`)
 - `mincache`
 - `qlog`, `delaylog`, `tmplog`
 - VERITAS cluster file system (requires VxVM cluster feature & VERITAS Cluster Server)

Cache advisories can be used, but the effect is observed on the given node only.

All other VxFS features and options that are supported in a cluster file system are supported by Sun Cluster 3.1 software. See VxFS documentation for details about VxFS options that are supported in a cluster configuration.

- **Nesting mount points** – Normally, you should not nest the mount points for cluster file systems. For example, do not set up one file system that is mounted on `/global/a` and another file system that is mounted on `/global/a/b`. To ignore this rule can cause availability and node boot-order problems. These problems would occur if the parent mount point is not present when the system attempts to mount a child of that file system. The only exception to this rule is if the devices for the two file systems have the same physical node connectivity. An example is different slices on the same disk.
- `forcedirectio` - Sun Cluster software does not support the execution of binaries off cluster file systems that are mounted by using the `forcedirectio` mount option.

Planning Volume Management

Add this planning information to the “[Disk Device Group Configurations Worksheet](#)” on page 260 and the “[Volume-Manager Configurations Worksheet](#)” on page 262. For Solstice DiskSuite or Solaris Volume Manager, also add this planning information to the “[Metadevices Worksheet \(Solstice DiskSuite or Solaris Volume Manager\)](#)” on page 264.

This section provides the following guidelines for planning volume management of your cluster configuration:

- “[Guidelines for Volume-Manager Software](#)” on page 34
- “[Guidelines for Solstice DiskSuite or Solaris Volume Manager Software](#)” on page 35
- “[SPARC: Guidelines for VERITAS Volume Manager Software](#)” on page 37
- “[File-System Logging](#)” on page 38
- “[Mirroring Guidelines](#)” on page 39

Sun Cluster software uses volume-manager software to group disks into disk device groups which can then be administered as one unit. Sun Cluster software supports Solstice DiskSuite or Solaris Volume Manager software and VERITAS Volume Manager (VxVM) software that you install or use in the following ways.

TABLE 1-5 Supported Use of Volume Managers with Sun Cluster Software

Volume-Manager Software	Requirements
Solstice DiskSuite or Solaris Volume Manager	You must install Solstice DiskSuite or Solaris Volume Manager software on all nodes of the cluster, regardless of whether you use VxVM on some nodes to manage disks.
SPARC: VxVM with the cluster feature	You must install and license VxVM with the cluster feature on all nodes of the cluster.
SPARC: VxVM without the cluster feature	You are only required to install and license VxVM on those nodes that are attached to storage devices which VxVM manages.
SPARC: Both Solstice DiskSuite or Solaris Volume Manager and VxVM	If you install both volume managers on the same node, you must use Solstice DiskSuite or Solaris Volume Manager software to manage disks that are local to each node. Local disks include the root disk. Use VxVM to manage all shared disks.

See your volume-manager documentation and [“Installing and Configuring Solstice DiskSuite or Solaris Volume Manager Software” on page 129](#) or [“SPARC: Installing and Configuring VxVM Software” on page 163](#) for instructions on how to install and configure the volume-manager software. For more information about volume management in a cluster configuration, see the *Sun Cluster Concepts Guide for Solaris OS*.

Guidelines for Volume-Manager Software

Consider the following general guidelines when you configure your disks with volume-manager software:

- **Software RAID** – Sun Cluster software does not support software RAID.
- **Mirrored multihost disks** – You must mirror all multihost disks across disk expansion units. See [“Guidelines for Mirroring Multihost Disks” on page 39](#) for guidelines on mirroring multihost disks. You do not need to use software mirroring if the storage device provides hardware RAID as well as redundant paths to devices.
- **Mirrored root** – Mirroring the root disk ensures high availability, but such mirroring is not required. See [“Mirroring Guidelines” on page 39](#) for guidelines about deciding whether to mirror the root disk.
- **Unique naming** – You might have local Solstice DiskSuite metadevices, local Solaris Volume Manager volumes, or VxVM volumes that are used as devices on which the `/global/.devices/node@nodeid` file systems are mounted. If so, the name of each local metadvice or local volume must be unique throughout the cluster.

- **Node lists** – To ensure high availability of a disk device group, make its node lists of potential masters and its failback policy identical to any associated resource group. Or, if a scalable resource group uses more nodes than its associated disk device group, make the scalable resource group's node list a superset of the disk device group's node list. See the resource group planning information in the *Sun Cluster Data Service Planning and Administration Guide for Solaris OS* for information about node lists.
- **multihost disks** – You must connect, or port, all devices used to construct a device group within the cluster to all of the nodes that are configured in the node list for that device group. Solstice DiskSuite or Solaris Volume Manager software can automatically check for this connection at the time that devices are added to a disk set. However, configured VxVM disk groups do not have an association to any particular set of nodes.
- **Hot spare disks** – You can use hot spare disks to increase availability, but hot spare disks are not required.

See your volume-manager documentation for disk layout recommendations and any additional restrictions.

Guidelines for Solstice DiskSuite or Solaris Volume Manager Software

Consider the following points when you plan Solstice DiskSuite or Solaris Volume Manager configurations:

- **Local metadvice names or volume names** – The name of each local Solstice DiskSuite metadvice or Solaris Volume Manager volume must be unique throughout the cluster. Also, the name cannot be the same as any device-ID name.
- **Dual-string mediators** – Each disk set configured with exactly two disk strings and mastered by exactly two nodes must have Solstice DiskSuite or Solaris Volume Manager mediators configured for the disk set. A disk string consists of a disk enclosure, its physical disks, cables from the enclosure to the node(s), and the interface adapter cards. Observe the following rules to configure dual-string mediators:
 - You must configure each disk set with exactly two nodes that act as mediator hosts.
 - You must use the same two nodes for all disk sets that require mediators. Those two nodes must master those disk sets.
 - Mediators cannot be configured for disk sets that do not meet the two-string and two-host requirements.

See the `mediator(7D)` man page for details.

- **/kernel/drv/md.conf settings** – All Solstice DiskSuite metadevices or Solaris Volume Manager volumes used by each disk set are created in advance, at reconfiguration boot time. This reconfiguration is based on the configuration

parameters that exist in the `/kernel/drv/md.conf` file.



Caution – All cluster nodes must have identical `/kernel/drv/md.conf` files, regardless of the number of disk sets that are served by each node. Failure to follow this guideline can result in serious Solstice DiskSuite or Solaris Volume Manager errors and possible loss of data.

You must modify the `nmd` and `md_nsets` fields as follows to support a Sun Cluster configuration:

- `md_nsets` – The `md_nsets` field defines the total number of disk sets that can be created for a system to meet the needs of the entire cluster. Set the value of `md_nsets` to the expected number of disk sets in the cluster plus one additional disk set. Solstice DiskSuite or Solaris Volume Manager software uses the additional disk set to manage the private disks on the local host. The private disks are those metadevices or volumes that are not in the local disk set.

The maximum number of disk sets that are allowed per cluster is 32. This number allows for 31 disk sets for general use plus one disk set for private disk management. The default value of `md_nsets` is 4.

- `nmd` – The `nmd` field defines the number of metadevices or volumes that are created for each disk set. Set the value of `nmd` to the predicted highest value of metadvice or volume name that is used by any one of the disk sets in the cluster. For example, if a cluster uses 10 metadevices or volumes in its first 15 disk sets, but 1000 metadevices or volumes in the 16th disk set, set the value of `nmd` to at least 1000. Also, the value of `nmd` must be large enough to ensure that enough numbers exist for each device-ID name. The number must also be large enough to ensure that each local metadvice name or local volume name can be unique throughout the cluster.

The highest allowed value of a metadvice or volume name per disk set is 8192. The default value of `nmd` is 128.

Set these fields at installation time to allow for all predicted future expansion of the cluster. To increase the value of these fields after the cluster is in production is time consuming. The value change requires a reconfiguration reboot for each node. To raise these values later also increases the possibility of inadequate space allocation in the root (`/`) file system to create all of the requested devices.

At the same time, keep the value of the `nmd` field and the `md_nsets` field as low as possible. Memory structures exist for all possible devices as determined by `nmd` and `md_nsets`, even if you have not created those devices. For optimal performance, keep the value of `nmd` and `md_nsets` only slightly higher than the number of metadevices or volumes you plan to use.

See “System and Startup Files” in *Solstice DiskSuite 4.2.1 Reference Guide* or “System Files and Startup Files” in *Solaris Volume Manager Administration Guide* for more information about the `md.conf` file.

SPARC: Guidelines for VERITAS Volume Manager Software

Consider the following points when you plan VERITAS Volume Manager (VxVM) configurations.

- **Enclosure-Based Naming** – If you use Enclosure-Based Naming of devices, ensure that you use consistent device names on all cluster nodes that share the same storage. VxVM does not coordinate these names, so the administrator must ensure that VxVM assigns the same names to the same devices from different nodes. Failure to assign consistent names does not interfere with correct cluster behavior. However, inconsistent names greatly complicate cluster administration and greatly increase the possibility of configuration errors, potentially leading to loss of data.
- **Root disk group** – If you use VxVM 3.5 or earlier, you must create a default root disk group on each node. For VxVM 4.0, the creation of a root disk group is optional.

A root disk group can be created on the following disks:

- The root disk, which must be encapsulated
- One or more local nonroot disks, which you can encapsulate or initialize
- A combination of root and local nonroot disks

The root disk group must be local to the node.

- **Simple root disk groups** – Simple root disk groups (`rootdg` created on a single slice of the root disk) are not supported as disk types with VxVM on Sun Cluster software. This is a general VxVM software restriction.
- **Encapsulation** – Disks to be encapsulated must have two disk-slice table entries free.
- **Number of volumes** – Estimate the maximum number of volumes any given disk device group can use at the time the disk device group is created.
 - If the number of volumes is less than 1000, you can use default minor numbering.
 - If the number of volumes is 1000 or greater, you must carefully plan the way in which minor numbers are assigned to disk device group volumes. No two disk device groups can have overlapping minor number assignments.
- **Dirty Region Logging** – The use of Dirty Region Logging (DRL) decreases volume recovery time after a node failure. Using DRL might decrease I/O throughput.
- **Dynamic Multipathing (DMP)** – The use of DMP alone to manage multiple I/O paths per node to the shared storage is not supported. The use of DMP is supported only in the following configurations:
 - A single I/O path per node to the cluster's shared storage.
 - A supported multipathing solution, such as Sun Traffic Manager, EMC PowerPath, or Hitachi HDLM, that manages multiple I/O paths per node to the shared cluster storage.

See your VxVM installation documentation for additional information.

File-System Logging

Logging is required for UFS and VxFS cluster file systems. This requirement does not apply to QFS shared file systems. Sun Cluster software supports the following choices of file-system logging:

- Solaris UFS logging – See the `mount_ufs(1M)` man page for more information.
- Solstice DiskSuite trans-metadevice logging or Solaris Volume Manager transactional-volume logging – See “Creating DiskSuite Objects” in *Solstice DiskSuite 4.2.1 User’s Guide* or “Transactional Volumes (Overview)” in *Solaris Volume Manager Administration Guide* for more information.
- SPARC: VERITAS File System (VxFS) logging – See the `mount_vxfs` man page provided with VxFS software for more information.

The following table lists the file-system logging supported by each volume manager.

TABLE 1-6 Supported File System Logging Matrix

Volume Manager	Supported File System Logging
Solstice DiskSuite or Solaris Volume Manager	<ul style="list-style-type: none">■ Solaris UFS logging■ Solstice DiskSuite trans-metadevice logging■ Solaris Volume Manager transactional-volume logging■ VxFS logging
SPARC: VERITAS Volume Manager	<ul style="list-style-type: none">■ Solaris UFS logging■ VxFS logging

Consider the following points when you choose between Solaris UFS logging and Solstice DiskSuite trans-metadevice logging/Solaris Volume Manager transactional-volume logging for UFS cluster file systems:

- Solaris Volume Manager transactional-volume logging (formerly Solstice DiskSuite trans-metadevice logging) is scheduled to be removed from the Solaris OS in an upcoming Solaris release. Solaris UFS logging provides the same capabilities but superior performance, as well as lower system administration requirements and overhead.
- **Solaris UFS log size** – Solaris UFS logging always allocates the log by using free space on the UFS file system, and depending on the size of the file system.
 - On file systems less than 1 Gbyte, the log occupies 1 Mbyte.
 - On file systems 1 Gbyte or greater, the log occupies 1 Mbyte per Gbyte on the file system, to a maximum of 64 Mbytes.

- **Log metadvice/transactional volume** – A Solstice DiskSuite trans metadvice or Solaris Volume Manager transactional volume manages UFS logging. The logging device component of a trans metadvice or transactional volume is a metadvice or volume that you can mirror and stripe. You can create a maximum 1-Gbyte log size, although 64 Mbytes is sufficient for most file systems. The minimum log size is 1 Mbyte.

Mirroring Guidelines

This section provides the following guidelines for planning the mirroring of your cluster configuration:

- [“Guidelines for Mirroring Multihost Disks”](#) on page 39
- [“Guidelines for Mirroring the Root Disk”](#) on page 39

Guidelines for Mirroring Multihost Disks

To mirror all multihost disks in a Sun Cluster configuration enables the configuration to tolerate single-device failures. Sun Cluster software requires that you mirror all multihost disks across expansion units. You do not need to use software mirroring if the storage device provides hardware RAID as well as redundant paths to devices.

Consider the following points when you mirror multihost disks.

- **Separate disk expansion units** – Each submirror of a given mirror or plex should reside in a different multihost expansion unit.
- **Disk space** – Mirroring doubles the amount of necessary disk space.
- **Three-way mirroring** – Solstice DiskSuite or Solaris Volume Manager software and VERITAS Volume Manager (VxVM) software support three-way mirroring. However, Sun Cluster software requires only two-way mirroring.
- **Number of metadvicees or volumes** – Under Solstice DiskSuite or Solaris Volume Manager software, mirrors consist of other Solstice DiskSuite metadvicees or Solaris Volume Manager volumes such as concatenations or stripes. Large configurations might contain a large number of metadvicees or volumes.
- **Differing device sizes** – If you mirror to a device of a different size, your mirror capacity is limited to the size of the smallest submirror or plex.

For more information about multihost disks, see *“Multihost Disk Storage”* in *Sun Cluster Overview for Solaris OS* and *Sun Cluster Concepts Guide for Solaris OS*.

Guidelines for Mirroring the Root Disk

Add this planning information to the [“Local File System Layout Worksheet”](#) on page 254.

For maximum availability, mirror root (/), /usr, /var, /opt, and swap on the local disks. Under VxVM, you encapsulate the root disk and mirror the generated subdisks. However, Sun Cluster software does not require that you mirror the root disk.

Before you decide whether to mirror the root disk, consider the risks, complexity, cost, and service time for the various alternatives that concern the root disk. No single mirroring strategy works for all configurations. You might want to consider your local Sun service representative's preferred solution when you decide whether to mirror root.

See your volume-manager documentation and [“Installing and Configuring Solstice DiskSuite or Solaris Volume Manager Software” on page 129](#) or [“SPARC: Installing and Configuring VxVM Software” on page 163](#) for instructions on how to mirror the root disk.

Consider the following points when you decide whether to mirror the root disk.

- **Boot disk** – You can set up the mirror to be a bootable root disk. You can then boot from the mirror if the primary boot disk fails.
- **Complexity** – To mirror the root disk adds complexity to system administration. To mirror the root disk also complicates booting in single-user mode.
- **Backups** – Regardless of whether you mirror the root disk, you also should perform regular backups of root. Mirroring alone does not protect against administrative errors. Only a backup plan enables you to restore files that have been accidentally altered or deleted.
- **Quorum devices** – Do not use a disk that was configured as a quorum device to mirror a root disk.
- **Quorum** – Under Solstice DiskSuite or Solaris Volume Manager software, in failure scenarios in which state database quorum is lost, you cannot reboot the system until maintenance is performed. See your Solstice DiskSuite or Solaris Volume Manager documentation for information about the state database and state database replicas.
- **Separate controllers** – Highest availability includes mirroring the root disk on a separate controller.
- **Secondary root disk** – With a mirrored root disk, the primary root disk can fail but work can continue on the secondary (mirror) root disk. Later, the primary root disk might return to service, for example, after a power cycle or transient I/O errors. Subsequent boots are then performed by using the primary root disk that is specified for the `eeprom(1M) boot-device` parameter. In this situation, no manual repair task occurs, but the drive starts working well enough to boot. With Solstice DiskSuite or Solaris Volume Manager software, a resync does occur. A resync requires a manual step when the drive is returned to service.

If changes were made to any files on the secondary (mirror) root disk, they would not be reflected on the primary root disk during boot time. This condition would cause a stale submirror. For example, changes to the `/etc/system` file would be lost. With Solstice DiskSuite or Solaris Volume Manager software, some administrative commands might have changed the `/etc/system` file while the primary root disk was out of service.

The boot program does not check whether the system is booting from a mirror or from an underlying physical device. The mirroring becomes active partway through the boot process, after the metadevices or volumes are loaded. Before this point, the system is therefore vulnerable to stale submirror problems.

Installing and Configuring Sun Cluster Software

This chapter provides procedures for how to install and configure your cluster. You can also use these procedures to add a new node to an existing cluster. This chapter also provides procedures to uninstall certain cluster software.

The following information and procedures are in this chapter.

- “Task Map: Installing the Software” on page 44
- “How to Prepare for Cluster Software Installation” on page 46
- “How to Install Cluster Control Panel Software on an Administrative Console” on page 48
- “How to Install Solaris Software” on page 52
- “How to Install Sun Cluster Software Packages” on page 56
- “How to Configure Sun Cluster Software on All Nodes (`scinstall`)” on page 57
- “Using SunPlex Installer to Install Sun Cluster Software” on page 63
- “How to Install SunPlex Manager Software” on page 65
- “How to Install and Configure Sun Cluster Software (SunPlex Installer)” on page 70
- “How to Install Solaris and Sun Cluster Software (JumpStart)” on page 76
- “How to Install Sun Cluster Software on a Single-Node Cluster” on page 87
- “How to Configure Sun Cluster Software on Additional Cluster Nodes (`scinstall`)” on page 90
- “SPARC: How to Install VERITAS File System Software” on page 97
- “How to Configure the Name-Service Switch” on page 98
- “How to Set Up the Root Environment” on page 99
- “How to Install Data-Service Software Packages (`installer`)” on page 100
- “How to Install Data-Service Software Packages (`scinstall`)” on page 102
- “How to Perform Postinstallation Setup and Configure Quorum Devices” on page 104
- “How to Verify the Quorum Configuration and Installation Mode” on page 107

- "Task Map: Configuring the Cluster" on page 107
- "How to Create Cluster File Systems" on page 108
- "How to Configure Internet Protocol (IP) Network Multipathing Groups" on page 113
- "How to Change Private Hostnames" on page 115
- "How to Configure Network Time Protocol (NTP)" on page 116
- "SPARC: Task Map: Installing the Sun Cluster Module for Sun Management Center" on page 117
- "SPARC: Installation Requirements for Sun Cluster Monitoring" on page 118
- "SPARC: How to Install the Sun Cluster Module for Sun Management Center" on page 119
- "SPARC: How to Start Sun Management Center" on page 120
- "SPARC: How to Add a Cluster Node as a Sun Management Center Agent Host Object" on page 121
- "SPARC: How to Load the Sun Cluster Module" on page 121
- "How to Uninstall Sun Cluster Software to Correct Installation Problems" on page 123
- "How to Uninstall the SUNWscrdt Package" on page 124
- "How to Unload the RSMRDT Driver Manually" on page 125

Task Map: Installing the Software

The following task map lists the tasks that you perform to install software on multinode or single-node clusters. Complete the procedures in the order that is indicated.

TABLE 2-1 Task Map: Installing the Software

Task	Instructions
1. Plan the layout of your cluster configuration and prepare to install software.	"How to Prepare for Cluster Software Installation" on page 46
2. (Optional) Install Cluster Control Panel (CCP) software on the administrative console.	"How to Install Cluster Control Panel Software on an Administrative Console" on page 48

TABLE 2-1 Task Map: Installing the Software (Continued)

Task	Instructions
3. Install the Solaris OS and Sun Cluster software. Optionally, install Sun StorEdge QFS software. Choose one of the following methods:	
<ul style="list-style-type: none"> ■ Method 1 – (New clusters only) Install Solaris software. Next, install Sun Cluster software on all nodes. Then use the <code>scinstall</code> utility to establish the cluster. 	<ol style="list-style-type: none"> 1. “How to Install Solaris Software” on page 52 2. “How to Install Sun Cluster Software Packages” on page 56 3. “How to Configure Sun Cluster Software on All Nodes (<code>scinstall</code>)” on page 57
<ul style="list-style-type: none"> ■ Method 2 – (New clusters only) Install Solaris software. Next, install SunPlex™ Manager software. Then use SunPlex Installer to install Sun Cluster software. 	<ol style="list-style-type: none"> 1. “How to Install Solaris Software” on page 52 2. “Using SunPlex Installer to Install Sun Cluster Software” on page 63
<ul style="list-style-type: none"> ■ Method 3 – (New clusters or added nodes) Install Solaris software and Sun Cluster software in one operation by using the <code>scinstall</code> utility’s custom JumpStart option. 	<p>“How to Install Solaris and Sun Cluster Software (JumpStart)” on page 76</p>
<ul style="list-style-type: none"> ■ Method 4 – (New single-node clusters) Install Solaris software. Then install Sun Cluster software by using the <code>scinstall -iFo</code> command. 	<ol style="list-style-type: none"> 1. “How to Install Solaris Software” on page 52 2. “How to Install Sun Cluster Software on a Single-Node Cluster” on page 87
<ul style="list-style-type: none"> ■ Method 5 – (Added nodes only) Install Solaris software on the new nodes. Next, install Sun Cluster software on the new node. Then configure Sun Cluster software on the new node by using the <code>scinstall</code> utility. 	<ol style="list-style-type: none"> 1. “How to Install Solaris Software” on page 52 2. “How to Install Sun Cluster Software Packages” on page 56 3. “How to Configure Sun Cluster Software on Additional Cluster Nodes (<code>scinstall</code>)” on page 90
4. (Optional) SPARC: Install VERITAS File System software.	“SPARC: How to Install VERITAS File System Software” on page 97
5. Configure the name-service look-up order.	“How to Configure the Name-Service Switch” on page 98
6. Set up directory paths.	“How to Set Up the Root Environment” on page 99
7. Install data-service software packages.	“How to Install Data-Service Software Packages (<code>installer</code>)” on page 100 or “How to Install Data-Service Software Packages (<code>scinstall</code>)” on page 102
8. Assign quorum votes and remove the cluster from installation mode, if this operation was not already performed during Sun Cluster installation.	“How to Perform Postinstallation Setup and Configure Quorum Devices” on page 104
9. Validate the quorum configuration.	“How to Verify the Quorum Configuration and Installation Mode” on page 107

TABLE 2-1 Task Map: Installing the Software (Continued)

Task	Instructions
10. Install and configure volume-manager software:	
<ul style="list-style-type: none"> ■ Install and configure Solstice DiskSuite or Solaris Volume Manager software. 	<ul style="list-style-type: none"> ■ “Installing and Configuring Solstice DiskSuite or Solaris Volume Manager Software” on page 129 ■ Solstice DiskSuite or Solaris Volume Manager documentation
<ul style="list-style-type: none"> ■ SPARC: Install and configure VERITAS Volume Manager software. 	<ul style="list-style-type: none"> ■ “SPARC: Installing and Configuring VxVM Software” on page 163 ■ VERITAS Volume Manager documentation
11. Configure the cluster.	“Configuring the Cluster” on page 108

Installing the Software

This section provides information and procedures to install software on the cluster nodes.

▼ How to Prepare for Cluster Software Installation

Before you begin to install software, make the following preparations.

1. **Read the following manuals for information that can help you plan your cluster configuration and prepare your installation strategy.**
 - *Sun Cluster Release Notes for Solaris OS* – Restrictions, bug workarounds, and other late-breaking information.
 - *Sun Cluster 3.x Release Notes Supplement* – Post-release documentation about additional restrictions, bug workarounds, new features, and other late-breaking information. This document is regularly updated and published online at the following Web site.
<http://docs.sun.com>
 - *Sun Cluster Overview for Solaris OS* and *Sun Cluster Concepts Guide for Solaris OS* – Overviews of the Sun Cluster product.
 - *Sun Cluster Software Installation Guide for Solaris OS* (this manual) – Planning guidelines and procedures for installing and configuring Solaris, Sun Cluster, and volume-manager software.
 - *Sun Cluster Data Service Planning and Administration Guide for Solaris OS* – Planning guidelines and procedures to install and configure data services.

- Documentation for all third-party software products.

2. Have available all related documentation, including third-party documents.

The following is a partial list of products whose documentation you might need to reference during cluster installation:

- Solaris OS
- Solstice DiskSuite or Solaris Volume Manager software
- Sun StorEdge QFS software
- SPARC: VERITAS Volume Manager
- SPARC: Sun Management Center
- Third-party applications

3. Plan your cluster configuration.



Caution – Plan your cluster installation completely. Identify requirements for all data services and third-party products **before** you begin Solaris and Sun Cluster software installation. Failure to do so might result in installation errors that require that you completely reinstall the Solaris and Sun Cluster software.

For example, the Oracle Real Application Clusters Guard option of Oracle Real Application Clusters has special requirements for the hostnames that you use in the cluster. Another example with special requirements is Sun Cluster HA for SAP. You must accommodate these requirements before you install Sun Cluster software because you cannot change hostnames after you install Sun Cluster software.

Also note that both Oracle Real Application Clusters and Sun Cluster HA for SAP are not supported for use in x86 based clusters.

- Use the planning guidelines in [Chapter 1](#) and in the *Sun Cluster Data Service Planning and Administration Guide for Solaris OS* to determine how to install and configure your cluster.
- Fill out the cluster framework and data-services configuration worksheets that are referenced in the planning guidelines. Use your completed worksheets for reference during the installation and configuration tasks.

4. Get all necessary patches for your cluster configuration.

See “Patches and Required Firmware Levels” in *Sun Cluster 3.1 Release Notes* for the location of patches and installation instructions.

a. Copy the patches that are required for Sun Cluster into a single directory.

The directory must be on a file system that is accessible by all nodes. The default patch directory is `/var/cluster/patches/`.

Tip – After you install Solaris software on a node, you can view the `/etc/release` file to see the exact version of Solaris software that is installed.

b. (Optional) If you are not using SunPlex Installer, you can create a patch list file.

If you specify a patch list file, SunPlex Installer only installs the patches that are listed in the patch list file. For information about creating a patch-list file, refer to the `patchadd(1M)` man page.

c. Record the path to the patch directory.

5. (Optional) Use Cluster Control Panel software to connect from an administrative console to your cluster nodes.

Go to “[How to Install Cluster Control Panel Software on an Administrative Console](#)” on page 48.

6. Choose the Solaris installation procedure to use.

- If you intend to install Sun Cluster software by using either the `scinstall(1M)` utility (text-based method) or SunPlex Installer (GUI-based method), go to “[How to Install Solaris Software](#)” on page 52 to first install Solaris software.
- If you intend to install Solaris and Sun Cluster software in the same operation (JumpStart based method), go to “[How to Install Solaris and Sun Cluster Software \(JumpStart\)](#)” on page 76.

▼ How to Install Cluster Control Panel Software on an Administrative Console

Note – You are not required to use an administrative console. If you do not use an administrative console, perform administrative tasks from one designated node in the cluster.

This procedure describes how to install the Cluster Control Panel (CCP) software on an administrative console. The CCP provides a launchpad for the `cconsole(1M)`, `ctelnet(1M)`, and `crlogin(1M)` tools. Each of these tools provides a multiple-window connection to a set of nodes, as well as a common window. You can use the common window to send input to all nodes at one time.

You can use any desktop machine that runs the Solaris 8 or Solaris 9 OS as an administrative console. In addition, you can also use the administrative console as a documentation server. If you are using Sun Cluster on a SPARC based system, you can

use the administrative console as a Sun Management Center console or server as well. See Sun Management Center documentation for information on how to install Sun Management Center software. See the *Sun Cluster Release Notes for Solaris OS* for additional information on how to install Sun Cluster documentation.

1. Become superuser on the administrative console.

2. Ensure that a supported version of the Solaris OS and any Solaris patches are installed on the administrative console.

All platforms require at least the End User Solaris Software Group.

3. Insert the Sun Cluster 3.1 9/04 CD-ROM into the CD-ROM drive of the administrative console.

If the volume management daemon `vold(1M)` is running and is configured to manage CD-ROM devices, the daemon automatically mounts the CD-ROM on the `/cdrom/cdrom0/` directory.

4. Change to the `Solaris_arch/Product/sun_cluster/Solaris_ver/Packages/` directory, where `arch` is `sparc` or `x86`, and where `ver` is 8 (for Solaris 8) or 9 (for Solaris 9) .

```
# cd /cdrom/cdrom0/Solaris_arch/Product/sun_cluster/Solaris_ver/Packages/
```

5. Start the installer program.

```
# ./installer
```

6. Choose Custom installation.

The utility displays a list of software packages.

7. If necessary, deselect any packages that you do not want to install on the administrative console.

8. Choose the menu item, Sun Cluster `cconsole` package.

9. (Optional) Choose the menu item, Sun Cluster `manpage` package.

10. (Optional) Choose the menu item, Sun Cluster documentation packages.

11. Follow onscreen instructions to continue package installation.

After installation is finished, you can view any available installation log.

12. Install the `SUNWcccon` package.

```
# pkgadd -d . SUNWcccon
```

13. (Optional) Install the `SUNWscman` package.

```
# pkgadd -d . SUNWscman
```

When you install the `SUNWscman` package on the administrative console, you can view Sun Cluster man pages from the administrative console before you install Sun Cluster software on the cluster nodes.

14. (Optional) Install the Sun Cluster documentation packages.

Note – If you do not install the documentation on your administrative console, you can still view HTML or PDF documentation directly from the CD-ROM. Use a web browser to view the `index.html` file at the top level of the CD-ROM.

a. Start the `pkgadd` utility in interactive mode.

```
# pkgadd -d .
```

b. Choose the Documentation Navigation for Solaris 9 package, if it has not already been installed on the administrative console.

c. Choose the Sun Cluster documentation packages to install.

The following documentation collections are available in both HTML and PDF format:

- Sun Cluster 3.1 9/04 Software Collection for Solaris OS (SPARC Platform Edition)
- Sun Cluster 3.1 9/04 Software Collection for Solaris OS (x86 Platform Edition)
- Sun Cluster 3.x Hardware Collection for Solaris OS (SPARC Platform Edition)
- Sun Cluster 3.x Hardware Collection for Solaris OS (x86 Platform Edition)
- Sun Cluster 3.1 9/04 Reference Collection for Solaris OS

d. Follow onscreen instructions to continue package installation.

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

```
# eject cdrom
```

16. Create an `/etc/clusters` file on the administrative console.

Add your cluster name and the physical node name of each cluster node to the file.

```
# vi /etc/clusters  
clustername node1 node2
```

See the `/opt/SUNWcluster/bin/clusters(4)` man page for details.

17. Create an `/etc/serialports` file.

Add an entry for each node in the cluster to the file. Specify the physical node name, the hostname of the console-access device, and the port number. Examples of a console-access device are a terminal concentrator (TC), a System Service Processor (SSP), and a Sun Fire system controller.

```
# vi /etc/serialports  
node1 ca-dev-hostname port  
node2 ca-dev-hostname port
```

node1, node2 Physical names of the cluster nodes
ca-dev-hostname Hostname of the console-access device
port Serial port number

Note these special instructions to create an `/etc/serialports` file:

- For a Sun Fire 15000 system controller, use `telnet(1)` port number 23 for the serial port number of each entry.
- For all other console-access devices, use the `telnet` serial port number, not the physical port number. To determine the telnet serial port number, add 5000 to the physical port number. For example, if a physical port number is 6, the telnet serial port number is 5006.
- For Sun Enterprise 10000 servers, also see the `/opt/SUNWcluster/bin/serialports(4)` man page for details and special considerations.

18. (Optional) For convenience, set the directory paths on the administrative console.

- Add the `/opt/SUNWcluster/bin/` directory to the `PATH`.
- Add the `/opt/SUNWcluster/man/` directory to the `MANPATH`.
- If you installed the `SUNWscman` package, also add the `/usr/cluster/man/` directory to the `MANPATH`.

19. Start the CCP utility.

```
# /opt/SUNWcluster/bin/ccp &
```

Click the `cconsole`, `crlogin`, or `ctelnet` button in the CCP window to launch that tool. Alternately, you can start any of these tools directly. For example, to start `ctelnet`, type the following command:

```
# /opt/SUNWcluster/bin/ctelnet &
```

See the procedure “How to Remotely Log In to Sun Cluster” in “Beginning to Administer the Cluster” in *Sun Cluster System Administration Guide for Solaris OS* for additional information about how to use the CCP utility. Also see the `ccp(1M)` man page.

20. Determine whether the Solaris OS is already installed on each cluster node to meet Sun Cluster software requirements.

- If the Solaris OS meets Sun Cluster requirements, go to [“How to Install Sun Cluster Software Packages” on page 56](#).
- If the Solaris OS does not meet Sun Cluster requirements, install, reconfigure, or reinstall the Solaris OS as needed. See [“Planning the Solaris OS” on page 16](#) for information about Sun Cluster installation requirements for the Solaris OS.
 - To install only Solaris software, go to [“How to Install Solaris Software” on page 52](#).

- To install both Solaris and Sun Cluster software by using the `scinstall` JumpStart option, go to [“How to Install Solaris and Sun Cluster Software \(JumpStart\)”](#) on page 76.

▼ How to Install Solaris Software

If you do not use the `scinstall(1M)` custom JumpStart installation method to install software, perform this task. Follow these procedures to install the Solaris OS on each node in the cluster.

Tip – To speed installation, you can install the Solaris OS on each node at the same time.

If your nodes are already installed with the Solaris OS but do not meet Sun Cluster installation requirements, you might need to reinstall the Solaris software. Follow the steps in this procedure to ensure subsequent successful installation of Sun Cluster software. See [“Planning the Solaris OS”](#) on page 16 for information about required root-disk partitioning and other Sun Cluster installation requirements.

- 1. Ensure that the hardware setup is complete and that connections are verified before you install Solaris software.**
See the *Sun Cluster Hardware Administration Collection* and your server and storage device documentation for details.
- 2. Ensure that your cluster configuration planning is complete.**
See [“How to Prepare for Cluster Software Installation”](#) on page 46 for requirements and guidelines.
- 3. Have available your completed “Local File System Layout Worksheet” on page 254.**
- 4. If you use a naming service, add address-to-name mappings for all public hostnames and logical addresses to any naming services that clients use for access to cluster services. You set up local hostname information in Step 11.**
See [“IP Addresses”](#) on page 22 for planning guidelines. See your Solaris system-administrator documentation for information about using Solaris naming services.
- 5. If you are using a cluster administrative console, display a console screen for each node in the cluster.**

- If Cluster Control Panel (CCP) software is installed and configured on your administrative console, you can use the `cconsole(1M)` utility to display the individual console screens. The `cconsole` utility also opens a master window from which you can send your input to all individual console windows at the same time. Use the following command to start `cconsole`:

```
# /opt/SUNWcluster/bin/cconsole clustername &
```

- If you do not use the `cconsole` utility, connect to the consoles of each node individually.

6. Install the Solaris OS as instructed in your Solaris installation documentation.

Note – You must install all nodes in a cluster with the same version of the Solaris OS.

You can use any method that is normally used to install Solaris software. During Solaris software installation, perform the following steps:

a. Install at least the End User Solaris Software Group.

See [“Solaris Software Group Considerations” on page 17](#) for information about additional Solaris software requirements.

b. Choose Manual Layout to set up the file systems.

- Create a file system of at least 512 Mbytes for use by the global-device subsystem. If you intend to use SunPlex Installer to install Sun Cluster software, you must create the file system with a mount-point name of `/globaldevices`. The `/globaldevices` mount-point name is the default that is used by `scinstall`.

Note – Sun Cluster software requires a global-devices file system for installation to succeed.

- Specify that slice 7 is at least 20 Mbytes in size. If you intend to use SunPlex Installer to install Solstice DiskSuite software (Solaris 8) or configure Solaris Volume Manager software (Solaris 9), also make this file system mounted on `/sds`.
- Create any other file-system partitions that you need, as described in [“System Disk Partitions” on page 18](#).

Note – If you intend to install Sun Cluster HA for NFS or Sun Cluster HA for Apache, you must also install Solstice DiskSuite software (Solaris 8) or configure Solaris Volume Manager software (Solaris 9).

- c. For ease of administration, set the same root password on each node.
7. If you are adding a node to an existing cluster, prepare the cluster to accept the new node.
 - a. On any active cluster member, start the `scsetup(1M)` utility.

```
# scsetup
```

The Main Menu is displayed.
 - b. Choose the menu item, New nodes.
 - c. Choose the menu item, Specify the name of a machine which may add itself.
 - d. Follow the prompts to add the node's name to the list of recognized machines.

The `scsetup` utility prints the message Command completed successfully if the task completes without error.
 - e. Quit the `scsetup` utility.
 - f. From the active cluster node, display the names of all cluster file systems.

```
% mount | grep global | egrep -v node@ | awk '{print $1}'
```
 - g. On the new node, create a mount point for each cluster file system in the cluster.

```
% mkdir -p mountpoint
```

For example, if the mount command returned the file-system name `/global/dg-schost-1`, run `mkdir -p /global/dg-schost-1` on the new node you are adding to the cluster.
 - h. Determine whether VERITAS Volume Manager (VxVM) is installed on any nodes that are already in the cluster.
 - i. If VxVM is installed on any existing cluster nodes, ensure that the same `vxio` number is used on the VxVM-installed nodes. Also ensure that the `vxio` number is available for use on each of the nodes that do not have VxVM installed.

```
# grep vxio /etc/name_to_major
vxio NNN
```

If the vxio number is already in use on a node that does not have VxVM installed, free the number on that node. Change the `/etc/name_to_major` entry to use a different number.

- 8. If you installed the End User Solaris Software Group, use the `pkgadd` command to manually install any additional Solaris software packages that you might need.**

The following Solaris packages are required to support some Sun Cluster functionality.

Feature	Required Solaris Software Packages (shown in installation order)
RSMAPI, RSMRDT drivers, or SCI-PCI adapters (SPARC based clusters only)	SUNWrsm SUNWrsmx SUNWrsmo SUNWrsmox
SunPlex Manager	SUNWapchr SUNWapchu

- 9. Install any hardware-related patches. Also download any needed firmware that is contained in the hardware patches.**

See “Patches and Required Firmware Levels” in *Sun Cluster 3.1 Release Notes* for the location of patches and installation instructions.

- 10. x86: Set the default boot file to `kadb`.**

```
# eeprom boot-file=kadb
```

The setting of this value enables you to reboot the node if you are unable to access a login prompt.

- 11. Update the `/etc/inet/hosts` file on each node with all public hostnames and logical addresses for the cluster.**

Perform this step regardless of whether you are using a naming service.

- 12. (Optional) On Sun Enterprise 10000 servers, configure the `/etc/system` file to use dynamic reconfiguration.**

Add the following entry to the `/etc/system` file on each node of the cluster:

```
set kernel_cage_enable=1
```

This entry becomes effective after the next system reboot.

See the *Sun Cluster System Administration Guide for Solaris OS* for procedures to perform dynamic reconfiguration tasks in a Sun Cluster configuration. See your server documentation for more information about dynamic reconfiguration.

- 13. Install Sun Cluster software packages.**

Go to “How to Install Sun Cluster Software Packages” on page 56.

▼ How to Install Sun Cluster Software Packages

Note – If you enable remote shell (`rsh(1M)`) or secure shell (`ssh(1)`) access for superuser to all cluster nodes, you do not have to perform this procedure. Instead, go to “[How to Configure Sun Cluster Software on All Nodes \(`scinstall`\)](#)” on page 57. In that procedure, the `scinstall(1M)` utility automatically installs Sun Cluster framework software on all cluster nodes.

However, if you need to install any Sun Cluster software packages in addition to the framework software, install those packages from the Sun Cluster 3.1 9/04 CD-ROM. Do this task before you start the `scinstall` utility. You can install those additional Sun Cluster packages by using the `pkgadd(1M)` command or by using the `installer(1M)` program as described in the following procedure.

Perform this procedure on each node in the cluster to install the Sun Cluster software packages.

1. Ensure that the Solaris OS is installed to support Sun Cluster software.

If Solaris software is already installed on the node, you must ensure that the Solaris installation meets the requirements for Sun Cluster software and any other software that you intend to install on the cluster. See “[How to Install Solaris Software](#)” on page 52 for more information about installing Solaris software to meet Sun Cluster software requirements.

2. Become superuser on the cluster node to install.

3. Install Sun Web Console packages.

These packages are required by Sun Cluster software, even if you do not use Sun Web Console.

a. Insert the Sun Cluster 3.1 9/04 CD-ROM in the CD-ROM drive.

b. Change to the
`/cdrom/cdrom0/Solaris_arch/Product/sun_web_console/2.1/`
directory, where *arch* is `sparc` or `x86`.

c. Run the `setup` command.

```
# ./setup
```

The `setup` command installs all packages to support Sun Web Console.

4. (Optional) To use the `installer` program with a GUI, ensure that the `DISPLAY` environment variable is set.

5. Change to the root directory of the CD-ROM, where the `installer` program resides.

```
# cd /cdrom/cdrom0/
```


6. Start the installer program.

```
# ./installer
```

7. Choose Typical or Custom installation.

- Choose Typical to install the default set of Sun Cluster framework software packages.
- Choose Custom to specify additional Sun Cluster software packages to install, such as packages that support other languages, the RSM API, and SCI-PCI adapters.

8. Follow instructions on the screen to install Sun Cluster software on the node.

After installation is finished, you can view any available installation log.

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

```
# eject cdrom
```

10. Configure Sun Cluster software on the cluster nodes.

- To establish a new cluster, go to “How to Configure Sun Cluster Software on All Nodes (`scinstall`)” on page 57.
- To add a new node to an existing cluster, go to “How to Configure Sun Cluster Software on Additional Cluster Nodes (`scinstall`)” on page 90.

▼ How to Configure Sun Cluster Software on All Nodes (`scinstall`)

Perform this procedure from one node of the cluster to configure Sun Cluster software on all nodes of the cluster.

1. Ensure that the Solaris OS is installed to support Sun Cluster software.

If Solaris software is already installed on the node, you must ensure that the Solaris installation meets the requirements for Sun Cluster software and any other software that you intend to install on the cluster. See “How to Install Solaris Software” on page 52 for more information about installing Solaris software to meet Sun Cluster software requirements.

2. If you disabled remote configuration during Sun Cluster software installation, re-enable remote configuration.

Enable remote shell (`rsh(1M)`) or secure shell (`ssh(1)`) access for superuser to all cluster nodes.

3. (Optional) To use the `scinstall(1M)` utility to install patches, download patches to a patch directory.

If you use Typical mode to install the cluster, use a directory named either `/var/cluster/patches/` or `/var/patches/` to contain the patches to install.

In Typical mode, the `scinstall` command checks both those directories for patches.

- If neither of those directories exist, then no patches are added.
- If both directories exist, then only the patches in the `/var/cluster/patches/` directory are added.

If you use Custom mode to install the cluster, you specify the path to the patch directory, so there is no requirement to use the patch directories that `scinstall` checks for in Typical mode.

You can include a patch-list file in the patch directory. The default patch-list file name is `patchlist`. For information about creating a patch-list file, refer to the `patchadd(1M)` man page.

4. Complete one of the following configuration worksheets:

- [Table 2-2](#) for Typical configuration
- [Table 2-3](#) for Custom configuration

See “[Planning the Sun Cluster Environment](#)” on page 21 for planning guidelines.

TABLE 2-2 Interactive `scinstall` Configuration Worksheet (Typical)

Component	Description/Example	Enter Answers Here	
Cluster Name	What is the name of the cluster that you want to establish?		
Cluster Nodes	What are the names of the other cluster nodes planned for the initial cluster configuration?		
Cluster-Transport Adapters and Cables	What are the names of the two cluster-transport adapters that attach the node to the private interconnect?	<i>First</i>	<i>Second</i>
Check	Do you want to interrupt installation for <code>sccheck</code> errors? (<i>sccheck verifies that preconfiguration requirements are met</i>)	Yes No	

For the Typical configuration of Sun Cluster software, the `scinstall` utility automatically specifies the following defaults.

Component	Default Value
Private-network address	172.16.0.0
Private-network netmask	255.255.0.0
Cluster-transport junctions	switch1 and switch2
Global-devices file-system name	/globaldevices
Installation security (DES)	Limited
Solaris and Sun Cluster patch directory	/var/cluster/patches/

TABLE 2-3 Interactive `scinstall` Configuration Worksheet (Custom)

Component	Description/Example	Enter Answers Here	
Cluster Name	What is the name of the cluster that you want to establish?		
Cluster Nodes	What are the names of the other cluster nodes planned for the initial cluster configuration?		
DES Authentication	Do you need to use DES authentication?	No Yes	
Network Address for the Cluster Transport	Do you want to accept the default network address (172.16.0.0)?	Yes No	
	If no, supply your own network address:	____.____.0.0	
	Do you want to accept the default netmask (255.255.0.0)?	Yes No	
	If no, supply your own netmask:	255.255.____.____	
Point-to-Point Cables	If this is a two-node cluster, does this cluster use transport junctions?	Yes No	
Cluster-Transport Junctions	If used, what are the names of the two transport junctions? Defaults: <code>switch1</code> and <code>switch2</code>	<i>First</i>	<i>Second</i>
Cluster-Transport Adapters and Cables	Node name (<i>the node from which you run scinstall</i>):		
	Transport adapters:	<i>First</i>	<i>Second</i>
	Where does each transport adapter connect to (<i>a transport junction or another adapter</i>)? Junction defaults: <code>switch1</code> and <code>switch2</code>		
	For transport junctions, do you want to use the default port name?	Yes No	Yes No
	If no, what is the name of the port that you want to use?		
	Do you want to use autodiscovery to list the available adapters for the other nodes? If no, supply the following information for each additional node:	Yes No	
<i>Specify for each additional node</i>	Node name:		
	Transport adapters:	<i>First</i>	<i>Second</i>
	Where does each transport adapter connect to (<i>a transport junction or another adapter</i>)? Defaults: <code>switch1</code> and <code>switch2</code>		
	For transport junctions, do you want to use the default port name?	Yes No	Yes No
	If no, what is the name of the port that you want to use?		

TABLE 2-3 Interactive `scinstall` Configuration Worksheet (Custom) (Continued)

Component	Description/Example	Enter Answers Here
Software Patch Installation	Do you want <code>scinstall</code> to install patches for you?	Yes No
	If yes, what is the name of the patch directory?	
	Do you want to use a patch list?	Yes No
Global-Devices File System (specify for each node)	Do you want to use the default name of the global-devices file system (<code>/globaldevices</code>)?	Yes No
	If no, do you want to use an already-existing file system?	Yes No
	What is the name of the file system that you want to use?	
Check	Do you want to interrupt installation for <code>sccheck</code> errors? (<i>sccheck verifies that preconfiguration requirements are met</i>)	Yes No

Note – You cannot change the private-network address and netmask after `scinstall` processing is finished. If you need to use a different private-network address or netmask and the node is still in installation mode, follow the procedures in “How to Uninstall Sun Cluster Software to Correct Installation Problems” on page 123. Then perform the procedures in “How to Install Sun Cluster Software Packages” on page 56 and in this procedure to reinstall the software and configure the node with the correct information.

5. **Become superuser on the cluster node from which you intend to configure the cluster.**
6. **Install additional packages if you intend to use any of the following features.**
 - Remote Shared Memory Application Programming Interface (RSM API)
 - SCI-PCI adapters for the interconnect transport
 - RSMRDT drivers

Note – Use of the RSMRDT driver is restricted to clusters that run an Oracle9i release 2 SCI configuration with RSM enabled. Refer to Oracle9i release 2 user documentation for detailed installation and configuration instructions.

- a. **Determine which packages you must install.**

The following table lists the Sun Cluster 3.1 9/04 packages that each feature requires and the order in which you must install each group of packages. The installer program does not automatically install these packages.

Feature	Additional Sun Cluster 3.1 9/04 Packages to Install
RSMAPI	SUNWscrif
SCI-PCI adapters	SUNWsci SUNWscid SUNWscidx
RSMRDT drivers	SUNWscrdt

b. Ensure that any dependency Solaris packages are already installed.

See Step 8 in “How to Install Solaris Software” on page 52.

c. Insert the Sun Cluster 3.1 9/04 CD-ROM into the CD-ROM drive of a node.

d. Change to the

Solaris_arch/Product/sun_cluster/Solaris_ver/Packages/ directory, where *arch* is **sparc** or **x86**, and where *ver* is **8** (for Solaris 8) or **9** (for Solaris 9) .

```
# cd /cdrom/cdrom0/Solaris_arch/Product/sun_cluster/Solaris_ver/Packages/
```

e. Install the additional packages.

```
# pkgadd -d . packages
```

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

```
# eject cdrom
```

g. Repeat for each additional node in the cluster.

7. On one node, start the `scinstall` utility.

```
# /usr/cluster/bin/scinstall
```

8. Follow these guidelines to use the interactive `scinstall` utility:

- Interactive `scinstall` enables you to type ahead. Therefore, do not press the Return key more than once if the next menu screen does not appear immediately.
- Unless otherwise noted, you can press Control-D to return to either the start of a series of related questions or to the Main Menu.
- Default answers or answers to previous sessions are displayed in brackets ([]) at the end of a question. Press Return to enter the response that is in brackets without typing it.

9. From the Main Menu, choose the menu item, Install a cluster or cluster node.

```
*** Main Menu ***
```

```
Please select from one of the following (*) options:
```

```
* 1) Install a cluster or cluster node
```

```
2) Configure a cluster to be JumpStarted from this install server
3) Add support for new data services to this cluster node
* 4) Print release information for this cluster node
* ?) Help with menu options
* q) Quit
```

Option: 1

10. From the Install Menu, choose the menu item, Install all nodes of a new cluster.

11. From the Type of Installation menu, choose either Typical or Custom.

12. Follow the menu prompts to supply your answers from the worksheet that you completed in Step 4.

The `scinstall` utility installs and configures all cluster nodes and reboots the cluster. The cluster is established when all nodes have successfully booted into the cluster. Sun Cluster installation output is logged in a `/var/cluster/logs/install/scinstall.log.N` file.

13. Install Sun StorEdge QFS file system software.

Follow the procedures for initial installation in the *Sun StorEdge QFS and Sun StorEdge SAM-FS Software Installation and Configuration Guide*.

14. (Optional) SPARC: To install VERITAS File System, go to “SPARC: How to Install VERITAS File System Software” on page 97.

15. Set up the name-service look-up order.

Go to “How to Configure the Name-Service Switch” on page 98.

Example – Configuring Sun Cluster Software on All Nodes

The following example shows the `scinstall` progress messages that are logged as `scinstall` completes configuration tasks on a two-node cluster. The cluster node names are `phys-schost-1` and `phys-schost-2`. The specified adapter names are `qfe2` and `hme2`.

Installation and Configuration

```
Log file - /var/cluster/logs/install/scinstall.log.834
```

```
Testing for "/globaldevices" on "phys-schost-1" ... done
```

```
Testing for "/globaldevices" on "phys-schost-2" ... done
```

```
Checking installation status ... done
```

```
The Sun Cluster software is already installed on "phys-schost-1".
```

```
The Sun Cluster software is already installed on "phys-schost-2".
```

```
Starting discovery of the cluster transport configuration.
```

```
Probing ..

The following connections were discovered:

    phys-schost-1:qfe2  switch1  phys-schost-2:qfe2
    phys-schost-1:hme2  switch2  phys-schost-2:hme2

Completed discovery of the cluster transport configuration.

Started sccheck on "phys-schost-1".
Started sccheck on "phys-schost-2".

sccheck completed with no errors or warnings for "phys-schost-1".
sccheck completed with no errors or warnings for "phys-schost-2".

Configuring "phys-schost-2" ... done
Rebooting "phys-schost-2" ... done

Configuring "phys-schost-1" ... done
Rebooting "phys-schost-1" ...

Log file - /var/cluster/logs/install/scinstall.log.834

Rebooting ...
```

Using SunPlex Installer to Install Sun Cluster Software

Note – To add a new node to an existing cluster, instead follow the procedures in [“How to Configure Sun Cluster Software on Additional Cluster Nodes \(scinstall\)”](#) on page 90.

This section describes how to install SunPlex Manager software. This section also describes how to use SunPlex Installer, the installation module of SunPlex Manager, to install Sun Cluster software and to establish new cluster nodes. You can also use SunPlex Installer to install or configure one or more of the following additional software products:

- (On Solaris 8 only) Solstice DiskSuite software – After it installs Solstice DiskSuite software, SunPlex Installer configures up to three metaset and associated metadevices. SunPlex Installer also creates and mounts cluster file systems for each metaset.
- (On Solaris 9 only) Solaris Volume Manager software – SunPlex Installer configures up to three Solaris Volume Manager volumes. SunPlex Installer also creates and mounts cluster file systems for each volume. Solaris Volume Manager software is already installed as part of Solaris software installation.

- Sun Cluster HA for NFS data service.
- Sun Cluster HA for Apache scalable data service.

Installation Requirements

The following table lists SunPlex Installer installation requirements for these additional software products.

TABLE 2-4 Requirements to Use SunPlex Installer to Install Software

Software Package	Installation Requirements
Solstice DiskSuite or Solaris Volume Manager	A partition that uses <code>/sds</code> as the mount-point name. The partition must be at least 20 Mbytes in size.
Sun Cluster HA for NFS data service	<ul style="list-style-type: none"> ■ At least two shared disks, of the same size, that are connected to the same set of nodes. ■ Solstice DiskSuite software installed, or Solaris Volume Manager software configured, by SunPlex Installer. ■ A logical hostname for use by Sun Cluster HA for NFS. The logical hostname must have a valid IP address that is accessible by all cluster nodes. The IP address must be on the same subnet as the base hostnames of the cluster nodes. ■ A test IP address for each node of the cluster. SunPlex Installer uses these test IP addresses to create Internet Protocol (IP) Network Multipathing (IP Network Multipathing) groups for use by Sun Cluster HA for NFS.
Sun Cluster HA for Apache scalable data service	<ul style="list-style-type: none"> ■ At least two shared disks of the same size that are connected to the same set of nodes. ■ Solstice DiskSuite software installed, or Solaris Volume Manager software configured, by SunPlex Installer. ■ A shared address for use by Sun Cluster HA for Apache. The shared address must have a valid IP address that is accessible by all cluster nodes. The IP address must be on the same subnet as the base hostnames of the cluster nodes. ■ A test IP address for each node of the cluster. SunPlex Installer uses these test IP addresses to create Internet Protocol (IP) Network Multipathing (IP Network Multipathing) groups for use by Sun Cluster HA for Apache.

Test IP Addresses

The test IP addresses that you supply must meet the following requirements:

- Test IP addresses for all adapters in the same multipathing group must belong to a single IP subnet.
- Test IP addresses must not be used by normal applications because the test IP addresses are not highly available.

The following table lists each metaset name and cluster-file-system mount point that is created by SunPlex Installer. The number of metasets and mount points that SunPlex Installer creates depends on the number of shared disks that are connected to the

node. For example, if a node is connected to four shared disks, SunPlex Installer creates the `mirror-1` and `mirror-2` metaset. However, SunPlex Installer does not create the `mirror-3` metaset, because the node does not have enough shared disks to create a third metaset.

TABLE 2-5 Metasets Installed by SunPlex Installer

Shared Disks	Metaset Name	Cluster File System Mount Point	Purpose
First pair	<code>mirror-1</code>	<code>/global/mirror-1</code>	Sun Cluster HA for NFS or Sun Cluster HA for Apache scalable data service, or both
Second pair	<code>mirror-2</code>	<code>/global/mirror-2</code>	Unused
Third pair	<code>mirror-3</code>	<code>/global/mirror-3</code>	Unused

Note – If the cluster does not meet the minimum shared-disk requirement, SunPlex Installer still installs the Solstice DiskSuite packages. However, without sufficient shared disks, SunPlex Installer cannot configure the metasets, metadevices, or volumes. SunPlex Installer then cannot configure the cluster file systems that are needed to create instances of the data service.

Character-Set Limitations

SunPlex Installer recognizes a limited character set to increase security. Characters that are not a part of the set are silently filtered out when HTML forms are submitted to the SunPlex Installer server. The following characters are accepted by SunPlex Installer:

```
()+, - . / 0-9 := @ A-Z ^ _ a-z { | } ~
```

This filter can cause problems in the following two areas:

- **Password entry for Sun Java™ System services** – If the password contains unusual characters, these characters are stripped out, resulting in one of the following problems:
 - The resulting password therefore fails because it has less than eight characters.
 - The application is configured with a different password than the user expects.
- **Localization** – Alternative character sets, such as accented characters or Asian characters, do not work for input.

▼ How to Install SunPlex Manager Software

This procedure describes how to install SunPlex Manager software on your cluster.

Perform this procedure on each node of the cluster.

1. **Ensure that Solaris software and patches are installed on each node of the cluster.**
 You must install Solaris software as described in [“How to Install Solaris Software” on page 52](#). Or, if Solaris software is already installed on the node, you must ensure that the Solaris installation meets the requirements for Sun Cluster software. You must also ensure that the installation meets the requirements for any other software that you intend to install on the cluster.
2. **Review the requirements and guidelines in [“Using SunPlex Installer to Install Sun Cluster Software” on page 63](#).**
3. **x86: Determine whether you are using the Netscape Navigator™ browser or the Microsoft Internet Explorer browser on your administrative console.**
 - If you are using Netscape Navigator, go to [Step 4](#).
 - If you are using Internet Explorer, go to [Step 5](#).
4. **x86: Ensure that the Java plug-in is installed and working on your administrative console.**
 - a. **Start the Netscape Navigator browser on the administrative console that you use to connect to the cluster.**
 - b. **From the Help menu, choose About Plug-ins.**
 - c. **Determine whether the Java plug-in is listed.**
 - If yes, skip to [Step 6](#).
 - If no, proceed to [Step d](#).
 - d. **Download the latest Java plug-in from <http://java.sun.com/products/plugin>.**
 - e. **Install the plug-in on your administrative console.**
 - f. **Create a symbolic link to the plug-in.**

```

% cd ~/.netscape/plugins/
% ln -s /usr/j2se/plugin/i386/ns4/javaplugin.so .

```
 - g. **Skip to [Step 6](#).**
5. **x86: Ensure that Java 2 Platform, Standard Edition (J2SE) for Windows is installed and working on your administrative console.**
 - a. **On your Microsoft Windows desktop, click Start, point to Settings, and then select Control Panel.**
 The Control Panel window appears.
 - b. **Determine whether the Java Plug-in is listed.**
 - If no, proceed to [Step c](#).

- If yes, double-click the Java Plug-in control panel. When the control panel window opens, click the About tab.
 - If version 1.4.1 or a later version is shown, skip to [Step 6](#).
 - If an earlier version is shown, proceed to [Step c](#).
 - c. **Download the latest version of J2SE for Windows from <http://java.sun.com/j2se/downloads.html>.**
 - d. **Install the J2SE for Windows software on your administrative console.**
 - e. **Restart the system on which your administrative console runs.**
The J2SE for Windows control panel is activated.
- 6. Become superuser on a cluster node.**
- 7. Ensure that Apache software packages are installed on the node.**
- ```
pkginfo SUNWapchr SUNWapchu SUNWapchd
```
- If necessary, install any missing Apache software packages by performing the following steps.
- a. **Insert the Solaris 8 or Solaris 9 Software 2 of 2 CD-ROM into the CD-ROM drive of the node.**  
If the volume management daemon `vold(1M)` is running and is configured to manage CD-ROM devices, the daemon automatically mounts the CD-ROM on the `/cdrom/cdrom0/` directory.
  - b. **Change to the `Product/` directory.**
    - For Solaris 8, change to the `/cdrom/sol_8_sparc/Solaris_8/Product/` directory.  

```
cd /cdrom/sol_8_sparc/Solaris_8/Product/
```
    - For Solaris 9, change to the `/cdrom/cdrom0/Solaris_9/Product/` directory.  

```
cd /cdrom/cdrom0/Solaris_9/Product/
```
  - c. **Install the Apache software packages in the order that is shown in this step.**  

```
pkgadd -d . SUNWapchr SUNWapchu SUNWapchd
```
  - d. **Change to a directory that does *not* reside on the CD-ROM and eject the CD-ROM.**  

```
eject cdrom
```
  - e. **Install any Apache software patches.**  
See “Patches and Required Firmware Levels” in *Sun Cluster 3.1 Release Notes* for the location of patches and installation instructions.

8. If not already installed, install the Java Dynamic Management Kit (JDMK) packages.

These packages are required by Sun Cluster software.

- a. Insert the Sun Cluster 3.1 9/04 CD-ROM.
- b. Change to the `Solaris_arch/Product/sun_cluster/Solaris_ver/Packages/` directory, where `arch` is `sparc` or `x86`, and where `ver` is 8 (for Solaris 8) or 9 (for Solaris 9) .

```
phys-schost-1# cd Solaris_arch/Product/sun_cluster/Solaris_ver/Packages/
```

- c. Install the JDMK packages.

```
phys-schost-1# pkgadd -d . SUNWjdmk*
```

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

```
eject cdrom
```

9. If not already installed, install the security files for the common agent container.

Perform the following steps to ensure that the common agent container security files are identical on all cluster nodes and that the copied files retain the correct file permissions. These files are required by Sun Cluster software.

- a. On all cluster nodes, stop the security file agent for the common agent container.

```
/opt/SUNWcacao/bin/cacaoadm stop
```

- b. On one node of the cluster, insert the Sun Cluster 3.1 9/04 CD-ROM.

- c. Change to the `Solaris_arch/Product/sun_cluster/Solaris_ver/Packages/` directory, where `arch` is `sparc` or `x86`, and where `ver` is 8 (for Solaris 8) or 9 (for Solaris 9) .

```
phys-schost-1# cd Solaris_arch/Product/sun_cluster/Solaris_ver/Packages/
```

- d. Install the common agent container packages.

```
phys-schost-1# pkgadd -d . SUNWcacao*
```

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

```
eject cdrom
```

- f. Change to the `/etc/opt/SUNWcacao/` directory.

```
phys-schost-1# cd /etc/opt/SUNWcacao/
```

- g. Create a tar file of the `/etc/opt/SUNWcacao/security/` directory.

```
phys-schost-1# tar cf /tmp/SECURITY.tar security
```

- h. Copy the `/tmp/SECURITY.tar` file to each of the other cluster nodes.
- i. On each node to which you copied the `/tmp/SECURITY.tar` file, extract the security files.

Any security files that already exist in the `/etc/opt/SUNWcacao/` directory are overwritten.

```
phys-schost-2# cd /etc/opt/SUNWcacao/
phys-schost-2# tar xf /tmp/SECURITY.tar
```

- j. Delete the `/tmp/SECURITY.tar` file from each node in the cluster.  
You must delete each copy of the tar file to avoid security risks.

```
phys-schost-1# rm /tmp/SECURITY.tar
phys-schost-2# rm /tmp/SECURITY.tar
```

- k. On all nodes, restart the security file agent.

```
phys-schost-1# /opt/SUNWcacao/bin/cacaoadm start
```

## 10. Install Sun Web Console packages.

These packages are required by Sun Cluster software, even if you do not use Sun Web Console.

- a. Insert the Sun Cluster 3.1 9/04 CD-ROM in the CD-ROM drive.
- b. Change to the `/cdrom/cdrom0/Solaris_arch/Product/sun_web_console/2.1/` directory, where *arch* is `sparc` or `x86`.
- c. Run the `setup` command.

```
./setup
```

The `setup` command installs all packages to support Sun Web Console.

## 11. Install the SunPlex Manager software packages.

- a. Insert the Sun Cluster 3.1 9/04 CD-ROM into the CD-ROM drive of the node.
- b. Change to the `Solaris_arch/Product/sun_cluster/Solaris_ver/Packages/` directory, where *arch* is `sparc` or `x86`, and where *ver* is 8 (for Solaris 8) or 9 (for Solaris 9) .

```
cd /cdrom/cdrom0/Solaris_arch/Product/sun_cluster/Solaris_ver/Packages/
```

- c. Install the SunPlex Manager software packages.

```
pkgadd -d . SUNWscva SUNWscspm SUNWscspmu SUNWscspmr
```

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

```
eject cdrom
```

## 12. Use SunPlex Installer to install and configure Sun Cluster software.

Go to [“How to Install and Configure Sun Cluster Software \(SunPlex Installer\)”](#) on page 70.

# ▼ How to Install and Configure Sun Cluster Software (SunPlex Installer)

---

**Note** – To add a new node to an existing cluster, instead follow the procedures in [“How to Configure Sun Cluster Software on Additional Cluster Nodes \(scinstall\)”](#) on page 90.

---

Perform this procedure to use SunPlex Installer to install and configure Sun Cluster software and patches on all nodes in the cluster in a single operation. In addition, you can use this procedure to install Solstice DiskSuite software and patches (Solaris 8) or to configure Solaris Volume Manager mirrored disk sets (Solaris 9).

If you use SunPlex Installer to install Solstice DiskSuite software or to configure Solaris Volume Manager disk sets, you can also install one or both of the following data services:

- Sun Cluster HA for NFS data service
- Scalable Sun Cluster HA for Apache data service

The installation process might take from 30 minutes to two or more hours. The actual length of time depends on the number of nodes that are in the cluster, your choice of data services to install, and the number of disks that are in your cluster configuration.

### 1. Ensure that the cluster configuration meets the requirements to use SunPlex Installer to install software.

See [“Using SunPlex Installer to Install Sun Cluster Software”](#) on page 63 for installation requirements and restrictions.

### 2. Ensure that the root password is the same on every node of the cluster.

To use the `root` password to access SunPlex Installer or SunPlex Manager, the `root` password must be the same on all nodes in the cluster.

If some nodes have a different `root` password than other nodes, set the `root` password to the same value on each node of the cluster. If necessary, also use the `chkey` command to update the RPC key pair. See the `chkey(1)` man page.

```
passwd
 Enter new password
chkey -p
```

3. **If you intend to install Sun Cluster HA for NFS or Sun Cluster HA for Apache, ensure that the cluster configuration meets all applicable requirements.**  
See “Using SunPlex Installer to Install Sun Cluster Software” on page 63.
4. **Ensure that SunPlex Manager software is installed on each node of the cluster.**  
See the installation procedures in “How to Install SunPlex Manager Software” on page 65.
5. **Prepare file-system paths to a CD-ROM image of each software product that you intend to install.**  
Follow these guidelines to prepare the file-system paths:
  - Provide each CD-ROM image in a location that is available to each node.
  - Ensure that the CD-ROM images are accessible to all nodes of the cluster from the same file-system path. These paths can be one or more of the following locations:
    - CD-ROM drives that are exported to the network from machines outside the cluster.
    - Exported file systems on machines outside the cluster.
    - CD-ROM images that are copied to local file systems on each node of the cluster. The local file system must use the same name on each node.
6. **Install additional packages if you intend to use one or more of the following features.**
  - Remote Shared Memory Application Programming Interface (RSMAPI)
  - SCI-PCI adapters for the interconnect transport
  - RSMRDT drivers

---

**Note** – Use of the RSMRDT driver is restricted to clusters that run an Oracle9i release 2 SCI configuration with RSM enabled. Refer to Oracle9i release 2 user documentation for detailed installation and configuration instructions.

---

**a. Determine which packages you must install.**

The following table lists the Sun Cluster 3.1 9/04 packages that each feature requires and the order in which you must install each group of packages. SunPlex Installer does not automatically install these packages.

| Feature          | Additional Sun Cluster 3.1 9/04 Packages to Install |
|------------------|-----------------------------------------------------|
| RSMAPI           | SUNWscrif                                           |
| SCI-PCI adapters | SUNWsci SUNWscid SUNWscidx                          |

| Feature        | Additional Sun Cluster 3.1 9/04 Packages to Install |
|----------------|-----------------------------------------------------|
| RSMRDT drivers | SUNWscrdt                                           |

- b. **Ensure that any dependency Solaris packages are already installed.**  
See Step 8 in “How to Install Solaris Software” on page 52.
  - c. **Insert the Sun Cluster 3.1 9/04 CD-ROM into the CD-ROM drive of a node.**
  - d. **Change to the**  
`Solaris_arch/Product/sun_cluster/Solaris_ver/Packages/`  
**directory, where *arch* is `sparc` or `x86`, and where *ver* is 8 (for Solaris 8) or 9 (for Solaris 9) .**  
  
`# cd /cdrom/cdrom0/Solaris_arch/Product/sun_cluster/Solaris_ver/Packages/`
  - e. **Install the additional packages.**  
  
`# pkgadd -d . packages`
  - f. **Change to a directory that does *not* reside on the CD-ROM and eject the CD-ROM.**  
  
`# eject cdrom`
  - g. **Repeat for each additional node in the cluster.**
7. **If patches exist that are required to support Sun Cluster or Solstice DiskSuite software, determine how to install those patches.**
    - To manually install patches, use the `pkgadd` command to install all patches *before* you use SunPlex Installer.
    - To use SunPlex Installer to install patches, copy patches into a single directory. Ensure that the patch directory meets the following requirements:
      - The patch directory resides on a file system that is available to each node.
      - Only one version of each patch is present in this patch directory.  
If the patch directory contains multiple versions of the same patch, SunPlex Installer cannot determine the correct patch dependency order.
      - The patches are uncompressed.
  8. **Complete the following installation worksheet.**



**TABLE 2-6** SunPlex Installer Installation and Configuration Worksheet

| Component                                                                                            | Description/Example                                                                                                                                                              | Enter Answers Here |
|------------------------------------------------------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|--------------------|
| Cluster Name                                                                                         | What is the name of the cluster that you want to establish?                                                                                                                      |                    |
|                                                                                                      | How many nodes are you installing in the cluster?                                                                                                                                |                    |
| Node Names                                                                                           | What are the names of the cluster nodes?                                                                                                                                         |                    |
| Cluster-Transport Adapters and Cables                                                                | What are the names of the two transport adapters to use, two adapters per node?                                                                                                  |                    |
| Solstice DiskSuite or Solaris Volume Manager                                                         | <ul style="list-style-type: none"> <li>■ Solaris 8: Do you want to install Solstice DiskSuite?</li> <li>■ Solaris 9: Do you want to configure Solaris Volume Manager?</li> </ul> | Yes   No           |
| Sun Cluster HA for NFS<br><i>Requires Solstice DiskSuite or Solaris Volume Manager</i>               | Do you want to install Sun Cluster HA for NFS?                                                                                                                                   | Yes   No           |
|                                                                                                      | If yes, also specify the following:                                                                                                                                              |                    |
|                                                                                                      | What is the logical hostname that the data service is to use?                                                                                                                    |                    |
|                                                                                                      | What are the test IP addresses to use?<br><i>Supply one test IP address for each node in the cluster.</i>                                                                        |                    |
| Sun Cluster HA for Apache (scalable)<br><i>Requires Solstice DiskSuite or Solaris Volume Manager</i> | Do you want to install scalable Sun Cluster HA for Apache?                                                                                                                       | Yes   No           |
|                                                                                                      | If yes, also specify the following:                                                                                                                                              |                    |
|                                                                                                      | What is the logical hostname that the data service is to use?                                                                                                                    |                    |
|                                                                                                      | What are the test IP addresses to use?<br><i>Supply one test IP address for each node in the cluster.</i>                                                                        |                    |
| CD-ROM Paths                                                                                         | What is the path for each of the following components that you want to install?<br><br><i>The CD-ROM path must end with the directory that contains the .cdtoc file.</i>         |                    |
|                                                                                                      | Solstice DiskSuite:                                                                                                                                                              |                    |
|                                                                                                      | Sun Cluster (framework):                                                                                                                                                         |                    |
|                                                                                                      | Sun Cluster data services (agents):                                                                                                                                              |                    |
|                                                                                                      | Patches:                                                                                                                                                                         |                    |
| Validation Checks                                                                                    | Do you want to run the <code>sccheck</code> utility to validate the cluster?                                                                                                     | Yes   No           |

---

**Note** – SunPlex Installer installation automatically specifies the default private-network address (172.16.0.0) and netmask (255.255.0.0). If you need to use a different address, do not use SunPlex Installer to install Sun Cluster software. Instead, follow procedures in “How to Install Sun Cluster Software Packages” on page 56 and in “How to Configure Sun Cluster Software on All Nodes (scinstall)” on page 57 to install and configure the cluster.

You cannot change the private-network address and netmask after `scinstall` processing has finished. If you need to use a different private-network address or netmask and the node is still in installation mode, follow the procedures in “How to Uninstall Sun Cluster Software to Correct Installation Problems” on page 123. Then repeat this procedure to reinstall and configure the node with the correct information.

---

See “Planning the Solaris OS” on page 16 and “Planning the Sun Cluster Environment” on page 21 for planning guidelines. See the *Sun Cluster Data Service Planning and Administration Guide for Solaris OS* for data-service planning guidelines.

## 9. Start SunPlex Installer.

- a. **From the administrative console or any other machine outside the cluster, launch a browser.**
- b. **Disable the browser’s Web proxy.**  
SunPlex Installer installation functionality is incompatible with Web proxies.
- c. **Ensure that disk caching and memory caching is enabled.**  
The disk cache and memory cache size must be greater than 0.
- d. **From the browser, connect to port 3000 on a node of the cluster.**

`https://node:3000`

The Sun Cluster Installation screen is displayed in the browser window.

---

**Note** – If SunPlex Installer displays the data services installation screen instead of the Sun Cluster Installation screen, Sun Cluster framework software is already installed and configured on that node. Check that the name of the node in the URL is the correct name of the cluster node to install.

---

- e. **If the browser displays a New Site Certification window, follow the onscreen instructions to accept the certificate.**

## 10. Log in as superuser.

**11. In the Sun Cluster Installation screen, verify that the cluster meets the listed requirements for using SunPlex Installer.**

If you meet all listed requirements, click Next to continue to the next screen.

**12. Follow the menu prompts to supply your answers from the worksheet that you completed in [Step 8](#).**

**13. Click [Begin Installation](#) to start the installation process.**

Follow these guidelines:

- Do *not* close the browser window nor change the URL during the installation process.
- If the browser displays a New Site Certification window, follow the onscreen instructions to accept the certificate.
- If the browser prompts for login information, type the appropriate superuser ID and password for the node that you connect to.

SunPlex Installer installs and configures all cluster nodes and reboots the cluster. The cluster is established when all nodes have successfully booted into the cluster. Sun Cluster installation output is logged in a `/var/cluster/logs/install/scinstall.log.N` file.

During installation, the screen displays brief messages about the status of the cluster installation process. When installation and configuration is complete, the browser displays the cluster monitoring and administration GUI.

SunPlex Installer installation output is logged in the `/var/cluster/spm/messages` file. Sun Cluster installation output is logged in a `/var/cluster/logs/install/scinstall.log.N` file.

**14. Verify the quorum assignments and modify those assignments, if necessary.**

For clusters with three or more nodes, the use of shared quorum devices is optional. SunPlex Installer might or might not have assigned quorum votes to any quorum devices, depending on whether appropriate shared disks were available. You can use SunPlex Manager to designate quorum devices and to reassign quorum votes in the cluster. See “Administering Quorum” in *Sun Cluster System Administration Guide for Solaris OS* for more information.

**15. Install Sun StorEdge QFS file system software.**

Follow the procedures for initial installation in the *Sun StorEdge QFS and Sun StorEdge SAM-FS Software Installation and Configuration Guide*.

**16. (Optional) SPARC: To install VERITAS File System, go to [“SPARC: How to Install VERITAS File System Software”](#) on page 97.**

**17. Set up the name-service look-up order.**

Go to [“How to Configure the Name-Service Switch”](#) on page 98.

## ▼ How to Install Solaris and Sun Cluster Software (JumpStart)

This procedure describes how to set up and use the `scinstall(1M)` custom JumpStart installation method. This method installs both Solaris OS and Sun Cluster software on all cluster nodes in the same operation and establishes the cluster. You can also use this procedure to add new nodes to an existing cluster.

1. **Ensure that the hardware setup is complete and that connections are verified before you install Solaris software.**

See the *Sun Cluster Hardware Administration Collection* and your server and storage device documentation for details on how to set up the hardware.

2. **Ensure that your cluster configuration planning is complete.**

See “How to Prepare for Cluster Software Installation” on page 46 for requirements and guidelines.

3. **If you use a naming service, add the following information to any naming services that clients use to access cluster services.**

- Address-to-name mappings for all public hostnames and logical addresses
- The IP address and hostname of the JumpStart server

See “IP Addresses” on page 22 for planning guidelines. See your Solaris system-administrator documentation for information about using Solaris naming services.

4. **If you are installing a new node to an existing cluster, add the node to the list of authorized cluster nodes.**

- a. **Run `scsetup(1M)` from another cluster node that is active.**
- b. **Use the `scsetup` utility to add the new node’s name to the list of authorized cluster nodes.**

For more information, see “How to Add a Cluster Node to the Authorized Node List” in “Adding and Removing a Cluster Node” in *Sun Cluster System Administration Guide for Solaris OS*.

5. **Set up your JumpStart installation server.**

- Follow the appropriate instructions for your software platform.

| Solaris OS platform | Procedures                                                                                                                                                                                                                                                                                                                                                                           |
|---------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| SPARC               | See "Creating a Profile Server for Networked Systems" in <i>Solaris 8 Advanced Installation Guide</i> or "Creating a Profile Server for Networked Systems" in <i>Solaris 9 Installation Guide</i> for instructions about how to set up a JumpStart installation server.<br><br>See also the <code>setup_install_server(1M)</code> and <code>add_install_client(1M)</code> man pages. |
| x86                 | See "Solaris 9 Software Installation From a PXE Server" in <i>Sun Fire V60x and Sun Fire V65x Server Solaris Operating Environment Installation Guide</i> for instructions about how to set up a JumpStart Dynamic Host Configuration Protocol (DHCP) server and a Solaris network for Preboot eXecution Environment (PXE) installations.                                            |

- Ensure that the JumpStart installation server meets the following requirements.
    - The installation server is on the same subnet as the cluster nodes.
    - The installation server is not itself a cluster node.
    - The installation server installs a release of the Solaris OS that is supported by the Sun Cluster software.
    - A custom JumpStart directory exists for JumpStart installation of Sun Cluster software. This *jumpstart-dir* directory must contain a copy of the `check(1M)` utility. The directory must also be NFS exported for reading by the JumpStart installation server.
    - Each new cluster node is configured as a custom JumpStart install client that uses the custom JumpStart directory that you set up for Sun Cluster installation.
6. On a cluster node or another machine of the same server platform, prepare a flash archive of the Solaris OS and Sun Web Console software.
- a. Install the Solaris OS as described in "How to Install Solaris Software" on page 52.
  - b. Insert the Sun Cluster 3.1 9/04 CD-ROM in the CD-ROM drive.
  - c. Change to the  
`/cdrom/cdrom0/Solaris_arch/Product/sun_web_console/2.1/`  
 directory, where *arch* is `sparc` or `x86`.
  - d. Run the `setup` command.
 

```
./setup
```

 The `setup` command installs all packages to support Sun Web Console.
  - e. Change to a directory that does *not* reside on the CD-ROM and eject the CD-ROM.
 

```
eject cdrom
```

**f. Create the flash archive of the installed system.**

```
flar create -n name archive
```

`-n name` Name to give the flash archive.

`archive` File name to give the flash archive, with the full path. By convention, the file name ends in `.flar`.

Follow procedures in “Creating Web Start Flash Archives” in *Solaris 8 Advanced Installation Guide* or “Creating Solaris Flash Archives (Tasks)” in *Solaris 9 9/04 Installation Guide*.

**7. Copy the flash archive to the JumpStart installation server.**

**8. Ensure that the flash archive on the JumpStart installation server is NFS exported for reading by the JumpStart installation server.**

See “Solaris NFS Environment” in *System Administration Guide, Volume 3* or “Managing Network File Systems (Overview)” in *System Administration Guide: Resource Management and Network Services* for more information about automatic file sharing. See also the `share(1M)` and `dfstab(4)` man pages.

**9. Create a directory on the JumpStart installation server to hold your copy of the Sun Cluster 3.1 9/04 CD-ROM.**

In the following example, the `/export/suncluster/` directory is created for this purpose.

```
mkdir -m 755 /export/suncluster/
```

**10. Copy the Sun Cluster CD-ROM to the JumpStart installation server.**

**a. Insert the Sun Cluster 3.1 9/04 CD-ROM into the CD-ROM drive on the JumpStart installation server.**

If the volume management daemon `vold(1M)` is running and is configured to manage CD-ROM devices, the daemon automatically mounts the CD-ROM on the `/cdrom/cdrom0/` directory.

**b. Change to the**

`Solaris_arch/Product/sun_cluster/Solaris_ver/Tools/` directory, where `arch` is `sparc` or `x86` and where `ver` is 8 (for Solaris 8) or 9 (for Solaris 9).

```
cd /cdrom/cdrom0/Solaris_arch/Product/sun_cluster/Solaris_ver/Tools/
```

**c. Copy the CD-ROM to a new directory on the JumpStart installation server.**

The `scinstall` command creates the new installation directory when the command copies the CD-ROM files. The following example uses the installation directory name `/export/suncluster/sc31/`.

```
./scinstall -a /export/suncluster/sc31/
```

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

```
eject cdrom
```

- 11. Ensure that the Sun Cluster 3.1 9/04 CD-ROM image on the JumpStart installation server is NFS exported for reading by the JumpStart installation server.

See “Solaris NFS Environment” in *System Administration Guide, Volume 3* or “Managing Network File Systems (Overview)” in *System Administration Guide: Resource Management and Network Services* for more information about automatic file sharing. See also the `share(1M)` and `dfstab(4)` man pages.

- 12. Have available the following information:

- The Ethernet address of each cluster node
- The following completed installation worksheet

TABLE 2-7 JumpStart Installation and Configuration Worksheet

| Component                                 | Description/Example                                                                                                    | Enter Answers Here |               |
|-------------------------------------------|------------------------------------------------------------------------------------------------------------------------|--------------------|---------------|
| JumpStart Directory                       | What is the name of the JumpStart directory to use?                                                                    |                    |               |
| Cluster Name                              | What is the name of the cluster that you want to establish?                                                            |                    |               |
| Cluster Nodes                             | What are the names of the cluster nodes that are planned for the initial cluster configuration?                        |                    |               |
| DES Authentication                        | Do you need to use DES authentication?                                                                                 | No   Yes           |               |
| Network Address for the Cluster Transport | Do you want to accept the default network address (172.16.0.0)?                                                        | Yes   No           |               |
|                                           | If no, supply your own network address:                                                                                | ____.____.0.0      |               |
|                                           | Do you want to accept the default netmask (255.255.0.0)?                                                               | Yes   No           |               |
|                                           | If no, supply your own netmask:                                                                                        | 255.255.____.____  |               |
| Point-to-Point Cables                     | Does this cluster use transport junctions?                                                                             | Yes   No           |               |
| Cluster-Transport Junctions               | If used, what are the names of the two transport junctions?<br>Defaults: <code>switch1</code> and <code>switch2</code> | <i>First</i>       | <i>Second</i> |

**TABLE 2-7** JumpStart Installation and Configuration Worksheet (Continued)

| Component                                                    | Description/Example                                                                                                                                                  | Enter Answers Here |               |
|--------------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------|--------------------|---------------|
| Cluster-Transport Adapters and Cables                        | First node name:                                                                                                                                                     |                    |               |
|                                                              |                                                                                                                                                                      | <i>First</i>       | <i>Second</i> |
|                                                              | Transport adapters:                                                                                                                                                  |                    |               |
|                                                              | Where does each transport adapter connect to ( <i>a transport junction or another adapter</i> )?<br>Junction defaults: <code>switch1</code> and <code>switch2</code> |                    |               |
|                                                              | For transport junctions, do you want to use the default port name?                                                                                                   | Yes   No           | Yes   No      |
|                                                              | If no, what is the name of the port that you want to use?                                                                                                            |                    |               |
| <i>Specify for each additional node</i>                      | Node name:                                                                                                                                                           |                    |               |
|                                                              |                                                                                                                                                                      | <i>First</i>       | <i>Second</i> |
|                                                              | Transport adapters:                                                                                                                                                  |                    |               |
|                                                              | Where does each transport adapter connect to ( <i>a transport junction or another adapter</i> )?<br>Junction defaults: <code>switch1</code> and <code>switch2</code> |                    |               |
|                                                              | For transport junctions, do you want to use the default port name?                                                                                                   | Yes   No           | Yes   No      |
| Global-Devices File System<br><i>(specify for each node)</i> | Do you want to use the default name of the global-devices file system ( <code>/globaldevices</code> )?                                                               | Yes   No           |               |
|                                                              | If no, do you want to use an already-existing file system?                                                                                                           | Yes   No           |               |
|                                                              | What is the name of the file system?                                                                                                                                 |                    |               |
| Software Patch Installation                                  | Do you want <code>scinstall</code> to install patches for you?                                                                                                       | Yes   No           |               |
|                                                              | If yes, what is the name of the patch directory?                                                                                                                     |                    |               |
|                                                              | Do you want to use a patch list?                                                                                                                                     | Yes   No           |               |

See “Planning the Solaris OS” on page 16 and “Planning the Sun Cluster Environment” on page 21 for planning guidelines.



---

**Note** – You cannot change the private-network address and netmask after `scinstall` processing has finished. If you need to use a different private-network address or netmask and the node is still in installation mode, follow the procedures in “[How to Uninstall Sun Cluster Software to Correct Installation Problems](#)” on page 123. Then repeat this procedure to reinstall and configure the node with the correct information.

---

**13. From the JumpStart installation server, start the `scinstall(1M)` utility.**

The path `/export/suncluster/sc31/` is used here as an example of the installation directory that you created. In the CD-ROM path, replace `arch` with `sparc` or `x86` and replace `ver` with 8 (for Solaris 8) or 9 (for Solaris 9).

```
cd /export/suncluster/sc31/Solaris_arch/Product/sun_cluster/ \
Solaris_ver/Tools/
./scinstall
```

**14. Follow these guidelines to use the interactive `scinstall` utility:**

- Interactive `scinstall` enables you to type ahead. Therefore, do not press the Return key more than once if the next menu screen does not appear immediately.
- Unless otherwise noted, you can press Control-D to return to either the start of a series of related questions or to the Main Menu.
- Default answers or answers to previous sessions are displayed in brackets ([ ]) at the end of a question. Press Return to enter the response that is in brackets without typing it.

**15. From the Main Menu, choose the menu item, Configure a cluster to be JumpStarted from this installation server.**

This option is used to configure custom JumpStart finish scripts. JumpStart uses these finish scripts to install the Sun Cluster software.

```
*** Main Menu ***
```

```
Please select from one of the following (*) options:
```

- \* 1) Install a cluster or cluster node
- \* 2) Configure a cluster to be JumpStarted from this install server
- 3) Add support for new data services to this cluster node
- \* 4) Print release information for this cluster node
  
- \* ?) Help with menu options
- \* q) Quit

```
Option: 2
```

---

**Note** – If the JumpStart option does not have an asterisk in front, the option is disabled. This condition indicates that JumpStart setup is not complete or that the setup has an error. To correct this condition, do the following:

- a. Quit the `scinstall` utility.
  - b. Repeat [Step 5](#) through [Step 10](#) to correct JumpStart setup.
  - c. Restart the `scinstall` utility.
- 

**16. Follow the menu prompts to supply your answers from the worksheet that you completed in [Step 12](#).**

The `scinstall` command stores your configuration information and creates the following `autoscinstall.class` default class file in the `jumpstart-dir/autoscinstall.d/3.1/` directory.

```
install_type initial_install
system_type standalone
partitioning explicit
filesystems rootdisk.s0 free /
 rootdisk.s1 750 swap
filesystems rootdisk.s3 512 /globaldevices
filesystems rootdisk.s7 20
cluster SUNWCuser add
package SUNWman add
```

**17. Make adjustments to the default `autoscinstall.class` file to configure JumpStart to install the flash archive.**

- a. Change the following entries in the `autoscinstall.class` file. In the last new entry in the table, *archive* represents the location of the flash archive file.

| Existing Entry            |                              | New Entry                     |                            |
|---------------------------|------------------------------|-------------------------------|----------------------------|
| <code>install_type</code> | <code>initial_install</code> | <code>install_type</code>     | <code>flash_install</code> |
| <code>system_type</code>  | <code>standalone</code>      | <code>archive_location</code> | <i>archive</i>             |

- b. Remove all entries that would install a specific package.

```
cluster SUNWCuser add
package SUNWman add
```

**18. Set up Solaris patch directories.**

---

**Note** – If you specified a patch directory to the `scinstall` utility, patches that are located in Solaris patch directories are not installed.

---

- a. **Create `jumpstart-dir/autoscinstall.d/nodes/node/patches/` directories on the JumpStart installation server.**

Create one directory for each node in the cluster, where *node* is the name of a cluster node. Alternately, use this naming convention to create symbolic links to a shared patch directory.

```
mkdir jumpstart-dir/autoscinstall.d/nodes/node/patches/
```

- b. **Place copies of any Solaris patches into each of these directories.**
- c. **Place copies of any hardware-related patches that you must install after Solaris software is installed into each of these directories.**

#### 19. Set up files to contain the necessary hostname information locally on each node.

- a. **On the JumpStart installation server, create files that are named `jumpstart-dir/autoscinstall.d/nodes/node/archive/etc/inet/hosts`.**

Create one file for each node, where *node* is the name of a cluster node. Alternately, use this naming convention to create symbolic links to a shared `hosts` file.

- b. **Add the following entries into each file.**
  - IP address and hostname of the NFS server that holds a copy of the Sun Cluster CD-ROM image. The NFS server could be the JumpStart installation server or another machine.
  - IP address and hostname of each node in the cluster.

#### 20. If you are using a cluster administrative console, display a console screen for each node in the cluster.

- If Cluster Control Panel (CCP) software is installed and configured on your administrative console, you can use the `cconsole(1M)` utility to display the individual console screens. The `cconsole` utility also opens a master window from which you can send your input to all individual console windows at the same time. Use the following command to start `cconsole`:

```
/opt/SUNWcluster/bin/cconsole clustername &
```

- If you do not use the `cconsole` utility, connect to the consoles of each node individually.

#### 21. Shut down each node.

```
shutdown -g0 -y -i0
```

#### 22. Boot each node to start the JumpStart installation.

- On SPARC based systems, do the following:

```
ok boot net - install
```

---

**Note** – Surround the dash (-) in the command with a space on each side.

---

- On x86 based systems, do the following:
  - a. When the BIOS information screen appears, press the Esc key.  
The Select Boot Device screen appears.
  - b. On the Select Boot Device screen, choose the listed IBA that is connected to the same network as the JumpStart DHCP installation server.  
The lowest number to the right of the IBA boot choices corresponds to the lower Ethernet port number. The higher number to the right of the IBA boot choices corresponds to the higher Ethernet port number.  
The node reboots and the Device Configuration Assistant appears.
  - c. On the Boot Solaris screen, choose Net.
  - d. At the following prompt, choose Custom JumpStart and press Enter:  
Select the type of installation you want to perform:  

```
1 Solaris Interactive
2 Custom JumpStart
```

  
Enter the number of your choice followed by the <ENTER> key.  
If you enter anything else, or if you wait for 30 seconds, an interactive installation will be started.
  - e. When prompted, answer the questions and follow the instructions on the screen.

JumpStart installs the Solaris OS and Sun Cluster software on each node.

---

**Note** – Unless you have installed your own `/etc/inet/ntp.conf` file, the `scinstall` command installs a default `ntp.conf` file for you. The default file is shipped with references to the maximum number of nodes. Therefore, the `xntpd(1M)` daemon might issue error messages regarding some of these references at boot time.

You can safely ignore these messages. See [“How to Configure Network Time Protocol \(NTP\)” on page 116](#) for information on how to suppress these messages under otherwise normal cluster conditions.

---

When the installation is successfully completed, each node is fully installed as a new cluster node. Sun Cluster installation output is logged in a `/var/cluster/logs/install/scinstall.log.N` file.

You cannot change the private-network address and netmask after `scinstall` processing has finished. If you need to use a different private-network address or

netmask and the node is still in installation mode, follow the procedures in “How to Uninstall Sun Cluster Software to Correct Installation Problems” on page 123. Then repeat this procedure to reinstall and configure the node with the correct information.

**23. If you are installing a new node to an existing cluster, create mount points on the new node for all existing cluster file systems.**

- a. From another cluster node that is active, display the names of all cluster file systems.

```
% mount | grep global | egrep -v node@ | awk '{print $1}'
```

- b. On the node that you added to the cluster, create a mount point for each cluster file system in the cluster.

```
% mkdir -p mountpoint
```

For example, if a file-system name that is returned by the mount command is `/global/dg-schost-1`, run `mkdir -p /global/dg-schost-1` on the node that is being added to the cluster.

---

**Note** – The mount points become active after you reboot the cluster in [Step 26](#).

---

- c. If VERITAS Volume Manager (VxVM) is installed on any nodes that are already in the cluster, view the `vxio` number on each VxVM-installed node.

```
grep vxio /etc/name_to_major
vxio NNN
```

- Ensure that the same `vxio` number is used on each of the VxVM-installed nodes.
- Ensure that the `vxio` number is available for use on each of the nodes that do not have VxVM installed.
- If the `vxio` number is already in use on a node that does not have VxVM installed, free the number on that node. Change the `/etc/name_to_major` entry to use a different number.

**24. (Optional) To use dynamic reconfiguration on Sun Enterprise 10000 servers, add the following entry to the `/etc/system` file. Add this entry on each node in the cluster.**

```
set kernel_cage_enable=1
```

This entry becomes effective after the next system reboot. See the *Sun Cluster System Administration Guide for Solaris OS* for procedures to perform dynamic reconfiguration tasks in a Sun Cluster configuration. See your server documentation for more information about dynamic reconfiguration.

**25. x86: Set the default boot file to `kadb`.**

```
eeprom boot-file=kadb
```

The setting of this value enables you to reboot the node if you are unable to access a login prompt.

**26. If you performed a task that requires a cluster reboot, follow these steps to perform a reconfiguration reboot of the cluster.**

The following are some of the tasks that require a reboot:

- Adding a new node to an existing cluster
- Installing patches that require a node or cluster reboot
- Making configuration changes that require a reboot to become active

**a. From one node, shut down the cluster.**

```
scshutdown
```

---

**Note** – Do not reboot the first-installed node of the cluster until *after* the cluster is shut down. Until cluster installation mode is disabled, only the first-installed node, which established the cluster, has a quorum vote. In an established cluster that is still in installation mode, if the cluster is not shut down before the first-installed node is rebooted, the remaining cluster nodes cannot obtain quorum. The entire cluster then shuts down.

Cluster nodes remain in installation mode until the first time that you run the `scsetup(1M)` command. You run this command during the procedure “[How to Perform Postinstallation Setup and Configure Quorum Devices](#)” on page 104.

---

**b. Reboot each node in the cluster.**

- On SPARC based systems, do the following:

```
ok boot
```

- On x86 based systems, do the following:

```
<<< Current Boot Parameters >>>
Boot path: /pci@0,0/pci-ide@7,1/ata@1/cmdk@0,0:b
Boot args:

Type b [file-name] [boot-flags] <ENTER> to boot with options
or i <ENTER> to enter boot interpreter
or <ENTER> to boot with defaults
```

```
<<< timeout in 5 seconds >>>
Select (b)oot or (i)nterpreter: b
```

The `scinstall` utility installs and configures all cluster nodes and reboots the cluster. The cluster is established when all nodes have successfully booted into the cluster. Sun Cluster installation output is logged in a `/var/cluster/logs/install/scinstall.log.N` file.

**27. Install Sun StorEdge QFS file system software.**

Follow the procedures for initial installation in the *Sun StorEdge QFS and Sun StorEdge SAM-FS Software Installation and Configuration Guide*.

**28. SPARC: To install VERITAS File System, go to “SPARC: How to Install VERITAS File System Software” on page 97.**

**29. Set up the name-service look-up order.**

Go to “How to Configure the Name-Service Switch” on page 98.

## How to Install Sun Cluster Software on a Single-Node Cluster

Perform this task to install Sun Cluster software and establish the cluster on a single node by using the `scinstall` command. See the `scinstall(1M)` man page for details.

---

**Note** – You cannot use SunPlex Installer or the interactive form of the `scinstall` utility to install Sun Cluster software on a single-node cluster.

---

The `scinstall -iFo` command establishes the following defaults during installation:

- Creates the required device IDs
- Uses the default `/globaldevices` file system, unless the `global-devices` file system is specified with the `-G` option
- Establishes a default cluster name, which is the same as the name of the node being installed, unless the `clustername` is specified with the `-C` option

Some steps that are required for multinode cluster installations are not necessary for single-node cluster installations. When you install a single-node cluster, you do not need to perform the following steps:

- You do not need to configure quorum.
- You do not need to configure interconnect adapters or junctions.

---

**Tip** – If you anticipate eventually adding a second node to your cluster, you can configure the transport interconnect during initial cluster installation. The transport interconnect is then available for later use. See the `scinstall(1M)` man page for details.

You can later expand a single-node cluster into a multinode cluster by following the appropriate procedures provided in [“How to Configure Sun Cluster Software on Additional Cluster Nodes \(scinstall\)”](#) on page 90.

---

**1. Ensure that the Solaris OS is installed to support Sun Cluster software.**

If Solaris software is already installed on the node, you must ensure that the Solaris installation meets the requirements for Sun Cluster software and any other software that you intend to install on the cluster. See [“How to Install Solaris Software”](#) on page 52 for more information about installing Solaris software to meet Sun Cluster software requirements.

**2. Become superuser on the cluster node to install.**

**3. Install Sun Web Console packages.**

These packages are required by Sun Cluster software, even if you do not use Sun Web Console.

**a. Insert the Sun Cluster 3.1 9/04 CD-ROM in the CD-ROM drive.**

**b. Change to the**  
`/cdrom/cdrom0/Solaris_arch/Product/sun_web_console/2.1/`  
**directory, where *arch* is `sparc` or `x86`.**

**c. Run the `setup` command.**

```
./setup
```

The `setup` command installs all packages to support Sun Web Console.

**4. On the Sun Cluster 3.1 9/04 CD-ROM, change to the**  
`Solaris_arch/Product/sun_cluster/Solaris_ver/Tools/` **directory, where**  
*arch* is `sparc` or `x86` and where *ver* is 8 (for Solaris 8) or 9 (for Solaris 9) .

```
cd /cdrom/cdrom0/Solaris_arch/Product/sun_cluster/Solaris_ver/Tools/
```

**5. Install the Sun Cluster software and necessary patches by using the `scinstall` command.**

```
./scinstall -iFo [-M patchdir=dirname]
```

**-i**

Specifies the install form of the `scinstall` command. The `scinstall` command installs Sun Cluster software and initializes the node as a new cluster.

**-F**

Establishes the node as the first node in a new cluster. All `-F` options can be used when installing a single-node cluster.



`-o`  
Specifies that only one node is being installed for a single-node cluster. The `-o` option is only legal when used with both the `-i` and the `-F` forms of the command. When the `-o` option is used, cluster installation mode is preset to the disabled state.

`-M patchdir=dirname [, patchlistfile=filename]`  
Specifies the path to patch information so that the specified patches can be installed by using the `scinstall` command. If you do not specify a patch-list file, the `scinstall` command installs all the patches in the directory *dirname*. This includes tarred, jarred, and zipped patches.

The `-M` option is not required with the `scinstall -iFo` command. The `-M` option is shown in this procedure because the use of this option is the most efficient method of installing patches during a single-node cluster installation. However, you can use any method that you prefer to install patches.

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

```
eject cdrom
```

**7. Reboot the node.**

This reboot after Sun Cluster software installation establishes the node as the cluster.

**8. (Optional) Change the cluster name.**

A single-node cluster is created with the same name as the cluster node. If you prefer, you can change the cluster name. Use either the `scsetup` utility or use the following `scconf` command:

```
/usr/cluster/bin/scconf -c -C cluster=newclustername
```

**9. Verify the installation by using the `scstat` command.**

```
/usr/cluster/bin/scstat -n
```

The command output should list the cluster node with the status of `Online`. See the `scstat(1M)` man page for details.

**10. Ensure that cluster installation mode is disabled.**

```
/usr/cluster/bin/scconf -pv | grep "install mode"
```

**11. (Optional) SPARC: To install VERITAS File System, go to [“SPARC: How to Install VERITAS File System Software”](#) on page 97.**

**12. Set up the name-service look-up order.**

Go to [“How to Configure the Name-Service Switch”](#) on page 98.

## Example—Installing Sun Cluster Software on a Single-Node Cluster

The following example shows how to use the `scinstall` and `scstat` commands to install and verify a single-node cluster. The example includes installation of all patches. See the `scinstall(1M)` and `scstat(1M)` man pages for details.

```
scinstall -iFo -M patchdir=/var/cluster/patches/

Checking device to use for global devices file system ... done
** Installing SunCluster 3.1 framework **
...
Installing patches ... done

Initializing cluster name to "phys-schost-1" ... done
Initializing authentication options ... done

Setting the node ID for "phys-schost-1" ... done (id=1)

Checking for global devices global file system ... done
Updating vfstab ... done

Verifying that "cluster" is set for "hosts" in nsswitch.conf ... done
Adding the "cluster" switch to "hosts" in nsswitch.conf ... done

Verifying that "cluster" is set for "netmasks" in nsswitch.conf ... done
Adding the "cluster" switch to "netmasks" in nsswitch.conf ... done

Verifying that power management is NOT configured ... done

Ensure that the EEPROM parameter "local-mac-address?" is set to "true" ... done

Ensure network routing is disabled ... done

Please reboot this machine.

reboot
scstat -n
-- Cluster Nodes --

 Node name Status
 ----- -
Cluster node: phys-schost-1 Online
sconfg -pv | grep "install mode"
Cluster install mode: disabled
```

## ▼ How to Configure Sun Cluster Software on Additional Cluster Nodes (`scinstall`)

Perform this procedure to add a new node to an existing cluster.

1. Ensure that all necessary hardware is installed.

- Ensure that the host adapter is installed on the new node.  
See the *Sun Cluster Hardware Administration Manual for Solaris OS*.
  - Verify that any existing cluster interconnects can support the new node.  
See the *Sun Cluster Hardware Administration Manual for Solaris OS*.
  - Ensure that any additional storage is installed.  
See the appropriate manual from the Sun Cluster 3.x Hardware Administration Collection.
2. **Ensure that the Solaris OS is installed to support Sun Cluster software.**  
If Solaris software is already installed on the node, you must ensure that the Solaris installation meets the requirements for Sun Cluster software and any other software that you intend to install on the cluster. See “[How to Install Solaris Software](#)” on page 52 for more information about installing Solaris software to meet Sun Cluster software requirements.
  3. **Ensure that Sun Cluster software packages are installed on the node.**  
See “[How to Install Sun Cluster Software Packages](#)” on page 56.
  4. **Complete the following configuration worksheet.**

**TABLE 2-8** Added Node Configuration Worksheet

| Component                          | Description/Example                                                                                                    | Enter Answers Here |               |
|------------------------------------|------------------------------------------------------------------------------------------------------------------------|--------------------|---------------|
| Software Patch Installation        | Do you want <code>scinstall</code> to install patches for you?                                                         | Yes   No           |               |
|                                    | If yes, what is the patch directory?                                                                                   |                    |               |
|                                    | Do you want to use a patch list?                                                                                       | Yes   No           |               |
| Sponsoring Node                    | What is the name of the sponsoring node?<br><i>Choose any node that is active in the cluster.</i>                      |                    |               |
| Cluster Name                       | What is the name of the cluster that you want the node to join?                                                        |                    |               |
| Check                              | Do you want to run the <code>sccheck</code> validation utility?                                                        | Yes   No           |               |
| Autodiscovery of Cluster Transport | Do you want to use autodiscovery to configure the cluster transport?                                                   | Yes   No           |               |
|                                    | If no, supply the following additional information:                                                                    |                    |               |
| Point-to-Point Cables              | Does the node that you are adding to the cluster make this a two-node cluster?                                         | Yes   No           |               |
|                                    | Does the cluster use transport junctions?                                                                              | Yes   No           |               |
| Cluster-Transport Junctions        | If used, what are the names of the two transport junctions?<br>Defaults: <code>switch1</code> and <code>switch2</code> | <i>First</i>       | <i>Second</i> |

**TABLE 2-8** Added Node Configuration Worksheet (Continued)

| Component                             | Description/Example                                                                                                                                                  | Enter Answers Here |               |
|---------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------|--------------------|---------------|
| Cluster-Transport Adapters and Cables | What are the names of the two transport adapters?                                                                                                                    | <i>First</i>       | <i>Second</i> |
|                                       | Where does each transport adapter connect to ( <i>a transport junction or another adapter</i> )?<br>Junction defaults: <code>switch1</code> and <code>switch2</code> |                    |               |
|                                       | For transport junctions, do you want to use the default port name?                                                                                                   | Yes   No           | Yes   No      |
|                                       | If no, what is the name of the port that you want to use?                                                                                                            |                    |               |
| Global-Devices File System            | What is the name of the global-devices file system?<br>Default: <code>/globaldevices</code>                                                                          |                    |               |
| Automatic Reboot                      | Do you want <code>scinstall</code> to automatically reboot the node after installation?                                                                              | Yes   No           |               |

See “Planning the Solaris OS” on page 16 and “Planning the Sun Cluster Environment” on page 21 for planning guidelines.

**5. If you are adding this node to a single-node cluster, determine whether two cluster interconnects already exist.**

You must have at least two cables or two adapters configured before you can add a node.

```
scconf -p | grep cable
scconf -p | grep adapter
```

- If the output shows configuration information for two cables or for two adapters, proceed to [Step 6](#).
- If the output shows no configuration information for either cables or adapters, or shows configuration information for only one cable or adapter, configure new cluster interconnects.

**a. On the existing cluster node, start the `scsetup(1M)` utility.**

```
scsetup
```

**b. Choose the menu item, Cluster interconnect.**

**c. Choose the menu item, Add a transport cable.**

Follow the instructions to specify the name of the node to add to the cluster, the name of a transport adapter, and whether to use a transport junction.

**d. If necessary, repeat [Step c](#) to configure a second cluster interconnect.**

When finished, quit the `scsetup` utility.

**e. Verify that the cluster now has two cluster interconnects configured.**

```
scconf -p | grep cable
scconf -p | grep adapter
```

The command output should show configuration information for at least two cluster interconnects.

**6. If you are adding this node to an existing cluster, add the new node to the cluster authorized-nodes list.**

**a. On any active cluster member, start the `scsetup(1M)` utility.**

```
scsetup
```

The Main Menu is displayed.

**b. Choose the menu item, New nodes.**

**c. Choose the menu item, Specify the name of a machine which may add itself.**

**d. Follow the prompts to add the node's name to the list of recognized machines.**

The `scsetup` utility prints the message `Command completed successfully` if the task completes without error.

**e. Quit the `scsetup` utility.**

**7. Become superuser on the cluster node to configure.**

**8. Install Sun Web Console packages.**

These packages are required by Sun Cluster software, even if you do not use Sun Web Console.

**a. Insert the Sun Cluster 3.1 9/04 CD-ROM in the CD-ROM drive.**

**b. Change to the `/cdrom/cdrom0/Solaris_arch/Product/sun_web_console/2.1/` directory, where `arch` is `sparc` or `x86`.**

**c. Run the `setup` command.**

```
./setup
```

The `setup` command installs all packages to support Sun Web Console.

**9. Install additional packages if you intend to use any of the following features.**

- Remote Shared Memory Application Programming Interface (RSMAPI)
- SCI-PCI adapters for the interconnect transport
- RSMRDT drivers

---

**Note** – Use of the RSMRDT driver is restricted to clusters that run an Oracle9i release 2 SCI configuration with RSM enabled. Refer to Oracle9i release 2 user documentation for detailed installation and configuration instructions.

---

**a. Determine which packages you must install.**

The following table lists the Sun Cluster 3.1 9/04 packages that each feature requires and the order in which you must install each group of packages. The `scinstall` utility does not automatically install these packages.

| Feature          | Additional Sun Cluster 3.1 9/04 Packages to Install |
|------------------|-----------------------------------------------------|
| RSMAPI           | SUNWscrif                                           |
| SCI-PCI adapters | SUNWsci SUNWscid SUNWscidx                          |
| RSMRDT drivers   | SUNWscrtd                                           |

**b. Ensure that any dependency Solaris packages are already installed.**

See Step 8 in “How to Install Solaris Software” on page 52.

**c. On the Sun Cluster 3.1 9/04 CD-ROM, change to the `Solaris_arch/Product/sun_cluster/Solaris_ver/Packages/` directory, where `arch` is `sparc` or `x86`, and where `ver` is 8 (for Solaris 8) or 9 (for Solaris 9).**

```
cd /cdrom/cdrom0/Solaris_arch/Product/sun_cluster/Solaris_ver/Packages/
```

**d. Install the additional packages.**

```
pkgadd -d . packages
```

**e. If you are adding a node to a single-node cluster, repeat these steps to add the same packages to the original cluster node.**

**10. On the Sun Cluster 3.1 9/04 CD-ROM, change to the `Solaris_arch/Product/sun_cluster/Solaris_ver/Tools/` directory, where `arch` is `sparc` or `x86` and where `ver` is 8 (for Solaris 8) or 9 (for Solaris 9).**

```
cd /cdrom/cdrom0/Solaris_arch/Product/sun_cluster/Solaris_ver/Tools/
```

**11. Start the `scinstall` utility.**

```
/usr/cluster/bin/scinstall
```

**12. Follow these guidelines to use the interactive `scinstall` utility:**

- Interactive `scinstall` enables you to type ahead. Therefore, do not press the Return key more than once if the next menu screen does not appear immediately.

- Unless otherwise noted, you can press Control-D to return to either the start of a series of related questions or to the Main Menu.
- Default answers or answers to previous sessions are displayed in brackets ([ ]) at the end of a question. Press Return to enter the response that is in brackets without typing it.

**13. From the Main Menu, choose the menu item, Install a cluster or cluster node.**

```
*** Main Menu ***
```

```
Please select from one of the following (*) options:
```

- ```

* 1) Install a cluster or cluster node
    2) Configure a cluster to be JumpStarted from this install server
    3) Add support for new data services to this cluster node
* 4) Print release information for this cluster node

* ?) Help with menu options
* q) Quit

```

```
Option: 1
```

14. From the Install Menu, choose the menu item, Add this machine as a node in an existing cluster.

15. Follow the menu prompts to supply your answers from the worksheet that you completed in [Step 4](#).

The `scinstall` utility configures the node and boots the node into the cluster.

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

```
# eject cdrom
```

17. Repeat this procedure on any other node to add to the cluster until all additional nodes are fully configured.

18. From an active cluster member, prevent any other nodes from joining the cluster.

```
# /usr/cluster/bin/scconf -a -T node=.
```

```
-a          Add
```

```
-T          Specifies authentication options
```

```
node=.     Specifies the node name of dot (.) to add to the authentication list, to prevent any other node from adding itself to the cluster
```

Alternately, you can use the `scsetup(1M)` utility. See “How to Add a Cluster Node to the Authorized Node List” in “Adding and Removing a Cluster Node” in *Sun Cluster System Administration Guide for Solaris OS* for procedures.

19. Update the quorum vote count.

When you increase or decrease the number of node attachments to a quorum device, the cluster does not automatically recalculate the quorum vote count. This step reestablishes the correct quorum vote.

Use the `scsetup` utility to remove each quorum device and then add it back into the configuration. Do this for one quorum device at a time.

If the cluster has only one quorum device, configure a second quorum device before you remove and readd the original quorum device. Then remove the second quorum device to return the cluster to its original configuration.

20. Install Sun StorEdge QFS file system software.

Follow the procedures for initial installation in the *Sun StorEdge QFS and Sun StorEdge SAM-FS Software Installation and Configuration Guide*.

21. (Optional) SPARC: To install VERITAS File System, go to [“SPARC: How to Install VERITAS File System Software”](#) on page 97.

22. Set up the name-service look-up order.

Go to [“How to Configure the Name-Service Switch”](#) on page 98.

Example – Configuring Sun Cluster Software on an Additional Node

The following example shows the `scinstall` command executed and the messages that the utility logs as `scinstall` completes configuration tasks on the node `phys-schost-3`. The sponsoring node is `phys-schost-1`.

```
>>> Confirmation <<<

Your responses indicate the following options to scinstall:

scinstall -ik \
  -C sc-cluster \
  -N phys-schost-1 \
  -A trtype=dlpi,name=hme1 -A trtype=dlpi,name=hme3 \
  -m endpoint=:hme1,endpoint=switch1 \
  -m endpoint=:hme3,endpoint=switch2

Are these the options you want to use (yes/no) [yes]?

Do you want to continue with the install (yes/no) [yes]?

Checking device to use for global devices file system ... done

Adding node "phys-schost-3" to the cluster configuration ... done
Adding adapter "hme1" to the cluster configuration ... done
Adding adapter "hme3" to the cluster configuration ... done
Adding cable to the cluster configuration ... done
Adding cable to the cluster configuration ... done
```



```

Copying the config from "phys-schost-1" ... done
Setting the node ID for "phys-schost-3" ... done (id=3)

Verifying the major number for the "did" driver with "phys-schost-1" ...done

Checking for global devices global file system ... done
Updating vfstab ... done

Verifying that NTP is configured ... done
Installing a default NTP configuration ... done
Please complete the NTP configuration after scinstall has finished.

Verifying that "cluster" is set for "hosts" in nsswitch.conf ... done
Adding the "cluster" switch to "hosts" in nsswitch.conf ... done

Verifying that "cluster" is set for "netmasks" in nsswitch.conf ... done
Adding the "cluster" switch to "netmasks" in nsswitch.conf ... done

Verifying that power management is NOT configured ... done
Unconfiguring power management ... done
/etc/power.conf has been renamed to /etc/power.conf.61501001054
Power management is incompatible with the HA goals of the cluster.
Please do not attempt to re-configure power management.

Ensure that the EEPROM parameter "local-mac-address?" is set to "true" ...done
Ensure network routing is disabled ... done
Network routing has been disabled on this node by creating /etc/notrouter.
Having a cluster node act as a router is not supported by Sun Cluster.
Please do not re-enable network routing.

Log file - /var/cluster/logs/install/scinstall.log.9853

Rebooting ...

```

▼ SPARC: How to Install VERITAS File System Software

Perform this procedure on each node of the cluster.

1. **Follow the procedures in your VxFS installation documentation to install VxFS software on each node of the cluster.**
2. **Install any Sun Cluster patches that are required to support VxFS.**
See "Patches and Required Firmware Levels" in *Sun Cluster 3.1 Release Notes* for the location of patches and installation instructions.
3. **In the `/etc/system` file on each node, set the following values.**

```
set rpcmod:svc_default_stksize=0x8000
set lwp_default_stksize=0x6000
```

- Sun Cluster software requires a minimum `rpcmod:svc_default_stksize` setting of `0x8000`. Because VxFS installation sets the value of the `rpcmod:svc_default_stksize` variable to `0x4000`, you must manually set the value to `0x8000` after VxFS installation is complete.
- You must set the `lwp_default_stksize` variable in the `/etc/system` file to override the VxFS default value of `0x4000`.

4. Set up the name-service look-up order.

Go to [“How to Configure the Name-Service Switch”](#) on page 98.

▼ How to Configure the Name-Service Switch

Perform this task on each node in the cluster.

1. Become superuser on the cluster node.

2. Edit the `/etc/nsswitch.conf` file.

a. Verify that `cluster` is the first source look-up for the `hosts` and `netmasks` database entries.

This order is necessary for Sun Cluster software to function properly. The `scinstall(1M)` command adds `cluster` to these entries during installation.

b. (Optional) To increase availability to data services if the naming service becomes unavailable, change the look-up order of the following entries:

- For the `hosts` and `netmasks` database entries, follow `cluster` with `files`.
- For Sun Cluster HA for NFS, also insert `[SUCCESS=return]` after `cluster files` and before name services.

```
hosts:      cluster files [SUCCESS=return] nis
```

This look-up order ensures that, if the node resolves a name locally, the node does not contact the listed name service(s). Instead, the node returns success immediately.

- For all other database entries, place `files` first in the look-up order.
- If the `[NOTFOUND=return]` criterion becomes the last item of an entry after you modify the lookup order, the criterion is no longer necessary. You can either delete the `[NOTFOUND=return]` criterion from the entry or leave the criterion in the entry. A `[NOTFOUND=return]` criterion at the end of an entry is ignored.

c. Make any other changes that are required by specific data services.

See each manual for the data services that you installed.

The following example shows partial contents of an `/etc/nsswitch.conf` file. The look-up order for the `hosts` and `netmasks` database entries is first `cluster`, then `files`. The look-up order for other entries begins with `files`. The `[NOTFOUND=return]` criterion is removed from the entries.

```
# vi /etc/nsswitch.conf
...
passwd:    files nis
group:    files nis
...
hosts:    cluster files nis
...
netmasks: cluster files nis
...
```

See the `nsswitch.conf(4)` man page for more information about `nsswitch.conf` file entries.

3. Set up your root user's environment.

Go to "How to Set Up the Root Environment" on page 99.

▼ How to Set Up the Root Environment

Note – In a Sun Cluster configuration, user initialization files for the various shells must verify that they are run from an interactive shell. The files must verify this before they attempt to output to the terminal. Otherwise, unexpected behavior or interference with data services might occur. See "Customizing a User's Work Environment" in *System Administration Guide, Volume 1* (Solaris 8) or "Customizing a User's Work Environment" in *System Administration Guide: Basic Administration* (Solaris 9) for more information.

Perform this procedure on each node in the cluster.

1. **Become superuser on a cluster node.**
2. **Modify `PATH` and `MANPATH` entries in the `.cshrc` or `.profile` file.**

 - a. **Set the `PATH` to include `/usr/sbin/` and `/usr/cluster/bin/`.**
 - b. **Set the `MANPATH` to include `/usr/cluster/man/`.**

See your volume manager documentation and other application documentation for additional file paths to set.

3. **(Optional) For ease of administration, set the same root password on each node, if you have not already done so.**
4. **Install Sun Cluster 3.1 9/04 data-service software packages.**
 - To use the `scinstall` utility, go to [“How to Install Data-Service Software Packages \(`scinstall`\)” on page 102.](#)
 - To use the `installer` program, go to [“How to Install Data-Service Software Packages \(`installer`\)” on page 100.](#)

▼ How to Install Data-Service Software Packages (`installer`)

To install data services from the Sun Cluster 3.1 9/04 release, you can use the `installer` program to install the packages. To install data services from the Sun Cluster 3.1 release or earlier, follow the procedures in [“How to Install Data-Service Software Packages \(`scinstall`\)” on page 102.](#)

You can run the `installer` program with a command-line interface (CLI) or with a graphical user interface (GUI). The content and sequence of instructions in the CLI and the GUI are similar. For more information about the `installer` program, see the `installer(1M)` man page.

Perform this procedure on each node in the cluster on which you want to run a data service.

1. **Become superuser on the cluster node.**
2. **(Optional) If you intend to use the `installer` program with a GUI, ensure that the `DISPLAY` environment variable is set.**
3. **Load the Sun Cluster 3.1 9/04 Agents CD-ROM into the CD-ROM drive.**

If the volume management daemon `vold(1M)` is running and is configured to manage CD-ROM devices, the daemon automatically mounts the CD-ROM on the `/cdrom/cdrom0/` directory.
4. **Change to the directory where the CD-ROM is mounted.**

```
# cd /cdrom/cdrom0/
```
5. **Start the `installer` program.**

```
# ./installer
```
6. **When you are prompted, select the type of installation.**

See the Sun Cluster Release Notes for a listing of the locales that are available for each data service.

- To install all data services on the CD-ROM, select Typical.
- To install only a subset of the data services on the CD-ROM, select Custom.

7. When you are prompted, select the locale to install.

- To install only the C locale, select Typical.
- To install other locales, select Custom.

8. Follow instructions on the screen to install the data-service packages on the node.

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.

9. Quit the `installer` program.

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

```
# eject cdrom
```

11. Install any Sun Cluster data-service patches.

See “Patches and Required Firmware Levels” in *Sun Cluster 3.1 Release Notes* for the location of patches and installation instructions.

You do not have to reboot after you install Sun Cluster data-service patches unless a reboot is specified by the patch special instructions. If a patch instruction requires that you reboot, perform the following steps:

- a. **From one node, shut down the cluster by using the `scshutdown(1M)` command.**
- b. **Reboot each node in the cluster.**

Note – Until cluster installation mode is disabled, only the first-installed node, which established the cluster, has a quorum vote. In an established multiple-node cluster which is still in installation mode, if the cluster is not shut down before the first-installed node is rebooted, the remaining cluster nodes cannot obtain quorum. The entire cluster then shuts down.

Cluster nodes remain in installation mode until you run the `scsetup(1M)` command, during the procedure “[How to Perform Postinstallation Setup and Configure Quorum Devices](#)” on page 104.

12. Determine your next step.

- If this is a multiple-node cluster, perform postinstallation setup and assign quorum votes.
Go to [“How to Perform Postinstallation Setup and Configure Quorum Devices” on page 104.](#)
- If this is a single-node cluster, cluster installation is complete. You are now ready to install volume management software and to configure the cluster.
 - To install Solstice DiskSuite software or configure Solaris Volume Manager software, go to [“Installing and Configuring Solstice DiskSuite or Solaris Volume Manager Software” on page 129.](#)
 - SPARC: To install VERITAS Volume Manager software, go to [“SPARC: Installing and Configuring VxVM Software” on page 163.](#)

▼ How to Install Data-Service Software Packages (scinstall)

Note – You do not need to perform this procedure if you used SunPlex Installer to install Sun Cluster HA for NFS or Sun Cluster HA for Apache or both and if you do not intend to install any other data services. Instead, go to [“How to Perform Postinstallation Setup and Configure Quorum Devices” on page 104.](#)

Perform this task on each cluster node to install data services. If you install data services from the earlier Sun Cluster 3.1 10/03 release or compatible, you can alternatively use the `installer` program to install the packages. See [“How to Install Data-Service Software Packages \(installer\)” on page 100.](#)

1. **Become superuser on the cluster node.**
2. **Load the Sun Cluster 3.1 9/04 Agents CD-ROM into the CD-ROM drive on the node.**
If the volume management daemon `vold(1M)` is running and is configured to manage CD-ROM devices, the daemon automatically mounts the CD-ROM on the `/cdrom/cdrom0/` directory.
3. **Change to the directory where the CD-ROM is mounted.**

```
# cd /cdrom/cdrom0/
```
4. **Start the `scinstall(1M)` utility.**

```
# scinstall
```
5. **Follow these guidelines to use the interactive `scinstall` utility:**

- Interactive `scinstall` enables you to type ahead. Therefore, do not press the Return key more than once if the next menu screen does not appear immediately.
 - Unless otherwise noted, you can press Control-D to return to either the start of a series of related questions or to the Main Menu.
 - Default answers or answers to previous sessions are displayed in brackets ([]) at the end of a question. Press Return to enter the response that is in brackets without typing it.
6. **From the Main Menu, choose the menu item, Add support for new data services to this cluster node.**
 7. **Follow the prompts to select the data services to install.**

You must install the same set of data-service packages on each node. This requirement applies even if a node is not expected to host resources for an installed data service.
 8. **After the data services are installed, quit the `scinstall` utility.**
 9. **Change to a directory that does *not* reside on the CD-ROM and eject the CD-ROM.**

```
# eject cdrom
```

10. **Install any Sun Cluster data-service patches.**

See “Patches and Required Firmware Levels” in *Sun Cluster 3.1 Release Notes* for the location of patches and installation instructions.

You do not have to reboot after you install Sun Cluster data-service patches unless a reboot is specified by the patch special instructions. If a patch instruction requires that you reboot, perform the following steps:

- a. **From one node, shut down the cluster by using the `scshutdown(1M)` command.**
- b. **Reboot each node in the cluster.**

Note – Until cluster installation mode is disabled, only the first-installed node, which established the cluster, has a quorum vote. In an established multiple-node cluster which is still in installation mode, if the cluster is not shut down before the first-installed node is rebooted, the remaining cluster nodes cannot obtain quorum. This inability to obtain quorum causes the entire cluster to shut down.

Cluster nodes remain in installation mode until you run the `scsetup(1M)` command, during the procedure “[How to Perform Postinstallation Setup and Configure Quorum Devices](#)” on page 104.

11. Determine your next step.

- If this is a multiple-node cluster, perform postinstallation setup and assign quorum votes.
Go to [“How to Perform Postinstallation Setup and Configure Quorum Devices”](#) on page 104.
- If this is a single-node cluster, cluster installation is complete. You are now ready to install volume management software and to configure the cluster.
 - To install Solstice DiskSuite software or configure Solaris Volume Manager software, go to [“Installing and Configuring Solstice DiskSuite or Solaris Volume Manager Software”](#) on page 129.
 - SPARC: To install VERITAS Volume Manager software, go to [“SPARC: Installing and Configuring VxVM Software”](#) on page 163.

▼ How to Perform Postinstallation Setup and Configure Quorum Devices

Note – You do not need to configure quorum devices in the following circumstances:

- You chose automatic quorum configuration during Sun Cluster software configuration.
- You used SunPlex Installer to install the cluster. SunPlex Installer assigns quorum votes and removes the cluster from installation mode for you.
- You installed a single-node cluster.
- You added a node to an existing cluster and already have sufficient quorum votes assigned.

Instead, proceed to [“How to Verify the Quorum Configuration and Installation Mode”](#) on page 107.

Perform this procedure one time only, after the cluster is fully formed. Use this procedure to assign quorum votes and then to remove the cluster from installation mode.

1. From one node, verify that all nodes have joined the cluster.

Run the `scstat(1M)` command to display a list of the cluster nodes. You do not need to be logged in as superuser to run this command.

```
% scstat -n
```

Output resembles the following.


```
-- Cluster Nodes --
                Node name      Status
                -----
Cluster node:   phys-schost-1  Online
Cluster node:   phys-schost-2  Online
```

2. On each node, verify device connectivity to the cluster nodes.

Run the `scdidadm(1M)` command to display a list of all the devices that the system checks. You do not need to be logged in as superuser to run this command.

```
% scdidadm -L
```

The list on each node should be the same. Output resembles the following:

```
1      phys-schost-1:/dev/rdisk/c0t0d0 /dev/did/rdsk/d1
2      phys-schost-1:/dev/rdisk/c1t1d0 /dev/did/rdsk/d2
2      phys-schost-2:/dev/rdisk/c1t1d0 /dev/did/rdsk/d2
3      phys-schost-1:/dev/rdisk/c1t2d0 /dev/did/rdsk/d3
3      phys-schost-2:/dev/rdisk/c1t2d0 /dev/did/rdsk/d3
...
```

3. If you are adding a new node to an existing cluster, determine whether you need to update the quorum configuration to accommodate your cluster's new configuration.

If this is a new cluster, proceed to [Step 4](#).

- a. See **"Quorum Devices"** in *Sun Cluster Overview for Solaris OS and Sun Cluster Concepts Guide for Solaris OS* for information about quorum.
 - b. If you need to change the quorum configuration, follow procedures in **"Administering Quorum"** in *Sun Cluster System Administration Guide for Solaris OS*.
 - c. When the modified quorum configuration is satisfactory, go to **"How to Verify the Quorum Configuration and Installation Mode"** on page 107.
- 4. Determine the global device-ID name of each shared disk that you are configuring as a quorum device.**

Note – Any shared disk that you choose must be qualified for use as a quorum device. See **"Quorum Devices"** on page 30 for further information about choosing quorum devices.

Use the `scdidadm` output from [Step 2](#) to identify the device-ID name of each shared disk that you are configuring as a quorum device. For example, the output in [Step 2](#) shows that global device `d2` is shared by `phys-schost-1` and `phys-schost-2`. You use this information in [Step 7](#).

5. Become superuser on one node of the cluster.

6. Start the `scsetup(1M)` utility.

```
# scsetup
```

The Initial Cluster Setup screen is displayed.

Note – If the Main Menu is displayed instead, initial cluster setup was already successfully performed. Skip to [Step 9](#).

7. Answer the prompt `Do you want to add any quorum disks?`.

- If your cluster is a two-node cluster, you must configure at least one shared quorum device. Type **Yes** and follow the prompts to configure one or more quorum devices.
- If your cluster has three or more nodes, quorum device configuration is optional. Type **No** if you do not want to configure additional quorum devices or type **Yes** to configure more quorum devices.

Tip – If you later increase or decrease the number of node attachments to a quorum device, the quorum vote count is not automatically recalculated. You can reestablish the correct quorum vote by removing each quorum device and then add it back into the configuration, one quorum device at a time.

For a two-node cluster, temporarily add a new quorum device before you remove and add back the original quorum device. Then remove the temporary quorum device.

See the procedure “How to Modify a Quorum Device Node List” in “Administering Quorum” in *Sun Cluster System Administration Guide for Solaris OS*.

8. At the prompt `Is it okay to reset "installmode"?`, type **Yes.**

After the `scsetup` utility sets the quorum configurations and vote counts for the cluster, the message `Cluster initialization is complete` is displayed. The utility returns you to the Main Menu.

Tip – If the quorum setup process is interrupted or fails to be completed successfully, rerun `scsetup`.

9. Quit the `scsetup` utility.

10. Verify the quorum configuration and that installation mode is disabled.

Go to [“How to Verify the Quorum Configuration and Installation Mode”](#) on page 107.

▼ How to Verify the Quorum Configuration and Installation Mode

Perform this procedure to verify that quorum configuration completed successfully and that cluster installation mode is disabled.

1. **From any node, verify the device and node quorum configurations.**

```
% scstat -q
```

2. **From any node, verify that cluster installation mode is disabled.**

You do not need to be superuser to run this command.

```
% sccnf -p | grep "install mode"  
Cluster install mode:                disabled
```

Cluster installation is complete. You are now ready to install volume management software and to configure the cluster.

- To install Solstice DiskSuite software or configure Solaris Volume Manager software, go to [“Installing and Configuring Solstice DiskSuite or Solaris Volume Manager Software”](#) on page 129.
- SPARC: To install VERITAS Volume Manager software, go to [“SPARC: Installing and Configuring VxVM Software”](#) on page 163.

Task Map: Configuring the Cluster

The following table lists the tasks to perform to configure your cluster. Before you start to perform these tasks, ensure that you completed the following tasks:

- Cluster framework installation as described in [“Installing the Software”](#) on page 46
- Volume manager installation and configuration as described in [“Installing and Configuring Solstice DiskSuite or Solaris Volume Manager Software”](#) on page 129 or [“SPARC: Installing and Configuring VxVM Software”](#) on page 163

TABLE 2-9 Task Map: Configuring the Cluster

Task	Instructions
Create and mount cluster file systems.	"How to Create Cluster File Systems" on page 108
Configure IP Network Multipathing groups.	"How to Configure Internet Protocol (IP) Network Multipathing Groups" on page 113
(Optional) Change a node's private hostname.	"How to Change Private Hostnames" on page 115
Create or modify the NTP configuration file.	"How to Configure Network Time Protocol (NTP)" on page 116
(Optional) SPARC: Install the Sun Cluster module to Sun Management Center software.	"SPARC: Installing the Sun Cluster Module for Sun Management Center" on page 118 Sun Management Center documentation
Install third-party applications and configure the applications, data services, and resource groups.	<i>Sun Cluster Data Service Planning and Administration Guide for Solaris OS</i> Third-party application documentation

Configuring the Cluster

This section provides information and procedures to configure the software that you installed on the cluster.

▼ How to Create Cluster File Systems

Perform this procedure to create a cluster file system. Unlike a local file system, a cluster file system is accessible from any node in the cluster. If you used SunPlex Installer to install data services, SunPlex Installer might have already created one or more cluster file systems.



Caution – Any data on the disks is destroyed when you create a file system. Be sure that you specify the correct disk device name. If you specify the wrong device name, you might erase data that you did not intend to delete.

Perform this procedure for each cluster file system that you want to create.

- 1. Ensure that volume-manager software is installed and configured.**

For volume-manager installation procedures, see “Installing and Configuring Solstice DiskSuite or Solaris Volume Manager Software” on page 129 or “SPARC: Installing and Configuring VxVM Software” on page 163.

2. Become superuser on any node in the cluster.

Tip – For faster file-system creation, become superuser on the current primary of the global device for which you create a file system.

3. Create a file system.

- For a UFS file system, use the `newfs(1M)` command.

`newfs raw-disk-device`

The following table shows examples of names for the *raw-disk-device* argument. Note that naming conventions differ for each volume manager.

Volume Manager	Sample Disk Device Name	Description
Solstice DiskSuite or Solaris Volume Manager	<code>/dev/md/nfs/rdisk/d1</code>	Raw disk device <code>d1</code> within the <code>nfs</code> disk set
SPARC: VERITAS Volume Manager	<code>/dev/vx/rdisk/oradg/vol01</code>	Raw disk device <code>vol01</code> within the <code>oradg</code> disk group
None	<code>/dev/global/rdisk/d1s3</code>	Raw disk device <code>d1s3</code>

- For a Sun StorEdge QFS file system, follow the procedures for defining the configuration in the *Sun StorEdge QFS and Sun StorEdge SAM-FS Software Installation and Configuration Guide*.
- SPARC: For a VERITAS File System (VxFS) file system, follow the procedures that are provided in your VxFS documentation.

4. On each node in the cluster, create a mount-point directory for the cluster file system.

A mount point is required *on each node*, even if the cluster file system is not accessed on that node.

Tip – For ease of administration, create the mount point in the `/global/device-group/` directory. This location enables you to easily distinguish cluster file systems, which are globally available, from local file systems.

```
# mkdir -p /global/device-group/mountpoint/
```

device-group Name of the directory that corresponds to the name of the device group that contains the device

mountpoint Name of the directory on which to mount the cluster file system

5. **On each node in the cluster, add an entry to the `/etc/vfstab` file for the mount point.**

See the `vfstab(4)` man page for details.

- a. **In each entry, specify the required mount options for the type of file system that you use. See [Table 2-10](#), [Table 2-11](#), or [Table 2-12](#) for the list of required mount options.**

Note – Do *not* use the `logging` mount option for Solstice DiskSuite trans metadevices or Solaris Volume Manager transactional volumes. Trans metadevices and transactional volumes provide their own logging.

In addition, Solaris Volume Manager transactional-volume logging (formerly Solstice DiskSuite trans-metadvice logging) is scheduled to be removed from the Solaris OS in an upcoming Solaris release. Solaris UFS logging provides the same capabilities but superior performance, as well as lower system administration requirements and overhead.

TABLE 2-10 Mount Options for UFS Cluster File Systems

Mount Option	Description
<code>global</code>	Required. This option makes the file system globally visible to all nodes in the cluster.
<code>logging</code>	Required. This option enables logging.
<code>forcedirectio</code>	Required for cluster file systems that will host Oracle Real Application Clusters RDBMS data files, log files, and control files. Note – Oracle Real Application Clusters is supported for use only in SPARC based clusters.

TABLE 2-10 Mount Options for UFS Cluster File Systems (Continued)

Mount Option	Description
<code>onerror=panic</code>	<p>Required. You do not have to explicitly specify the <code>onerror=panic</code> mount option in the <code>/etc/vfstab</code> file. This mount option is already the default value if no other <code>onerror</code> mount option is specified.</p> <p>Note – Only the <code>onerror=panic</code> mount option is supported by Sun Cluster software. Do not use the <code>onerror=umount</code> or <code>onerror=lock</code> mount options. These mount options are not supported on cluster file systems for the following reasons:</p> <ul style="list-style-type: none"> ■ Use of the <code>onerror=umount</code> or <code>onerror=lock</code> mount option might cause the cluster file system to lock or become inaccessible. This condition might occur if the cluster file system experiences file corruption. ■ The <code>onerror=umount</code> or <code>onerror=lock</code> mount option might cause the cluster file system to become unmountable. This condition might thereby cause applications that use the cluster file system to hang or prevent the applications from being killed. <p>A node might require rebooting to recover from these states.</p>
<code>syncdir</code>	<p>Optional. If you specify <code>syncdir</code>, you are guaranteed POSIX-compliant file system behavior for the <code>write()</code> system call. If a <code>write()</code> succeeds, then this mount option ensures that sufficient space is on the disk.</p> <p>If you do not specify <code>syncdir</code>, the same behavior occurs that is seen with UFS file systems. When you do not specify <code>syncdir</code>, performance of writes that allocate disk blocks, such as when appending data to a file, can significantly improve. However, in some cases, without <code>syncdir</code> you would not discover an out-of-space condition (<code>ENOSPC</code>) until you close a file.</p> <p>You see <code>ENOSPC</code> on close only during a very short time after a failover. With <code>syncdir</code>, as with POSIX behavior, the out-of-space condition would be discovered before the close.</p>

See the `mount_ufs(1M)` man page for more information about UFS mount options.

TABLE 2-11 SPARC: Mount Parameters for Sun StorEdge QFS Shared File Systems

Mount Parameter	Description
<code>shared</code>	<p>Required. This option specifies that this is a shared file system, therefore globally visible to all nodes in the cluster.</p>



Caution – Ensure that settings in the `/etc/vfstab` file do not conflict with settings in the `/etc/opt/SUNWsamfs/samfs.cmd` file. Settings in the `/etc/vfstab` file override settings in the `/etc/opt/SUNWsamfs/samfs.cmd` file.

Certain data services such as Sun Cluster Support for Oracle Real Application Clusters have additional requirements and guidelines for QFS mount parameters. See your data service manual for any additional requirements.

See the `mount_samfs(1M)` man page for more information about QFS mount parameters.

Note – Logging is not enabled by an `/etc/vfstab` mount parameter. To enable logging, follow procedures in the *Sun StorEdge QFS and Sun StorEdge SAM-FS Software Installation and Configuration Guide*.

TABLE 2-12 SPARC: Mount Options for VxFS Cluster File Systems

Mount Option	Description
<code>global</code>	Required. This option makes the file system globally visible to all nodes in the cluster.
<code>log</code>	Required. This option enables logging.

See the VxFS `mount_vxfs` man page and “Administering Cluster File Systems Overview” in *Sun Cluster System Administration Guide for Solaris OS* for more information about VxFS mount options.

- b. **To automatically mount the cluster file system, set the `mount at boot` field to `yes`.**
- c. **Ensure that, for each cluster file system, the information in its `/etc/vfstab` entry is identical on each node.**
- d. **Ensure that the entries in each node’s `/etc/vfstab` file list devices in the same order.**
- e. **Check the boot order dependencies of the file systems.**

For example, consider the scenario where `phys-schost-1` mounts disk device `d0` on `/global/oracle/`, and `phys-schost-2` mounts disk device `d1` on `/global/oracle/logs/`. With this configuration, `phys-schost-2` can boot and mount `/global/oracle/logs/` only after `phys-schost-1` boots and mounts `/global/oracle/`.

6. On any node in the cluster, run the `sccheck(1M)` utility.

The `sccheck` utility verifies that the mount points exist. The utility also verifies that `/etc/vfstab` file entries are correct on all nodes of the cluster.

```
# sccheck
```

If no errors occur, nothing is returned.

7. Mount the cluster file system.

```
# mount /global/device-group/mountpoint/
```

- For UFS and QFS, mount the cluster file system from any node in the cluster.
- SPARC: For VxFS, mount the cluster file system from the current master of *device-group* to ensure that the file system mounts successfully. In addition, unmount a VxFS file system from the current master of *device-group* to ensure

that the file system unmounts successfully.

Note – To manage a VxFS cluster file system in a Sun Cluster environment, run administrative commands only from the primary node on which the VxFS cluster file system is mounted.

8. On each node of the cluster, verify that the cluster file system is mounted.

You can use either the `df(1M)` or `mount(1M)` command to list mounted file systems.

9. Configure IP Network Multipathing groups.

Go to “[How to Configure Internet Protocol \(IP\) Network Multipathing Groups](#)” on page 113.

Example – Creating a Cluster File System

The following example creates a UFS cluster file system on the Solstice DiskSuite metadvice `/dev/md/oracle/rdisk/d1`.

```
# newfs /dev/md/oracle/rdisk/d1
...

      (on each node)
# mkdir -p /global/oracle/d1
# vi /etc/vfstab
#device          device          mount   FS      fsck    mount  mount
#to mount        to fsck      point   type   ; pass  at boot options
#
/dev/md/oracle/dsk/d1 /dev/md/oracle/rdisk/d1 /global/oracle/d1 ufs 2 yes global,logging
      (save and exit)

      (on one node)
# sctest
# mount /global/oracle/d1
# mount
...
/global/oracle/d1 on /dev/md/oracle/dsk/d1 read/write/setuid/global/logging/largefiles
on Sun Oct 3 08:56:16 2000
```

▼ How to Configure Internet Protocol (IP) Network Multipathing Groups

Perform this task on each node of the cluster. If you used SunPlex Installer to install Sun Cluster HA for Apache or Sun Cluster HA for NFS, SunPlex Installer configured IP Network Multipathing groups for the public-network adapters those data services use. You must configure IP Network Multipathing groups for the remaining public-network adapters.

Note – All public-network adapters *must* belong to an IP Network Multipathing group.

1. **Have available your completed “Public Networks Worksheet” on page 256.**

2. **Configure IP Network Multipathing groups.**

Perform procedures for IPv4 addresses in “Deploying Network Multipathing” in *IP Network Multipathing Administration Guide* (Solaris 8) or “Administering Network Multipathing (Task)” in *System Administration Guide: IP Services* (Solaris 9).

Follow these additional requirements to configure IP Network Multipathing groups in a Sun Cluster configuration:

- Each public network adapter must belong to a multipathing group.
- For multipathing groups that contain two or more adapters, you must configure a test IP address for each adapter in the group. If a multipathing group contains only one adapter, you do not need to configure a test IP address.
- Test IP addresses for all adapters in the same multipathing group must belong to a single IP subnet.
- Test IP addresses must not be used by normal applications because the test IP addresses are not highly available.
- In the `/etc/default/mpathd` file, the value of `TRACK_INTERFACES_ONLY_WITH_GROUPS` must be `yes`.
- The name of a multipathing group has no requirements or restrictions.

3. **If you want to change any private hostnames, go to “How to Change Private Hostnames” on page 115.**

4. **If you did not install your own `/etc/inet/ntp.conf` file before you installed Sun Cluster software, install or create the NTP configuration file.**

Go to “How to Configure Network Time Protocol (NTP)” on page 116.

5. **If you are using Sun Cluster on a SPARC based system and you want to use Sun Management Center to monitor the cluster, install the Sun Cluster module for Sun Management Center.**

Go to “SPARC: Installing the Sun Cluster Module for Sun Management Center” on page 118.

6. **Install third-party applications, register resource types, set up resource groups, and configure data services.**

Follow procedures in the *Sun Cluster Data Services Planning and Administration Guide for Solaris OS* as well as in the documentation that is supplied with your application software.

▼ How to Change Private Hostnames

Perform this task if you do not want to use the default private hostnames, `clusternodenodeid-priv`, that are assigned during Sun Cluster software installation.

Note – Do *not* perform this procedure after applications and data services have been configured and have been started. Otherwise, an application or data service might continue to use the old private hostname after the hostname is renamed, which would cause hostname conflicts. If any applications or data services are running, stop them before you perform this procedure.

Perform this procedure on one active node of the cluster.

1. **Become superuser on a node in the cluster.**
2. **Start the `scsetup(1M)` utility.**
3. **From the Main Menu, choose the menu item, Private hostnames.**
4. **From the Private Hostname Menu, choose the menu item, Change a private hostname.**

```
# scsetup
```

5. **Follow the prompts to change the private hostname.**
Repeat for each private hostname to change.
6. **Verify the new private hostnames.**

```
# sccnf -pv | grep "private hostname"
(phys-schost-1) Node private hostname:      phys-schost-1-priv
(phys-schost-3) Node private hostname:      phys-schost-3-priv
(phys-schost-2) Node private hostname:      phys-schost-2-priv
```

7. **If you did not install your own `/etc/inet/ntp.conf` file before you installed Sun Cluster software, install or create the NTP configuration file.**
Go to “How to Configure Network Time Protocol (NTP)” on page 116.
8. **(Optional) SPARC: Configure Sun Management Center to monitor the cluster.**
Go to “SPARC: Installing the Sun Cluster Module for Sun Management Center” on page 118.
9. **Install third-party applications, register resource types, set up resource groups, and configure data services.**

See the documentation that is supplied with the application software and the *Sun Cluster Data Services Planning and Administration Guide for Solaris OS*.

▼ How to Configure Network Time Protocol (NTP)

Perform this task to create or modify the NTP configuration file after you install Sun Cluster software. You must also modify the NTP configuration file when you add a node to an existing cluster or when you change the private hostname of a node in the cluster. If you added a node to a single-node cluster, you must ensure that the NTP configuration file that you use is copied to the original cluster node.

The primary requirement when you configure NTP, or any time synchronization facility within the cluster, is that all cluster nodes must be synchronized to the same time. Consider accuracy of time on individual nodes to be of secondary importance to the synchronization of time among nodes. You are free to configure NTP as best meets your individual needs if this basic requirement for synchronization is met.

See the *Sun Cluster Concepts Guide for Solaris OS* for further information about cluster time. See the `/etc/inet/ntp.cluster` template file for additional guidelines on how to configure NTP for a Sun Cluster configuration.

1. **If you installed your own `/etc/inet/ntp.conf` file before you installed Sun Cluster software, you do not need to modify your `ntp.conf` file.**
Skip to [Step 8](#).
2. **Become superuser on a cluster node.**
3. **If you have your own file, copy your file to each node of the cluster.**
4. **If you do not have your own `/etc/inet/ntp.conf` file to install, use the `/etc/inet/ntp.conf.cluster` file as your NTP configuration file.**

Note – Do not rename the `ntp.conf.cluster` file as `ntp.conf`.

If the `/etc/inet/ntp.conf.cluster` file does not exist on the node, you might have an `/etc/inet/ntp.conf` file from an earlier installation of Sun Cluster software. Sun Cluster software creates the `/etc/inet/ntp.conf.cluster` file as the NTP configuration file if an `/etc/inet/ntp.conf` file is not already present on the node. If so, perform the following edits instead on that `ntp.conf` file.

- a. **Use your preferred text editor to open the `/etc/inet/ntp.conf.cluster` file on one node of the cluster for editing.**
- b. **Ensure that an entry exists for the private hostname of each cluster node.**
If you changed any node's private hostname, ensure that the NTP configuration file contains the new private hostname.
- c. **Remove any unused private hostnames.**
The `ntp.conf.cluster` file might contain nonexistent private hostnames. When a node is rebooted, the system generates error messages as the node

attempts to contact those nonexistent private hostnames.

d. If necessary, make other modifications to meet your NTP requirements.

5. Copy the NTP configuration file to all nodes in the cluster.

The contents of the NTP configuration file must be identical on all cluster nodes.

6. Stop the NTP daemon on each node.

Wait for the stop command to complete successfully on each node before you proceed to [Step 7](#).

```
# /etc/init.d/xntpd stop
```

7. Restart the NTP daemon on each node.

■ If you use the `ntp.conf.cluster` file, run the following command:

```
# /etc/init.d/xntpd.cluster start
```

The `xntpd.cluster` startup script first looks for the `/etc/inet/ntp.conf` file. If that file exists, the script exits immediately without starting the NTP daemon. If the `ntp.conf` file does not exist but the `ntp.conf.cluster` file does exist, the script starts the NTP daemon. In this case, the script uses the `ntp.conf.cluster` file as the NTP configuration file.

■ If you use the `ntp.conf` file, run the following command:

```
# /etc/init.d/xntpd start
```

8. (Optional) SPARC: Configure Sun Management Center to monitor the cluster.

Go to [“SPARC: Installing the Sun Cluster Module for Sun Management Center”](#) on page 118.

9. Install third-party applications, register resource types, set up resource groups, and configure data services.

See the documentation that is supplied with the application software and the *Sun Cluster Data Services Planning and Administration Guide for Solaris OS*.

SPARC: Task Map: Installing the Sun Cluster Module for Sun Management Center

The Sun Cluster module for Sun Management Center enables you to use Sun Management Center to monitor the cluster. The following table lists the tasks to perform to install the Sun Cluster–module software for Sun Management Center.

TABLE 2-13 Task Map: Installing the Sun Cluster Module for Sun Management Center

Task	Instructions
Install Sun Management Center server, help-server, agent, and console packages.	Sun Management Center documentation “SPARC: Installation Requirements for Sun Cluster Monitoring” on page 118
Install Sun Cluster-module packages.	“SPARC: How to Install the Sun Cluster Module for Sun Management Center” on page 119
Start Sun Management Center server, console, and agent processes.	“SPARC: How to Start Sun Management Center” on page 120
Add each cluster node as a Sun Management Center agent host object.	“SPARC: How to Add a Cluster Node as a Sun Management Center Agent Host Object” on page 121
Load the Sun Cluster module to begin to monitor the cluster.	“SPARC: How to Load the Sun Cluster Module” on page 121

SPARC: Installing the Sun Cluster Module for Sun Management Center

This section provides information and procedures to install the Sun Cluster module to Sun Management Center software.

SPARC: Installation Requirements for Sun Cluster Monitoring

The Sun Cluster module for Sun Management Center is used to monitor a Sun Cluster configuration. Perform the following tasks before you install the Sun Cluster module packages.

- **Space requirements** – Ensure that 25 Mbytes of space is available on each cluster node for Sun Cluster-module packages.
- **Sun Management Center packages** – You must install the Sun Management Center server, help-server, and console packages on noncluster nodes, and you must install the Sun Management Center agent package on each cluster node.

If you have an administrative console or other dedicated machine, you can run the console process on the administrative console and the server process on a separate machine. This installation approach improves Sun Management Center performance.

Follow procedures in your Sun Management Center documentation to install the Sun Management Center packages.

- **Simple Network Management Protocol (SNMP) port** – When you install Sun Management Center on an agent machine (cluster node), choose whether to use the default of 161 for the agent (SNMP) communication port or another number. This port number enables the server to communicate with this agent. Record the port number that you choose for reference later when you configure the cluster nodes for monitoring.

▼ SPARC: How to Install the Sun Cluster Module for Sun Management Center

Perform this procedure to install the Sun Cluster-module server and help-server packages.

Note – The Sun Cluster-module agent packages, `SUNWscsa1` and `SUNWscsam`, are already added to cluster nodes during Sun Cluster software installation.

1. Ensure that all Sun Management Center core packages are installed on the appropriate machines.

This step includes installing Sun Management Center agent packages on each cluster node. See your Sun Management Center documentation for installation instructions.

2. On the server machine, install the Sun Cluster-module server package `SUNWscssv`.

a. Become superuser.

- b. Insert the Sun Cluster 3.1 9/04 CD-ROM into the CD-ROM drive. If the volume management daemon `vold(1M)` is running and is configured to manage CD-ROM devices, the daemon automatically mounts the CD-ROM on the `/cdrom/cdrom0/` directory.**

- c. Change to the `Solaris_arch/Product/sun_cluster/Solaris_ver/Packages/` directory, where `arch` is `sparc` or `x86`, and where `ver` is 8 (for Solaris 8) or 9 (for Solaris 9).**

```
# cd /cdrom/cdrom0/Solaris_arch/Product/sun_cluster/Solaris_ver/Packages/
```

- d. Install the Sun Cluster-module server package.**

```
# pkgadd -d . SUNWscssv
```

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

```
# eject cdrom
```

3. **On the help-server machine, install the Sun Cluster–module help–server package `SUNWscsh1`.**

Use the same procedure as in the previous step.

4. **Install any Sun Cluster–module patches.**

See “Patches and Required Firmware Levels” in *Sun Cluster 3.1 Release Notes* for the location of patches and installation instructions.

5. **Start Sun Management Center.**

Go to “[SPARC: How to Start Sun Management Center](#)” on page 120.

▼ SPARC: How to Start Sun Management Center

Perform this procedure to start the Sun Management Center server, agent, and console processes.

1. **As superuser, on the Sun Management Center server machine, start the Sun Management Center server process.**

```
# /opt/SUNWsymon/sbin/es-start -s
```

2. **As superuser, on each Sun Management Center agent machine (cluster node), start the Sun Management Center agent process.**

```
# /opt/SUNWsymon/sbin/es-start -a
```

3. **On each Sun Management Center agent machine (cluster node), ensure that the `scsymon_srv` daemon is running.**

```
# ps -ef | grep scsymon_srv
```

If any cluster node is not already running the `scsymon_srv` daemon, start the daemon on that node.

```
# /usr/cluster/lib/scsymon/scsymon_srv
```

4. **On the Sun Management Center console machine (administrative console), start the Sun Management Center console.**

You do not need to be superuser to start the console process.

```
% /opt/SUNWsymon/sbin/es-start -c
```

5. **Type your login name, password, and server hostname, then click Login.**

6. **Add cluster nodes as monitored host objects.**

Go to “[SPARC: How to Add a Cluster Node as a Sun Management Center Agent Host Object](#)” on page 121.

▼ SPARC: How to Add a Cluster Node as a Sun Management Center Agent Host Object

Perform this procedure to create a Sun Management Center agent host object for a cluster node.

Note – You need only one cluster node host object to use Sun Cluster–module monitoring and configuration functions for the entire cluster. However, if that cluster node becomes unavailable, connection to the cluster through that host object also becomes unavailable. Then you need another cluster-node host object to reconnect to the cluster.

1. **From the Sun Management Center main window, select a domain from the Sun Management Center Administrative Domains pull-down list.**

This domain contains the Sun Management Center agent host object that you create. During Sun Management Center software installation, a Default Domain was automatically created for you. You can use this domain, select another existing domain, or create a new domain.

See your Sun Management Center documentation for information about how to create Sun Management Center domains.

2. **Choose Edit⇒Create an Object from the pull-down menu.**
3. **Click the Node tab.**
4. **From the Monitor Via pull-down list, select Sun Management Center Agent - Host.**
5. **Fill in the name of the cluster node, for example, `phys-schost-1`, in the Node Label and Hostname text fields.**
Leave the IP text field blank. The Description text field is optional.
6. **In the Port text field, type the port number that you chose when you installed the Sun Management Center agent machine.**
7. **Click OK.**
A Sun Management Center agent host object is created in the domain.
8. **Load the Sun Cluster module.**
Go to [“SPARC: How to Load the Sun Cluster Module”](#) on page 121.

▼ SPARC: How to Load the Sun Cluster Module

Perform this procedure to start cluster monitoring.

1. **In the Sun Management Center main window, right-click the icon of a cluster node.**
The pull-down menu is displayed.
2. **Choose Load Module.**
The Load Module window lists each available Sun Management Center module and whether the module is currently loaded.
3. **Choose Sun Cluster: Not Loaded and click OK.**
The Module Loader window shows the current parameter information for the selected module.
4. **Click OK.**
After a few moments, the module is loaded. A Sun Cluster icon is then displayed in the Details window.
5. **In the Details window under the Operating System category, expand the Sun Cluster subtree in either of the following ways:**
 - In the tree hierarchy on the left side of the window, place the cursor over the Sun Cluster-module icon and single-click the left mouse button.
 - In the topology view on the right side of the window, place the cursor over the Sun Cluster -module icon and double-click the left mouse button.
6. **See the Sun Cluster-module online help for information about how to use Sun Cluster-module features.**
 - To view online help for a specific Sun Cluster-module item, place the cursor over the item. Then click the right mouse button and select Help from the pop-up menu.
 - To access the home page for the Sun Cluster-module online help, place the cursor over the Cluster Info icon. Then click the right mouse button and select Help from the pop-up menu.
 - To directly access the home page for the Sun Cluster-module online help, click the Sun Management Center Help button to launch the help browser. Then go to the following URL:
`file:/opt/SUNWsymon/lib/locale/C/help/main.top.html`

Note – The Help button in the Sun Management Center browser accesses online help for Sun Management Center, not the topics specific to the Sun Cluster module.

See Sun Management Center online help and your Sun Management Center documentation for information about how to use Sun Management Center.

7. **Install third-party applications, register resource types, set up resource groups, and configure data services.**

See the documentation that is supplied with the application software and the *Sun Cluster Data Services Planning and Administration Guide for Solaris OS*.

Uninstalling the Software

This section provides the following procedures to uninstall or remove Sun Cluster software:

- [“How to Uninstall Sun Cluster Software to Correct Installation Problems”](#) on page 123
- [“How to Uninstall the SUNWscrdt Package”](#) on page 124
- [“How to Unload the RSMRDT Driver Manually”](#) on page 125

▼ How to Uninstall Sun Cluster Software to Correct Installation Problems

Perform this procedure if the installed node cannot join the cluster or if you need to correct configuration information. For example, perform this procedure to reconfigure the transport adapters or the private-network address.

Note – If the node has already joined the cluster and is no longer in installation mode (see [Step 2 of “How to Verify the Quorum Configuration and Installation Mode”](#) on page 107), do not perform this procedure. Instead, go to [“How to Uninstall Sun Cluster Software From a Cluster Node”](#) in [“Adding and Removing a Cluster Node”](#) in *Sun Cluster System Administration Guide for Solaris OS*.

1. Attempt to reinstall the node.

You can correct certain failed installations by repeating Sun Cluster software installation on the node. If you have already tried to reinstall the node without success, proceed to [Step 2](#) to uninstall Sun Cluster software from the node.

2. Become superuser on an active cluster member other than the node that you are uninstalling.

3. From an active cluster member, add the node that you intend to uninstall to the cluster node-authentication list.

Skip this step if you are uninstalling a single-node cluster.

```
# /usr/cluster/bin/scconf -a -T node=nodename
```

```
-a          Add
```

-T Specifies authentication options
node=*nodename* Specifies the name of the node to add to the authentication list
Alternately, you can use the `scsetup(1M)` utility. See “How to Add a Cluster Node to the Authorized Node List” in “Adding and Removing a Cluster Node” in *Sun Cluster System Administration Guide for Solaris OS* for procedures.

4. Become superuser on the node that you intend to uninstall.

5. Shut down the node that you intend to uninstall.

```
# shutdown -g0 -y -i0
```

6. Reboot the node into noncluster mode.

- On SPARC based systems, do the following:

```
ok boot -x
```

- On x86 based systems, do the following:

```
<<< Current Boot Parameters >>>
Boot path: /pci@0,0/pci-ide@7,1/ata@1/cmdk@0,0:b
Boot args:

Type   b [file-name] [boot-flags] <ENTER>  to boot with options
or     i <ENTER>                          to enter boot interpreter
or     <ENTER>                             to boot with defaults

<<< timeout in 5 seconds >>>
Select (b)oot or (i)nterpreter: b -x
```

7. Change to a directory, such as the root (/) directory, that does not contain any files that are delivered by the Sun Cluster packages.

```
# cd /
```

8. Uninstall Sun Cluster software from the node.

```
# /usr/cluster/bin/scinstall -r
```

See the `scinstall(1M)` man page for more information.

9. Reinstall and reconfigure Sun Cluster software on the node.

Refer to [Table 2-1](#) for the list of all installation tasks and the order in which to perform the tasks.

▼ How to Uninstall the SUNWscrdt Package

Perform this procedure on each node in the cluster.

1. Verify that no applications are using the RSMRDT driver before performing this procedure.
2. Become superuser on the node to which you want to uninstall the `SUNWscrtd` package.
3. Uninstall the `SUNWscrtd` package.

```
# pkgrm SUNWscrtd
```

▼ How to Unload the RSMRDT Driver Manually

If the driver remains loaded in memory after completing “[How to Uninstall the SUNWscrtd Package](#)” on page 124, perform this procedure to unload the driver manually.

1. Start the `adb` utility.

```
# adb -kw
```

2. Set the kernel variable `clifrsmrtd_modunload_ok` to 1.

```
physmem NNNN
clifrsmrtd_modunload_ok/W 1
```

3. Exit the `adb` utility by pressing Control-D.

4. Find the `clif_rsmrtd` and `rsmrtd` module IDs.

```
# modinfo | grep rdt
```

5. Unload the `clif_rsmrtd` module.

You must unload the `clif_rsmrtd` module before you unload the `rsmrtd` module.

```
# modunload -i clif_rsmrtd_id
```

Tip – If the `modunload` command fails, applications are probably still using the driver. Terminate the applications before you run `modunload` again.

clif_rsmrtd_id Specifies the numeric ID for the module being unloaded.

6. Unload the `rsmrtd` module.

```
# modunload -i rsmrtd_id
```

rsmrtd_id Specifies the numeric ID for the module being unloaded.

7. Verify that the module was successfully unloaded.

```
# modinfo | grep rdt
```

Example—Unloading the RSMRDT Driver

The following example shows the console output after the RSMRDT driver is manually unloaded.

```
# adb -kw
physmem fc54
clifrsrmt_modunload_ok/W 1
clifrsrmt_modunload_ok: 0x0 = 0x1
^D
# modinfo | grep rsm
 88 f064a5cb 974 - 1 rsmops (RSMOPS module 1.1)
 93 f08e07d4 b95 - 1 clif_rsrmt (CLUSTER-RSMRDT Interface module)
 94 f0d3d000 13db0 194 1 rsrmt (Reliable Datagram Transport dri)
# modunload -i 93
# modunload -i 94
# modinfo | grep rsm
 88 f064a5cb 974 - 1 rsmops (RSMOPS module 1.1)
#
```

Installing and Configuring Solstice DiskSuite or Solaris Volume Manager Software

Install and configure your local and multihost disks for Solstice DiskSuite or Solaris Volume Manager software by using the procedures in this chapter, along with the planning information in “[Planning Volume Management](#)” on page 33. See your Solstice DiskSuite or Solaris Volume Manager documentation for additional details.

Note – DiskSuite Tool (Solstice DiskSuite `metatool`) and the Enhanced Storage module of Solaris Management Console (Solaris Volume Manager) are not compatible with Sun Cluster software. Use the command-line interface or Sun Cluster utilities to configure Solstice DiskSuite or Solaris Volume Manager software.

The following information and procedures are in this chapter:

- “[Task Map: Installing and Configuring Solstice DiskSuite or Solaris Volume Manager Software](#)” on page 128
- “[Solstice DiskSuite or Solaris Volume Manager Configuration Example](#)” on page 130
- “[How to Install Solstice DiskSuite Software](#)” on page 131
- “[How to Set the Number of Metadevice or Volume Names and Disk Sets](#)” on page 133
- “[How to Create State Database Replicas](#)” on page 135
- “[How to Mirror the Root \(/\) File System](#)” on page 136
- “[How to Mirror the Global Namespace](#)” on page 140
- “[How to Mirror File Systems Other Than Root \(/\) That Cannot Be Unmounted](#)” on page 143
- “[How to Mirror File Systems That Can Be Unmounted](#)” on page 147
- “[How to Create a Disk Set](#)” on page 151
- “[How to Add Drives to a Disk Set](#)” on page 154
- “[How to Repartition Drives in a Disk Set](#)” on page 155
- “[How to Create an `md.tab` File](#)” on page 156
- “[How to Activate Metadevices or Volumes](#)” on page 157
- “[Requirements for Dual-String Mediators](#)” on page 159
- “[How to Add Mediator Hosts](#)” on page 160
- “[How to Check the Status of Mediator Data](#)” on page 160

- [“How to Fix Bad Mediator Data”](#) on page 161

Task Map: Installing and Configuring Solstice DiskSuite or Solaris Volume Manager Software

The following table lists the tasks that you perform to install and configure Solstice DiskSuite or Solaris Volume Manager software for Sun Cluster configurations. You can skip certain procedures under the following conditions:

- If you used SunPlex Installer to install Solstice DiskSuite software (Solaris 8), the procedures [“How to Install Solstice DiskSuite Software”](#) on page 131 through [“How to Create State Database Replicas”](#) on page 135 are already completed. Go to [“Mirroring the Root Disk”](#) on page 136 or [“Creating Disk Sets in a Cluster”](#) on page 150 to continue to configure Solstice DiskSuite software.
- If you installed Solaris 9 software, Solaris Volume Manager is already installed. You can start configuration at [“How to Set the Number of Metadevice or Volume Names and Disk Sets”](#) on page 133.
- If you plan to create a Solaris Volume Manager for Sun Cluster multi-owner disk set for use by Oracle Real Application Clusters, do not use the procedures in [“How to Create a Disk Set”](#) on page 151 through [“How to Activate Metadevices or Volumes”](#) on page 157 to create the multi-owner disk set.

Instead, perform the procedures in [“Creating a Multi-Owner Disk Set in Solaris Volume Manager for Sun Cluster for the Oracle Real Application Clusters Database”](#) in *Sun Cluster Data Service for Oracle Real Application Clusters Guide for Solaris OS*.

TABLE 3-1 Task Map: Installing and Configuring Solstice DiskSuite or Solaris Volume Manager Software

Task	Instructions
1. Install and Configure Solstice DiskSuite or Solaris Volume Manager software	
1a. Plan the layout of your Solstice DiskSuite or Solaris Volume Manager configuration.	<ul style="list-style-type: none"> ■ “Planning Volume Management” on page 33 ■ “Solstice DiskSuite or Solaris Volume Manager Configuration Example” on page 130

TABLE 3-1 Task Map: Installing and Configuring Solstice DiskSuite or Solaris Volume Manager Software (Continued)

Task	Instructions
1b. (<i>Solaris 8 only</i>) Install Solstice DiskSuite software.	"How to Install Solstice DiskSuite Software" on page 131
1c. Calculate the number of metadevice names and disk sets needed for your configuration, and modify the <code>/kernel/drv/md.conf</code> file.	"How to Set the Number of Metadevice or Volume Names and Disk Sets" on page 133
1d. Create state database replicas on the local disks.	"How to Create State Database Replicas" on page 135
1e. (<i>Optional</i>) Mirror file systems on the root disk.	"Mirroring the Root Disk" on page 136
2. Create disk sets	
2a. Create disk sets by using the <code>metaset</code> command.	"How to Create a Disk Set" on page 151
2b. Add drives to the disk sets.	"How to Add Drives to a Disk Set" on page 154
2c. (<i>Optional</i>) Repartition drives in a disk set to allocate space to slices 1 through 6.	"How to Repartition Drives in a Disk Set" on page 155
2d. List DID pseudo-driver mappings and define metadevices or volumes in the <code>/etc/lvm/md.tab</code> files.	"How to Create an <code>md.tab</code> File" on page 156
2e. Initialize the <code>md.tab</code> files.	"How to Activate Metadevices or Volumes" on page 157
3. (<i>Dual-string configurations only</i>) Configure dual-string mediator hosts, check the status of mediator data, and, if necessary, fix bad mediator data.	<ol style="list-style-type: none"> "How to Add Mediator Hosts" on page 160 "How to Check the Status of Mediator Data" on page 160
4. Configure the cluster.	"Configuring the Cluster" on page 108

Installing and Configuring Solstice DiskSuite or Solaris Volume Manager Software

This section provides the following information and procedures to install and configure Solstice DiskSuite or Solaris Volume Manager software:

- "Solstice DiskSuite or Solaris Volume Manager Configuration Example" on page 130

- “How to Install Solstice DiskSuite Software” on page 131
- “How to Set the Number of Metadevice or Volume Names and Disk Sets” on page 133
- “How to Create State Database Replicas” on page 135
- “How to Mirror the Root (/) File System” on page 136
- “How to Mirror the Global Namespace” on page 140
- “How to Mirror File Systems Other Than Root (/) That Cannot Be Unmounted” on page 143
- “How to Mirror File Systems That Can Be Unmounted” on page 147

Solstice DiskSuite or Solaris Volume Manager Configuration Example

The following example helps to explain the process for determining the number of drives to place in each disk set. In this example, three storage devices are used. Existing applications are run over NFS (two file systems of 5 Gbytes each) and two ORACLE databases (one 5 Gbytes and one 10 Gbytes).

The following table shows the calculations that are used to determine the number of drives that are needed in the sample configuration. In a configuration with three storage devices, you would need 28 drives, which would be divided as evenly as possible among each of the three storage devices. Note that the 5-Gbyte file systems were given an additional 1 Gbyte of disk space because the number of drives needed was rounded up.

TABLE 3-2 Determining the Number of Drives Needed for a Configuration

Use	Data	Disk Storage Needed	Drives Needed
nfs1	5 Gbytes	3x2.1 Gbyte disks * 2 (Mirror)	6
nfs2	5 Gbytes	3x2.1 Gbyte disks * 2 (Mirror)	6
SPARC: oracle1	5 Gbytes	3x2.1 Gbyte disks * 2 (Mirror)	6
SPARC: oracle2	10 Gbytes	5x2.1 Gbyte disks * 2 (Mirror)	10

The following table shows the allocation of drives among the two disk sets and four data services.

TABLE 3-3 Division of Disk Sets

Disk Set	Data Services	Drives	Storage Device 1	Storage Device 2	Storage Device 3
dg-schost-1	nfs1, oracle1	12	4	4	4
dg-schost-2	nfs2, oracle2	16	5	6	5

Initially, four drives on each storage device (a total of 12 drives) are assigned to dg-schost-1, and five or six drives on each (a total of 16 drives) are assigned to dg-schost-2.

No hot spare disks are assigned to either disk set. A minimum of one hot spare disk per storage device per disk set enables one drive to be hot spared, which restores full two-way mirroring.

▼ How to Install Solstice DiskSuite Software

Note – Do not perform this procedure under the following circumstances:

- You used SunPlex Installer to install Solstice DiskSuite software.
 - If you plan to create additional disk sets, go to [“How to Set the Number of Metadevice or Volume Names and Disk Sets”](#) on page 133.
 - If you do not plan to create additional disk sets, go to [“Mirroring the Root Disk”](#) on page 136.
- You installed Solaris 9 software.

Solaris Volume Manager software is automatically installed with Solaris 9 software. Instead, go to [“How to Set the Number of Metadevice or Volume Names and Disk Sets”](#) on page 133.

Perform this task on each node in the cluster.

1. Have available the following information.

- Mappings of your storage drives.
- The following completed configuration planning worksheets. See [“Planning Volume Management”](#) on page 33 for planning guidelines.
 - [“Local File System Layout Worksheet”](#) on page 254
 - [“Disk Device Group Configurations Worksheet”](#) on page 260
 - [“Volume-Manager Configurations Worksheet”](#) on page 262
 - [“Metadevices Worksheet \(Solstice DiskSuite or Solaris Volume Manager\)”](#) on page 264

2. **Become superuser on the cluster node.**
3. **If you install from the CD-ROM, insert the Solaris 8 Software 2 of 2 CD-ROM into the CD-ROM drive on the node.**
This step assumes that the Volume Management daemon `vold(1M)` is running and configured to manage CD-ROM devices.
4. **Install the Solstice DiskSuite software packages in the order that is shown in the following example.**

```
# cd /cdrom/sol_8_sparc_2/Solaris_8/EA/products/DiskSuite_4.2.1/sparc/Packages
# pkgadd -d . SUNWmdr SUNWmdu [SUNWmdx] optional-pkgs
```

Note – If you have Solstice DiskSuite software patches to install, do not reboot after you install the Solstice DiskSuite software.

The `SUNWmdr` and `SUNWmdu` packages are required for all Solstice DiskSuite installations. The `SUNWmdx` package is also required for the 64-bit Solstice DiskSuite installation.

See your Solstice DiskSuite installation documentation for information about optional software packages.

5. **If you installed from a CD-ROM, eject the CD-ROM.**
6. **Install any Solstice DiskSuite patches.**
See “Patches and Required Firmware Levels” in *Sun Cluster 3.1 Release Notes* for the location of patches and installation instructions.
7. **Repeat [Step 1](#) through [Step 6](#) on the other nodes of the cluster.**
8. **From one node of the cluster, manually populate the global-device namespace for Solstice DiskSuite.**

```
# scgdevs
```

Note – The `scgdevs` command might return a message similar to the following:

```
Could not open /dev/rdisk/c0t6d0s2 to verify device id, Device busy
```

If the listed device is a CD-ROM device, you can safely ignore the message.

9. **Set the number of metadevice names and disk sets that are expected in the cluster.**
Go to “[How to Set the Number of Metadevice or Volume Names and Disk Sets](#)” on page 133.

▼ How to Set the Number of Metadevice or Volume Names and Disk Sets

Note – If you used SunPlex Installer to install Solstice DiskSuite software, do not perform this procedure. Instead, go to [“Mirroring the Root Disk” on page 136](#).

This procedure describes how to determine the number of Solstice DiskSuite metadevice or Solaris Volume Manager volume names and disk sets that are needed for your configuration. This procedure also describes how to modify the `/kernel/drv/md.conf` file to specify these numbers.

Tip – The default number of metadevice or volume names per disk set is 128, but many configurations need more than the default. Increase this number before you implement a configuration, to save administration time later.

At the same time, keep the value of the `nmdfield` and the `md_nsets` field as low as possible. Memory structures exist for all possible devices as determined by `nmdand` and `md_nsets`, even if you have not created those devices. For optimal performance, keep the value of `nmd` and `md_nsets` only slightly higher than the number of metadevices or volumes you plan to use.

1. **Have available the [“Disk Device Group Configurations Worksheet” on page 260](#).**

2. **Determine the total number of disk sets that you expect to need in the cluster, then add one more disk set for private disk management.**

The cluster can have a maximum of 32 disk sets, 31 disk sets for general use plus one disk set for private disk management. The default number of disk sets is 4. You supply this value for the `md_nsets` field in [Step 4](#).

3. **Determine the largest metadevice or volume name that you expect to need for any disk set in the cluster.**

Each disk set can have a maximum of 8192 metadevice or volume names. You supply this value for the `nmd` field in [Step 4](#).

a. **Determine the quantity of metadevice or volume names that you expect to need for each disk set.**

If you use local metadevices or volumes, ensure that each local metadevice or volume name is unique throughout the cluster and does not use the same name as any device-ID name in the cluster.

Tip – Choose a range of numbers to use exclusively for device-ID names and a range for each node to use exclusively for its local metadvice or volume names. For example, device-ID names might use the range from d1 to d100. Local metadvice or volumes on node 1 might use names in the range from d100 to d199. And local metadvice or volumes on node 2 might use d200 to d299.

- b. **Determine the highest of the metadvice or volume names that you expect to use in any disk set.**

The quantity of metadvice or volume names to set is based on the metadvice or volume name *value* rather than on the *actual quantity*. For example, if your metadvice or volume names range from d950 to d1000, Solstice DiskSuite or Solaris Volume Manager software requires that you set the value at 1000 names, not 50.

4. **On each node, become superuser and edit the `/kernel/drv/md.conf` file.**



Caution – All cluster nodes (or cluster pairs in the cluster-pair topology) must have identical `/kernel/drv/md.conf` files, regardless of the number of disk sets served by each node. Failure to follow this guideline can result in serious Solstice DiskSuite or Solaris Volume Manager errors and possible loss of data.

- a. **Set the `md_nsets` field to the value that you determined in [Step 2](#).**

- b. **Set the `nmd` field to the value that you determined in [Step 3](#).**

5. **On each node, perform a reconfiguration reboot.**

```
# touch /reconfigure
# shutdown -g0 -y -i6
```

Changes to the `/kernel/drv/md.conf` file become operative after you perform a reconfiguration reboot.

6. **Create local state database replicas.**

Go to “[How to Create State Database Replicas](#)” on page 135.

▼ How to Create State Database Replicas

Note – If you used SunPlex Installer to install Solstice DiskSuite software, do not perform this procedure. Instead, go to [“Mirroring the Root Disk” on page 136](#).

Perform this procedure on each node in the cluster.

1. **Become superuser on the cluster node.**
2. **Create state database replicas on one or more local devices for each cluster node.**

Use the physical name (cNtXdYsZ), not the device-ID name (dN), to specify the slices to use.

```
# metadb -af slice-1 slice-2 slice-3
```

Tip – To provide protection of state data, which is necessary to run Solstice DiskSuite or Solaris Volume Manager software, create at least three replicas for each node. Also, you can place replicas on more than one device to provide protection if one of the devices fails.

See the `metadb(1M)` man page and your Solstice DiskSuite or Solaris Volume Manager documentation for details.

3. **Verify the replicas.**

```
# metadb
```

The `metadb` command displays the list of replicas.

4. **To mirror file systems on the root disk, go to [“Mirroring the Root Disk” on page 136](#).**

Otherwise, go to [“Creating Disk Sets in a Cluster” on page 150](#) to create Solstice DiskSuite or Solaris Volume Manager disk sets.

Example—Creating State Database Replicas

The following example shows three Solstice DiskSuite state database replicas. Each replica is created on a different device. For Solaris Volume Manager, the replica size would be larger.

```
# metadb -af c0t0d0s7 c0t1d0s7 c1t0d0s7
# metadb
flags          first blk      block count
a              u             16          1034      /dev/dsk/c0t0d0s7
a              u             16          1034      /dev/dsk/c0t1d0s7
a              u             16          1034      /dev/dsk/c1t0d0s7
```

Mirroring the Root Disk

Mirroring the root disk prevents the cluster node itself from shutting down because of a system disk failure. Four types of file systems can reside on the root disk. Each file-system type is mirrored by using a different method.

Use the following procedures to mirror each type of file system.

- “How to Mirror the Root (/) File System” on page 136
- “How to Mirror the Global Namespace” on page 140
- “How to Mirror File Systems Other Than Root (/) That Cannot Be Unmounted” on page 143
- “How to Mirror File Systems That Can Be Unmounted” on page 147

Note – Some of the steps in these mirroring procedures can cause an error message similar to the following, which is harmless and can be ignored.

```
metainit: dg-schost-1: d1s0: not a metadvice
```



Caution – For local disk mirroring, do not use `/dev/global` as the path when you specify the disk name. If you specify this path for anything other than cluster file systems, the system cannot boot.

▼ How to Mirror the Root (/) File System

Use this procedure to mirror the root (/) file system.

1. **Become superuser on the node.**
2. **Use the `metainit(1M)` command to put the root slice in a single-slice (one-way) concatenation.**

Specify the physical disk name of the root-disk slice (`cNtXdYsZ`).

```
# metainit -f submirror1 1 1 root-disk-slice
```

3. **Create a second concatenation.**

```
# metainit submirror2 1 1 submirror-disk-slice
```

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

```
# metainit mirror -m submirror1
```

Note – The metadvice or volume name for the mirror *must* be unique throughout the cluster.

5. Run the `metaroot(1M)` command.

This command edits the `/etc/vfstab` and `/etc/system` files so the system can be booted with the root (`/`) file system on a metadvice or volume.

```
# metaroot mirror
```

6. Run the `lockfs(1M)` command.

This command flushes all transactions out of the log and writes the transactions to the master file system on all mounted UFS file systems.

```
# lockfs -fa
```

7. Move any resource groups or device groups from the node.

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

`-S` Moves all resource groups and device groups

`-h from-node` Specifies the name of the node from which to move resource or device groups

8. Reboot the node.

This command remounts the newly mirrored root (`/`) file system.

```
# shutdown -g0 -y -i6
```

9. Use the `metattach(1M)` command to attach the second submirror to the mirror.

```
# metattach mirror submirror2
```

10. If the disk that is used to mirror the root disk is physically connected to more than one node (multihosted), enable the `localonly` property.

Perform the following steps to enable the `localonly` property of the raw-disk device group for the disk that is used to mirror the root disk. You must enable the `localonly` property to prevent unintentional fencing of a node from its boot device if the boot device is connected to multiple nodes.

a. If necessary, use the `scdidadm(1M) -L` command to display the full device-ID path name of the raw-disk device group.

In the following example, the raw-disk device-group name `dsk/d2` is part of the third column of output, which is the full device-ID path name.

```
# scdidadm -L
...
1      phys-schost-3:/dev/rdisk/c1t1d0    /dev/did/rdisk/d2
```

b. View the node list of the raw-disk device group.

Output looks similar to the following:

```
# scconf -pvv | grep dsk/d2
Device group name:                dsk/d2
...
  (dsk/d2) Device group node list:  phys-schost-1, phys-schost-3
...
```

- c. **If the node list contains more than one node name, remove all nodes from the node list except the node whose root disk you mirrored.**

Only the node whose root disk you mirrored should remain in the node list for the raw-disk device group.

```
# scconf -r -D name=dsk/dN,nodeList=node

-D name=dsk/dN    Specifies the cluster-unique name of the raw-disk device
                  group
nodeList=node    Specifies the name of the node or nodes to remove from
                  the node list
```

- d. **Use the `scconf(1M)` command to enable the `localonly` property.**

When the `localonly` property is enabled, the raw-disk device group is used exclusively by the node in its node list. This usage prevents unintentional fencing of the node from its boot device if the boot device is connected to multiple nodes.

```
# scconf -c -D name=rawdisk-groupname,localonly=true

-D name=rawdisk-groupname    Specifies the name of the raw-disk device
                              group
```

For more information about the `localonly` property, see the `scconf_dg_rawdisk(1M)` man page.

11. Record the alternate boot path for possible future use.

If the primary boot device fails, you can then boot from this alternate boot device. See “Troubleshooting the System” in *Solstice DiskSuite 4.2.1 User’s Guide* or “Mirroring root (/) Special Considerations” in *Solaris Volume Manager Administration Guide* for more information about alternate boot devices.

```
# ls -l /dev/rdsk/root-disk-slice
```

12. Repeat [Step 1](#) through [Step 11](#) on each remaining node of the cluster.

Ensure that each metadevice or volume name for a mirror is unique throughout the cluster.

13. (Optional) To mirror the global namespace, `/global/.devices/node@nodeid`, go to [“How to Mirror the Global Namespace” on page 140](#).

14. (Optional) To mirror file systems than cannot be unmounted, go to [“How to Mirror File Systems Other Than Root \(/\) That Cannot Be Unmounted”](#) on page 143.
15. (Optional) To mirror user-defined file systems, go to [“How to Mirror File Systems That Can Be Unmounted”](#) on page 147.
16. Go to [“Creating Disk Sets in a Cluster”](#) on page 150 to create a disk set.

Example—Mirroring the Root (/) File System

The following example shows the creation of mirror `d0` on the node `phys-schost-1`, which consists of submirror `d10` on partition `c0t0d0s0` and submirror `d20` on partition `c2t2d0s0`. Device `c2t2d0` is a multihost disk, so the `localonly` property is enabled.

```
(Create the mirror)
# metainit -f d10 1 1 c0t0d0s0
d11: Concat/Stripe is setup
# metainit d20 1 1 c2t2d0s0
d12: Concat/Stripe is setup
# metainit d0 -m d10
d10: Mirror is setup
# metaroot d0
# lockfs -fa

(Move resource groups and device groups from phys-schost-1)
# scswitch -S -h phys-schost-1

(Reboot the node)
# shutdown -g0 -y -i6

(Attach the second submirror)
# metattach d0 d20
d0: Submirror d20 is attached

(Display the device-group node list)
# sccnf -pvv | grep dsk/d2
Device group name:                dsk/d2
...
(dsk/d2) Device group node list:  phys-schost-1, phys-schost-3
...

(Remove phys-schost-3 from the node list)
# sccnf -r -D name=dsk/d2,nodelist=phys-schost-3

(Enable the localonly property)
# sccnf -c -D name=dsk/d2,localonly=true

(Record the alternate boot path)
# ls -l /dev/rdsk/c2t2d0s0
```

```
lrwxrwxrwx 1 root      root          57 Apr 25 20:11 /dev/rdisk/c2t2d0s0
-> ../../devices/node@1/pci@1f,0/pci@1/scsi@3,1/disk@2,0:a,raw
```

▼ How to Mirror the Global Namespace

Use this procedure to mirror the global namespace,
`/global/.devices/node@nodeid/`.

1. **Become superuser on a node of the cluster.**
2. **Put the global namespace slice in a single-slice (one-way) concatenation.**
Use the physical disk name of the disk slice (`cNtXdYsZ`).

```
# metainit -f submirror1 1 1 diskslice
```

3. **Create a second concatenation.**

```
# metainit submirror2 1 1 submirror-diskslice
```

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

```
# metainit mirror -m submirror1
```

Note – The metadevice or volume name for the mirror *must* be unique throughout the cluster.

5. **Attach the second submirror to the mirror.**

This attachment starts a synchronization of the submirrors.

```
# metattach mirror submirror2
```

6. **Edit the `/etc/vfstab` file entry for the `/global/.devices/node@nodeid` file system.**

Replace the names in the device to mount and device to fsck columns with the mirror name.

```
#
vi /etc/vfstab
#device      device      mount      FS      fsck      mount      mount
#to mount    to fsck     point      type    pass     at boot    options
#
/dev/md/dsk/mirror /dev/md/rdisk/mirror /global/.devices/node@nodeid ufs 2 no global
```

7. **Repeat [Step 1](#) through [Step 6](#) on each remaining node of the cluster.**

8. **Wait for the synchronization of the mirrors, started in [Step 5](#), to complete.**

Use the `metastat(1M)` command to view mirror status and to verify that mirror synchronization is complete.

```
# metastat mirror
```

9. If the disk that is used to mirror the global namespace is physically connected to more than one node (multihosted), enable the `localonly` property.

Perform the following steps to enable the `localonly` property of the raw-disk device group for the disk that is used to mirror the global namespace. You must enable the `localonly` property to prevent unintentional fencing of a node from its boot device if the boot device is connected to multiple nodes.

a. If necessary, use the `scdidadm(1M)` command to display the full device-ID path name of the raw-disk device group.

In the following example, the raw-disk device-group name `dsk/d2` is part of the third column of output, which is the full device-ID path name.

```
# scdidadm -L
...
1          phys-schost-3:/dev/rdisk/clt1d0    /dev/did/rdisk/d2
```

b. View the node list of the raw-disk device group.

Output looks similar to the following.

```
# scconf -pvv | grep dsk/d2
Device group name:                dsk/d2
...
(dsk/d2) Device group node list:   phys-schost-1, phys-schost-3
...
```

c. If the node list contains more than one node name, remove all nodes from the node list except the node whose disk is mirrored.

Only the node whose disk is mirrored should remain in the node list for the raw-disk device group.

```
# scconf -r -D name=dsk/dN,nodelist=node

-D name=dsk/dN    Specifies the cluster-unique name of the raw-disk device
                  group
nodelist=node     Specifies the name of the node or nodes to remove from
                  the node list
```

d. Use the `scconf(1M)` command to enable the `localonly` property.

When the `localonly` property is enabled, the raw-disk device group is used exclusively by the node in its node list. This usage prevents unintentional fencing of the node from its boot device if the boot device is connected to multiple nodes.

```
# scconf -c -D name=rawdisk-groupname,localonly=true
```

-D name=*rawdisk-groupname* Specifies the name of the raw-disk device group

For more information about the `localonly` property, see the `scconf_dg_rawdisk(1M)` man page.

10. (Optional) To mirror file systems other than root (/) that cannot be unmounted, go to [“How to Mirror File Systems Other Than Root \(/\) That Cannot Be Unmounted”](#) on page 143.
11. (Optional) To mirror user-defined file systems, go to [“How to Mirror File Systems That Can Be Unmounted”](#) on page 147
12. Go to [“Creating Disk Sets in a Cluster”](#) on page 150 to create a disk set.

Example—Mirroring the Global Namespace

The following example shows creation of mirror `d101`, which consists of submirror `d111` on partition `c0t0d0s3` and submirror `d121` on partition `c2t2d0s3`. The `/etc/vfstab` file entry for `/global/.devices/node@1` is updated to use the mirror name `d101`. Device `c2t2d0` is a multihost disk, so the `localonly` property is enabled.

(Create the mirror)

```
# metainit -f d111 1 1 c0t0d0s3
d111: Concat/Stripe is setup
# metainit d121 1 1 c2t2d0s3
d121: Concat/Stripe is setup
# metainit d101 -m d111
d101: Mirror is setup
# metattach d101 d121
d101: Submirror d121 is attached
```

(Edit the /etc/vfstab file)

```
# vi /etc/vfstab
#device      device      mount      FS      fsck      mount      mount
#to mount    to fsck     point      type     pass      at boot    options
#
/dev/md/dsk/d101 /dev/md/rdisk/d101 /global/.devices/node@1 ufs 2 no global
```

(View the sync status)

```
# metastat d101
d101: Mirror
  Submirror 0: d111
    State: Okay
  Submirror 1: d121
    State: Resyncing
  Resync in progress: 15 % done
...
```

(Identify the device-ID name of the mirrored disk's raw-disk device group)

```

# scdidadm -L
...
1          phys-schost-3:/dev/rdisk/c2t2d0      /dev/did/rdsk/d2

      (Display the device-group node list)
# sccconf -pvv | grep dsk/d2
Device group name:                          dsk/d2
...
      (dsk/d2) Device group node list:        phys-schost-1, phys-schost-3
...

      (Remove phys-schost-3 from the node list)
# sccconf -r -D name=dsk/d2,nodelist=phys-schost-3

      (Enable the localonly property)
# sccconf -c -D name=dsk/d2,localonly=true

```

▼ How to Mirror File Systems Other Than Root (/) That Cannot Be Unmounted

Use this procedure to mirror file systems other than root (/) that cannot be unmounted during normal system usage, such as /usr, /opt, or swap.

1. Become superuser on a node of the cluster.
2. Put the slice on which an unmountable file system resides in a single-slice (one-way) concatenation.

Specify the physical disk name of the disk slice (cNtXdYsZ).

```
# metainit -f submirror1 1 1 diskslice
```

3. Create a second concatenation.

```
# metainit submirror2 1 1 submirror-diskslice
```

4. Create a one-way mirror with one submirror.

```
# metainit mirror -m submirror1
```

Note – The metadvice or volume name for this mirror does *not* need to be unique throughout the cluster.

5. Repeat [Step 1](#) through [Step 4](#) for each remaining unmountable file system that you want to mirror.
6. On each node, edit the /etc/vfstab file entry for each unmountable file system you mirrored.

Replace the names in the device to mount and device to fsck columns with the mirror name.

```
# vi /etc/vfstab
#device      device      mount   FS      fsck    mount   mount
#to mount    to fsck     point   type    pass    at boot options
#
/dev/md/dsk/mirror /dev/md/rdisk/mirror /filesystem ufs 2 no global
```

7. Move any resource groups or device groups from the node.

```
# scswitch -S -h from-node
-S          Moves all resource groups and device groups
-h from-node Specifies the name of the node from which to move resource or
           device groups
```

8. Reboot the node.

```
# shutdown -g0 -y -i6
```

9. Attach the second submirror to each mirror.

This attachment starts a synchronization of the submirrors.

```
# metattach mirror submirror2
```

10. Wait for the synchronization of the mirrors, started in Step 9, to complete.

Use the `metastat(1M)` command to view mirror status and to verify that mirror synchronization is complete.

```
# metastat mirror
```

11. If the disk that is used to mirror the unmountable file system is physically connected to more than one node (multihosted), enable the `localonly` property.

Perform the following steps to enable the `localonly` property of the raw-disk device group for the disk that is used to mirror the unmountable file system. You must enable the `localonly` property to prevent unintentional fencing of a node from its boot device if the boot device is connected to multiple nodes.

a. If necessary, use the `scdidadm -L` command to display the full device-ID path name of the raw-disk device group.

In the following example, the raw-disk device-group name `dsk/d2` is part of the third column of output, which is the full device-ID path name.

```
# scdidadm -L
...
1          phys-schost-3:/dev/rdisk/c1t1d0    /dev/did/rdisk/d2
```

b. View the node list of the raw-disk device group.

Output looks similar to the following.


```
# scconf -pvv | grep dsk/d2
Device group name:                dsk/d2
...
(dsk/d2) Device group node list:  phys-schost-1, phys-schost-3
...
```

- c. **If the node list contains more than one node name, remove all nodes from the node list except the node whose root disk is mirrored.**

Only the node whose root disk is mirrored should remain in the node list for the raw-disk device group.

```
# scconf -r -D name=dsk/dN,nodelist=node
-D name=dsk/dN    Specifies the cluster-unique name of the raw-disk device
                  group
nodelist=node     Specifies the name of the node or nodes to remove from
                  the node list
```

- d. **Use the `scconf(1M)` command to enable the `localonly` property.**

When the `localonly` property is enabled, the raw-disk device group is used exclusively by the node in its node list. This usage prevents unintentional fencing of the node from its boot device if the boot device is connected to multiple nodes.

```
# scconf -c -D name=rawdisk-groupname,localonly=true
-D name=rawdisk-groupname    Specifies the name of the raw-disk device
                              group
```

For more information about the `localonly` property, see the `scconf_dg_rawdisk(1M)` man page.

12. (Optional) To mirror user-defined file systems, go to [“How to Mirror File Systems That Can Be Unmounted”](#) on page 147.

13. Go to [“Creating Disk Sets in a Cluster”](#) on page 150 to create a disk set.

Example—Mirroring File Systems That Cannot Be Unmounted

The following example shows the creation of mirror `d1` on the node `phys-schost-1` to mirror `/usr`, which resides on `c0t0d0s1`. Mirror `d1` consists of submirror `d11` on partition `c0t0d0s1` and submirror `d21` on partition `c2t2d0s1`. The `/etc/vfstab` file entry for `/usr` is updated to use the mirror name `d1`. Device `c2t2d0` is a multihost disk, so the `localonly` property is enabled.

```
(Create the mirror)
# metainit -f d11 1 1 c0t0d0s1
d11: Concat/Stripe is setup
# metainit d21 1 1 c2t2d0s1
```

```

d21: Concat/Stripe is setup
# metainit d1 -m d11
d1: Mirror is setup

    (Edit the /etc/vfstab file)
# vi /etc/vfstab
#device      device      mount   FS    fsck   mount   mount
#to mount    to fsck     point   type  pass   at boot options
#
/dev/md/dsk/d1 /dev/md/rdisk/d1 /usr ufs 2      no global

    (Move resource groups and device groups from phys-schost-1)
# scswitch -S -h phys-schost-1

    (Reboot the node)
# shutdown -g0 -y -i6

    (Attach the second submirror)
# metattach d1 d21
d1: Submirror d21 is attached

    (View the sync status)
# metastat d1
d1: Mirror
    Submirror 0: d11
        State: Okay
    Submirror 1: d21
        State: Resyncing
    Resync in progress: 15 % done
...

    (Identify the device-ID name of the mirrored disk's raw-disk device group)
# scdidadm -L
...
1          phys-schost-3:/dev/rdisk/c2t2d0      /dev/did/rdsk/d2

    (Display the device-group node list)
# scconf -pvv | grep dsk/d2
Device group name:                                dsk/d2
...
(dsk/d2) Device group node list:                  phys-schost-1, phys-schost-3
...

    (Remove phys-schost-3 from the node list)
# scconf -r -D name=dsk/d2,nodelist=phys-schost-3

    (Enable the localonly property)
# scconf -c -D name=dsk/d2,localonly=true

```

▼ How to Mirror File Systems That Can Be Unmounted

Use this procedure to mirror user-defined file systems that can be unmounted. In this procedure, the nodes do not need to be rebooted.

1. Become superuser on a node of the cluster.

2. Unmount the file system to mirror.

Ensure that no processes are running on the file system.

```
# umount /mount-point
```

See the `umount(1M)` man page and “Mounting and Unmounting File Systems” in *System Administration Guide: Basic Administration* for more information.

3. Put in a single-slice (one-way) concatenation the slice that contains a user-defined file system that can be unmounted.

Specify the physical disk name of the disk slice (`cNtXdYsZ`).

```
# metainit -f submirror1 1 1 diskslice
```

4. Create a second concatenation.

```
# metainit submirror2 1 1 submirror-diskslice
```

5. Create a one-way mirror with one submirror.

```
# metainit mirror -m submirror1
```

Note – The metadvice or volume name for this mirror does *not* need to be unique throughout the cluster.

6. Repeat [Step 1](#) through [Step 5](#) for each mountable file system to be mirrored.

7. On each node, edit the `/etc/vfstab` file entry for each file system you mirrored.

Replace the names in the device to mount and device to fsck columns with the mirror name.

```
# vi /etc/vfstab
#device          device          mount   FS    fsck   mount   mount
#to mount        to fsck         point   type  pass   at boot options
#
/dev/md/dsk/mirror /dev/md/rdisk/mirror /filesystem ufs 2 no global
```

8. Attach the second submirror to the mirror.

This attachment starts a synchronization of the submirrors.

```
# metattach mirror submirror2
```

9. Wait for the synchronization of the mirrors, started in Step 8, to be completed.

Use the `metastat(1M)` command to view mirror status.

```
# metastat mirror
```

10. If the disk that is used to mirror the user-defined file system is physically connected to more than one node (multihosted), enable the `localonly` property.

Perform the following steps to enable the `localonly` property of the raw-disk device group for the disk that is used to mirror the user-defined file system. You must enable the `localonly` property to prevent unintentional fencing of a node from its boot device if the boot device is connected to multiple nodes.

a. If necessary, use the `scdidadm -L` command to display the full device-ID path name of the raw-disk device group.

In the following example, the raw-disk device-group name `dsk/d4` is part of the third column of output, which is the full device-ID path name.

```
# scdidadm -L
...
1          phys-schost-3:/dev/rdisk/ctl1d0      /dev/did/rdisk/d2
```

b. View the node list of the raw-disk device group.

Output looks similar to the following.

```
# scconf -pvv | grep dsk/d2
Device group name:                dsk/d2
...
(dsk/d2) Device group node list:  phys-schost-1, phys-schost-3
...
```

c. If the node list contains more than one node name, remove all nodes from the node list except the node whose root disk you mirrored.

Only the node whose root disk you mirrored should remain in the node list for the raw-disk device group.

```
# scconf -r -D name=dsk/dN,nodelist=node

-D name=dsk/dN    Specifies the cluster-unique name of the raw-disk device
                  group
nodelist=node    Specifies the name of the node or nodes to remove from
                  the node list
```

d. Use the `scconf(1M)` command to enable the `localonly` property.

When the `localonly` property is enabled, the raw-disk device group is used exclusively by the node in its node list. This usage prevents unintentional fencing of the node from its boot device if the boot device is connected to multiple nodes.

```
# scconf -c -D name=rawdisk-groupname,localonly=true
```

-D name=*rawdisk-groupname* Specifies the name of the raw-disk device group

For more information about the `localonly` property, see the `scconf_dg_rawdisk(1M)` man page.

11. Mount the mirrored file system.

```
# mount /mount-point
```

See the `mount(1M)` man page and “Mounting and Unmounting File Systems” in *System Administration Guide: Basic Administration* for more information.

12. Create a disk set.

Go to “Creating Disk Sets in a Cluster” on page 150.

Example—Mirroring File Systems That Can Be Unmounted

The following example shows creation of mirror `d4` to mirror `/export`, which resides on `c0t0d0s4`. Mirror `d4` consists of submirror `d14` on partition `c0t0d0s4` and submirror `d24` on partition `c2t2d0s4`. The `/etc/vfstab` file entry for `/export` is updated to use the mirror name `d4`. Device `c2t2d0` is a multihost disk, so the `localonly` property is enabled.

(Unmount the file system)

```
# umount /export
```

(Create the mirror)

```
# metainit -f d14 1 1 c0t0d0s4
d14: Concat/Stripe is setup
# metainit d24 1 1 c2t2d0s4
d24: Concat/Stripe is setup
# metainit d4 -m d14
d4: Mirror is setup
```

(Edit the /etc/vfstab file)

```
# vi /etc/vfstab
#device      device      mount      FS      fsck      mount      mount
#to mount    to fsck     point      type    pass      at boot   options
#
/dev/md/dsk/d4 /dev/md/rdisk/d4 /export ufs 2 no global
```

(Attach the second submirror)

```
# metattach d4 d24
d4: Submirror d24 is attached
```

(View the sync status)

```
# metastat d4
d4: Mirror
    Submirror 0: d14
```

```

    State: Okay
  Submirror 1: d24
    State: Resyncing
  Resync in progress: 15 % done
...

  (Identify the device-ID name of the mirrored disk's raw-disk device group)
# scdidadm -L
...
1          phys-schost-3:/dev/rdisk/c2t2d0      /dev/did/rdsk/d2

  (Display the device-group node list)
# scconf -pvv | grep dsk/d2
Device group name:                dsk/d2
...
(dsk/d2) Device group node list:   phys-schost-1, phys-schost-3
...

  (Remove phys-schost-3 from the node list)
# scconf -r -D name=dsk/d2,nodelist=phys-schost-3

  (Enable the localonly property)
# scconf -c -D name=dsk/d2,localonly=true

  (Mount the file system)
# mount /export

```

Creating Disk Sets in a Cluster

This section describes how to create disk sets for a cluster configuration. You might not need to create disk sets under the following circumstances:

- If you used SunPlex Installer to install Solstice DiskSuite, one to three disk sets might already exist. See [“Using SunPlex Installer to Install Sun Cluster Software” on page 63](#) for information about the metaset that were created by SunPlex Installer.
- To create a Solaris Volume Manager for Sun Cluster disk set for use by Oracle Real Application Clusters, do not use these procedures. Instead, perform the procedures in [“Creating a Multi-Owner Disk Set in Solaris Volume Manager for Sun Cluster for the Oracle Real Application Clusters Database”](#) in *Sun Cluster Data Service for Oracle Real Application Clusters Guide for Solaris OS*.

The following procedures are in this section:

- [“How to Create a Disk Set” on page 151](#)
- [“How to Add Drives to a Disk Set” on page 154](#)
- [“How to Repartition Drives in a Disk Set” on page 155](#)

- “How to Create an `md.tab` File” on page 156
- “How to Activate Metadevices or Volumes” on page 157

▼ How to Create a Disk Set

Perform this procedure to create disk sets.

1. Determine whether the cluster will have more than three disk sets after you create the new disk sets.

- If the cluster will have no more than three disk sets, skip to [Step 2](#).
- If the cluster will have four or more disk sets, perform the following steps to prepare the cluster.

You must perform this task whether you are installing disk sets for the first time or whether you are adding more disk sets to a fully configured cluster.

- a. **On any node of the cluster, check the value of the `md_nsets` variable in the `/kernel/drv/md.conf` file.**
- b. **If the total number of disk sets in the cluster will be greater than the existing value of `md_nsets` minus one, increase the value of `md_nsets` to the desired value.**

The maximum permissible number of disk sets is one less than the configured value of `md_nsets`. The maximum possible value of `md_nsets` is 32, therefore the maximum permissible number of disk sets that you can create is 31.

- c. **Ensure that the `/kernel/drv/md.conf` file is identical on each node of the cluster.**



Caution – Failure to follow this guideline can result in serious Solstice DiskSuite or Solaris Volume Manager errors and possible loss of data.

- d. **If you made changes to the `md.conf` file on any node, perform the following steps to make those changes active.**

- i. **From one node, shut down the cluster.**

```
# scshutdow -g0 -y
```

- ii. **Reboot each node of the cluster.**

```
ok> boot
```

- e. **On each node in the cluster, run the `devfsadm(1M)` command.**

You can run this command on all nodes in the cluster at the same time.

- f. From one node of the cluster, run the `scgdevs(1M)` command to update the `global-devices` namespace.
- g. On each node, verify that the `scgdevs` command has completed processing before you attempt to create any disk sets.

The `scgdevs` command calls itself remotely on all nodes, even when the command is run from just one node. To determine whether the `scgdevs` command has completed processing, run the following command on each node of the cluster.

```
% ps -ef | grep scgdevs
```

2. Ensure that the disk set you intend to create meets one of the following requirements.

- If the disk set is configured with exactly two disk strings, the disk set must connect to exactly two nodes and use exactly two mediator hosts. These mediator hosts must be the same two hosts used for the disk set. See [“Configuring Dual-String Mediators” on page 159](#) for details on how to configure dual-string mediators.
- If the disk set is configured with more than two disk strings, ensure that for any two disk strings S1 and S2, the sum of the number of drives on those strings exceeds the number of drives on the third string S3. Stated as a formula, the requirement is that $\text{count}(S1) + \text{count}(S2) > \text{count}(S3)$.

3. Ensure that the local state database replicas exist.

For instructions, see [“How to Create State Database Replicas” on page 135](#).

4. Become superuser on the cluster node that will master the disk set.

5. Create the disk set.

The following command creates the disk set and registers the disk set as a Sun Cluster disk device group.

```
# metaset -s setname -a -h node1 node2
```

`-s setname` Specifies the disk set name

`-a` Adds (creates) the disk set

`-h node1` Specifies the name of the primary node to master the disk set

`node2` Specifies the name of the secondary node to master the disk set

Note – When you run the `metaset` command to configure a Solstice DiskSuite or Solaris Volume Manager device group on a cluster, the command designates one secondary node by default. You can change the desired number of secondary nodes in the device group by using the `scsetup(1M)` utility after the device group is created. Refer to “Administering Disk Device Groups” in *Sun Cluster System Administration Guide for Solaris OS* for more information about how to change the `numsecondaries` property.

6. Verify the status of the new disk set.

```
# metaset -s setname
```

7. Add drives to the disk set.

Go to “[Adding Drives to a Disk Set](#)” on page 153.

Example—Creating a Disk Set

The following command creates two disk sets, `dg-schost-1` and `dg-schost-2`, with the nodes `phys-schost-1` and `phys-schost-2` specified as the potential primaries.

```
# metaset -s dg-schost-1 -a -h phys-schost-1 phys-schost-2
# metaset -s dg-schost-2 -a -h phys-schost-1 phys-schost-2
```

Adding Drives to a Disk Set

When you add a drive to a disk set, the volume management software repartitions the drive as follows so that the state database for the disk set can be placed on the drive.

- A small portion of each drive is reserved in slice 7 for use by Solstice DiskSuite or Solaris Volume Manager software. The remainder of the space on each drive is placed into slice 0.
- Drives are repartitioned when they are added to the disk set only if slice 7 is not configured correctly.
- Any existing data on the drives is lost by the repartitioning.
- If slice 7 starts at cylinder 0, and the drive partition is large enough to contain a state database replica, the drive is not repartitioned.

▼ How to Add Drives to a Disk Set

1. Become superuser on the node.

2. Ensure that the disk set has been created.

For instructions, see “How to Create a Disk Set” on page 151.

3. List the DID mappings.

```
# sddidadm -L
```

- Choose drives that are shared by the cluster nodes that will master or potentially master the disk set.
- Use the full device-ID path names when you add drives to a disk set.

The first column of output is the DID instance number, the second column is the full physical path name, and the third column is the full device-ID path name (pseudo path). A shared drive has more than one entry for the same DID instance number.

In the following example, the entries for DID instance number 2 indicate a drive that is shared by `phys-schost-1` and `phys-schost-2`, and the full device-ID path name is `/dev/did/rdisk/d2`.

```
1      phys-schost-1:/dev/rdisk/c0t0d0 /dev/did/rdisk/d1
2      phys-schost-1:/dev/rdisk/c1t1d0 /dev/did/rdisk/d2
2      phys-schost-2:/dev/rdisk/c1t1d0 /dev/did/rdisk/d2
3      phys-schost-1:/dev/rdisk/c1t2d0 /dev/did/rdisk/d3
3      phys-schost-2:/dev/rdisk/c1t2d0 /dev/did/rdisk/d3
...
```

4. Take ownership of the disk set.

```
# metaset -s setname -t
```

`-s setname` Specifies the disk set name
`-t` Takes ownership of the disk set

5. Add the drives to the disk set.

Use the full device-ID path name.

```
# metaset -s setname -a drivename
```

`-a` Adds the drive to the disk set
`drivename` Full device-ID path name of the shared drive

Note – Do *not* use the lower-level device name (`cNtXdY`) when you add a drive to a disk set. Because the lower-level device name is a local name and not unique throughout the cluster, using this name might prevent the metaset from being able to switch over.

6. Verify the status of the disk set and drives.

```
# metaset -s setname
```

7. (Optional) To repartition drives for use in metadevices or volumes, go to [“How to Repartition Drives in a Disk Set”](#) on page 155.
8. Go to [“How to Create an md.tab File”](#) on page 156 to define metadevices or volumes by using an `md.tab` file.

Example—Adding Drives to a Disk Set

The `metaset` command adds the drives `/dev/did/rdisk/d1` and `/dev/did/rdisk/d2` to the disk set `dg-schost-1`.

```
# metaset -s dg-schost-1 -a /dev/did/rdisk/d1 /dev/did/rdisk/d2
```

▼ How to Repartition Drives in a Disk Set

The `metaset(1M)` command repartitions drives in a disk set so that a small portion of each drive is reserved in slice 7 for use by Solstice DiskSuite or Solaris Volume Manager software. The remainder of the space on each drive is placed into slice 0. To make more effective use of the drive, use this procedure to modify the disk layout. If you allocate space to slices 1 through 6, you can use these slices when you set up Solstice DiskSuite metadevices or Solaris Volume Manager volumes.

1. Become superuser on the cluster node.
2. Use the `format` command to change the disk partitioning for each drive in the disk set.

When you repartition a drive, you must meet the following conditions to prevent the `metaset(1M)` command from repartitioning the drive.

- Create slice 7 starting at cylinder 0, large enough to hold a state database replica. See your Solstice DiskSuite or Solaris Volume Manager administration guide to determine the size of a state database replica for your version of the volume-manager software.
- Set the `Flag` field in slice 7 to `wu` (read-write, unmountable). Do not set it to `read-only`.
- Do not allow slice 7 to overlap any other slice on the drive.

See the `format(1M)` man page for details.

3. Define metadevices or volumes by using an `md.tab` file.

Go to [“How to Create an md.tab File”](#) on page 156.

▼ How to Create an `md.tab` File

Create an `/etc/lvm/md.tab` file on each node in the cluster. Use the `md.tab` file to define Solstice DiskSuite metadevices or Solaris Volume Manager volumes for the disk sets that you created.

Note – If you are using local metadevices or volumes, ensure that local metadevices or volume names are distinct from the device-ID names used to form disk sets. For example, if the device-ID name `/dev/did/dsk/d3` is used in a disk set, do not use the name `/dev/md/dsk/d3` for a local metadevice or volume. This requirement does not apply to shared metadevices or volumes, which use the naming convention `/dev/md/setname/{r}dsk/d#`.

Tip – To avoid possible confusion between local metadevices or volumes in a cluster environment, use a naming scheme that makes each local metadevice or volume name unique throughout the cluster. For example, for node 1 choose names from `d100-d199`. And for node 2 use `d200-d299`.

1. Become superuser on the cluster node.

2. List the DID mappings for reference when you create your `md.tab` file.

Use the full device-ID path names in the `md.tab` file in place of the lower-level device names (`cNtXdY`).

```
# sddidadm -L
```

In the following example, the first column of output is the DID instance number, the second column is the full physical path name, and the third column is the full device-ID path name (pseudo path).

```
1      phys-schost-1:/dev/rdisk/c0t0d0 /dev/did/rdisk/d1
2      phys-schost-1:/dev/rdisk/c1t1d0 /dev/did/rdisk/d2
2      phys-schost-2:/dev/rdisk/c1t1d0 /dev/did/rdisk/d2
3      phys-schost-1:/dev/rdisk/c1t2d0 /dev/did/rdisk/d3
3      phys-schost-2:/dev/rdisk/c1t2d0 /dev/did/rdisk/d3
...
```

3. Create an `/etc/lvm/md.tab` file and edit it by hand with your preferred text editor.

See your Solstice DiskSuite or Solaris Volume Manager documentation and the `md.tab(4)` man page for details on how to create an `md.tab` file.

Note – If you have existing data on the drives that will be used for the submirrors, you must back up the data before metadevice or volume setup. Then restore the data onto the mirror.

4. Activate the metadevices or volumes that are defined in the `md.tab` files.

Go to “[How to Activate Metadevices or Volumes](#)” on page 157.

Example—Sample `md.tab` File

The following sample `md.tab` file defines the disk set that is named `dg-schost-1`. The ordering of lines in the `md.tab` file is not important.

```
dg-schost-1/d0 -m dg-schost-1/d10 dg-schost-1/d20
dg-schost-1/d10 1 1 /dev/did/rdisk/d1s0
dg-schost-1/d20 1 1 /dev/did/rdisk/d2s0
```

The following example uses Solstice DiskSuite terminology. For Solaris Volume Manager, a trans metadevice is instead called a *transactional volume* and a metadevice is instead called a *volume*. Otherwise, the following process is valid for both volume managers.

The sample `md.tab` file is constructed as follows.

1. The first line defines the device `d0` as a mirror of metadevices `d10` and `d20`. The `-m` signifies that this device is a mirror device.

```
dg-schost-1/d0 -m dg-schost-1/d10 dg-schost-1/d20
```

2. The second line defines metadevice `d10`, the first submirror of `d0`, as a one-way stripe.

```
dg-schost-1/d10 1 1 /dev/did/rdisk/d1s0
```

3. The third line defines metadevice `d20`, the second submirror of `d0`, as a one-way stripe.

```
dg-schost-1/d20 1 1 /dev/did/rdisk/d2s0
```

▼ How to Activate Metadevices or Volumes

Perform this procedure to activate Solstice DiskSuite metadevices or Solaris Volume Manager volumes that are defined in `md.tab` files.

1. **Become superuser on the cluster node.**
2. **Ensure that `md.tab` files are located in the `/etc/lvm` directory.**

3. Ensure that you have ownership of the disk set on the node where the command will be executed.

4. Take ownership of the disk set.

```
# metaset -s setname -t
```

-s *setname* Specifies the disk set name

-t Takes ownership of the disk set

5. Activate the disk set's metadevices or volumes, which are defined in the `md.tab` file.

```
# metainit -s setname -a
```

-a Activates all metadevices in the `md.tab` file

6. For each master and log device, attach the second submirror (*submirror2*).

When the metadevices or volumes in the `md.tab` file are activated, only the first submirror (*submirror1*) of the master and log devices is attached, so *submirror2* must be attached by hand.

```
# metattach mirror submirror2
```

7. Repeat [Step 3](#) through [Step 6](#) for each disk set in the cluster.

If necessary, run the `metainit(1M)` command from another node that has connectivity to the drives. This step is required for cluster-pair topologies, where the drives are not accessible by all nodes.

8. Check the status of the metadevices or volumes.

```
# metastat -s setname
```

See the `metastat(1M)` man page for more information.

9. If your cluster contains disk sets that are configured with exactly two disk enclosures and two nodes, add dual-string mediators.

Go to [“Configuring Dual-String Mediators”](#) on page 159.

10. Go to [“How to Create Cluster File Systems”](#) on page 108 to create a cluster file system.

Example—Activating Metadevices or Volumes in the `md.tab` File

In the following example, all metadevices that are defined in the `md.tab` file for disk set `dg-schost-1` are activated. Then the second submirrors of master device `dg-schost-1/d1` and log device `dg-schost-1/d4` are activated.

```
# metainit -s dg-schost-1 -a
# metattach dg-schost-1/d1 dg-schost-1/d3
# metattach dg-schost-1/d4 dg-schost-1/d6
```

Configuring Dual-String Mediators

This section contains the following information and procedures:

- [“Requirements for Dual-String Mediators” on page 159](#)
- [“How to Add Mediator Hosts” on page 160](#)
- [“How to Check the Status of Mediator Data” on page 160](#)
- [“How to Fix Bad Mediator Data” on page 161](#)

Requirements for Dual-String Mediators

A *dual-string mediator*, or mediator host, is a cluster node that stores mediator data. Mediator data provides information on the location of other mediators and contains a commit count that is identical to the commit count stored in the database replicas. This commit count is used to confirm that the mediator data is in sync with the data in the database replicas.

Dual-string mediators are required for all Solstice DiskSuite or Solaris Volume Manager disk sets that are configured with exactly two disk strings and two cluster nodes. A *disk string* consists of a disk enclosure, its physical drives, cables from the enclosure to the node(s), and the interface adapter cards. The use of mediators enables the Sun Cluster software to ensure that the most current data is presented in the instance of a single-string failure in a dual-string configuration. The following rules apply to dual-string configurations that use mediators.

- Disk sets must be configured with exactly two mediator hosts. Those two mediator hosts must be the same two cluster nodes that are used for the disk set.
- A disk set cannot have more than two mediator hosts.
- Mediators cannot be configured for disk sets that do not meet the two-string and two-host criteria.

These rules do not require that the entire cluster must have exactly two nodes. Rather, only those disk sets that have two disk strings must be connected to exactly two nodes. An N+1 cluster and many other topologies are permitted under these rules.

▼ How to Add Mediator Hosts

Perform this procedure if your configuration requires dual-string mediators.

1. **Become superuser on the node that currently masters the disk set to which you intend to add mediator hosts.**
2. **Run the `metaset(1M)` command to add each node with connectivity to the disk set as a mediator host for that disk set.**

```
# metaset -s setname -a -m mediator-host-list
```

-s *setname* Specifies the disk set name

-a Adds to the disk set

-m *mediator-host-list* Specifies the name of the node to add as a mediator host for the disk set

See the `mediator(7D)` man page for details about mediator-specific options to the `metaset` command.

3. **Check the status of mediator data.**

Go to [“How to Check the Status of Mediator Data”](#) on page 160.

Example—Adding Mediator Hosts

The following example adds the nodes `phys-schost-1` and `phys-schost-2` as mediator hosts for the disk set `dg-schost-1`. Both commands are run from the node `phys-schost-1`.

```
# metaset -s dg-schost-1 -a -m phys-schost-1
# metaset -s dg-schost-1 -a -m phys-schost-2
```

▼ How to Check the Status of Mediator Data

1. **Add mediator hosts as described in [“How to Add Mediator Hosts”](#) on page 160.**

2. **Run the `medstat` command.**

```
# medstat -s setname
```

-s *setname* Specifies the disk set name

See the `medstat(1M)` man page for more information.

3. **If `Bad` is the value in the Status field of the `medstat` output, repair the affected mediator host.**

Go to [“How to Fix Bad Mediator Data”](#) on page 161.

4. Go to [“How to Create Cluster File Systems” on page 108](#) to create a cluster file system.

▼ How to Fix Bad Mediator Data

Perform this procedure to repair bad mediator data.

1. Identify all mediator hosts with bad mediator data as described in the procedure [“How to Check the Status of Mediator Data” on page 160](#).
2. Become superuser on the node that owns the affected disk set.
3. Remove all mediator hosts with bad mediator data from all affected disk sets.

```
# metaset -s setname -d -m mediator-host-list
-s setname           Specifies the disk set name
-d                  Deletes from the disk set
-m mediator-host-list Specifies the name of the node to remove as a
                    mediator host for the disk set
```

4. Restore each mediator host that you removed in [Step 3](#).

```
# metaset -s setname -a -m mediator-host-list
-a                  Adds to the disk set
-m mediator-host-list Specifies the name of the node to add as a mediator
                    host for the disk set
```

See the `mediator(7D)` man page for details about mediator-specific options to the `metaset` command.

5. Create cluster file systems.

Go to [“How to Create Cluster File Systems” on page 108](#).

SPARC: Installing and Configuring VERITAS Volume Manager

Install and configure your local and multihost disks for VERITAS Volume Manager (VxVM) by using the procedures in this chapter, along with the planning information in “Planning Volume Management” on page 33. See your VxVM documentation for additional details.

The following information and procedures are in this chapter:

- “SPARC: Task Map: Installing and Configuring VxVM Software” on page 164
- “SPARC: Setting Up a Root Disk Group Overview” on page 164
- “SPARC: How to Install VERITAS Volume Manager Software and Encapsulate the Root Disk” on page 165
- “SPARC: How to Mirror the Encapsulated Root Disk” on page 168
- “SPARC: How to Install VERITAS Volume Manager Software Only” on page 170
- “SPARC: How to Create a Root Disk Group on a Nonroot Disk” on page 172
- “SPARC: How to Create and Register a Disk Group” on page 173
- “SPARC: How to Assign a New Minor Number to a Disk Device Group” on page 175
- “SPARC: How to Verify the Disk Group Configuration” on page 176
- “SPARC: How to Unencapsulate the Root Disk” on page 177

SPARC: Installing and Configuring VxVM Software

This section provides information and procedures to install and configure VxVM software on a Sun Cluster configuration.

SPARC: Task Map: Installing and Configuring VxVM Software

The following table lists the tasks to perform to install and configure VxVM software for Sun Cluster configurations.

TABLE 4-1 SPARC: Task Map: Installing and Configuring VxVM Software

Task	Instructions
1. Plan the layout of your VxVM configuration.	“Planning Volume Management” on page 33
2. Determine how you will create the root disk group on each node.	“SPARC: Setting Up a Root Disk Group Overview” on page 164
3. For VxVM 3.5 and earlier, install VxVM software and create the root disk group. For VxVM 4.0, the creation of a root disk group is optional.	
<ul style="list-style-type: none"> ■ Method 1 – Install VxVM software and encapsulate the root disk by using the <code>scvxinstall</code> command, and optionally mirror the encapsulated root disk. 	<ol style="list-style-type: none"> 1. “SPARC: How to Install VERITAS Volume Manager Software and Encapsulate the Root Disk” on page 165 2. “SPARC: How to Mirror the Encapsulated Root Disk” on page 168
<ul style="list-style-type: none"> ■ Method 2 – Install VxVM software and, if required, create the root disk group on local, nonroot disks. 	<ol style="list-style-type: none"> 1. “SPARC: How to Install VERITAS Volume Manager Software Only” on page 170 2. “SPARC: How to Create a Root Disk Group on a Nonroot Disk” on page 172
4. Create disk groups and volumes.	“SPARC: How to Create and Register a Disk Group” on page 173
5. If necessary, resolve any minor-number conflicts between disk device groups by assigning a new minor number.	“SPARC: How to Assign a New Minor Number to a Disk Device Group” on page 175
6. Verify the disk groups and volumes.	“SPARC: How to Verify the Disk Group Configuration” on page 176
7. Configure the cluster.	“Configuring the Cluster” on page 108

SPARC: Setting Up a Root Disk Group Overview

For VxVM 4.0, the creation of a root disk group is optional. If you do not intend to create a root disk group, proceed to [“SPARC: How to Install VERITAS Volume Manager Software Only” on page 170](#).

For VxVM 3.5 and earlier, each cluster node requires the creation of a root disk group after VxVM is installed. This disk group is used by VxVM to store configuration information, and has the following restrictions.

- Access to a node's root disk group must be restricted to only that node.
- Remote nodes must never access data stored in another node's root disk group.
- Do not use the `scconf(1M)` command to register the root disk group as a disk device group.
- Whenever possible, configure the root disk group for each node on a nonshared disk.

Sun Cluster software supports the following methods to configure the root disk group.

- **Encapsulate the node's root disk** – This method enables the root disk to be mirrored, which provides a boot alternative if the root disk is corrupted or damaged. To encapsulate the root disk you need two free disk slices as well as free cylinders, preferably at the beginning or the end of the disk.
- **Use local nonroot disks** – This method provides an alternative to encapsulating the root disk. If a node's root disk is encapsulated, certain tasks you might later perform, such as upgrade the Solaris OS or perform disaster recovery procedures, could be more complicated than if the root disk is not encapsulated. To avoid this potential added complexity, you can instead initialize or encapsulate local nonroot disks for use as root disk groups.

A root disk group that is created on local nonroot disks is local to that node, neither globally accessible nor highly available. As with the root disk, to encapsulate a nonroot disk you need two free disk slices as well as free cylinders at the beginning or the end of the disk.

See your VxVM installation documentation for more information.

SPARC: Where to Go From Here

Install VxVM by using one of the following installation methods, depending on how you intend to create the root disk group.

- If you intend to encapsulate the root disk, go to [“SPARC: How to Install VERITAS Volume Manager Software and Encapsulate the Root Disk”](#) on page 165.
- If you intend to create the root disk group on local nonroot disks, go to [“SPARC: How to Install VERITAS Volume Manager Software Only”](#) on page 170.

▼ SPARC: How to Install VERITAS Volume Manager Software and Encapsulate the Root Disk

This procedure uses the `scvxinstall(1M)` command to install VxVM software and encapsulate the root disk in one operation.

Note – If you intend to create the root disk group on local, nonroot disks, go instead to [“SPARC: How to Install VERITAS Volume Manager Software Only” on page 170.](#)

Perform this procedure on each node that you intend to install with VxVM. You can install VERITAS Volume Manager (VxVM) on all nodes of the cluster, or on only those nodes that are physically connected to the storage devices that VxVM will manage.

If you later need to unencapsulate the root disk, follow the procedures in [“SPARC: How to Unencapsulate the Root Disk” on page 177.](#)

1. Ensure that the cluster meets the following prerequisites.

- All nodes in the cluster are running in cluster mode.
- The root disk of the node you install has two free (unassigned) partitions.

2. Have available the following information.

- Mappings of your storage disk drives.
- The following completed configuration planning worksheets. See [“Planning Volume Management” on page 33](#) for planning guidelines.
 - [“Local File System Layout Worksheet” on page 254](#)
 - [“Disk Device Group Configurations Worksheet” on page 260](#)
 - [“Volume-Manager Configurations Worksheet” on page 262](#)

3. Become superuser on a node that you intend to install with VxVM.

4. Insert the VxVM CD-ROM into the CD-ROM drive on the node.

5. Start `scvxinstall` in interactive mode.

Press Ctrl-C at any time to abort the `scvxinstall` command.

```
# scvxinstall
```

See the `scvxinstall(1M)` man page for more information.

6. When prompted whether to encapsulate root, type `yes`.

```
Do you want Volume Manager to encapsulate root [no]? y
```

7. When prompted, provide the location of the VxVM CD-ROM.

- If the appropriate VxVM CD-ROM is found, the location is displayed as part of the prompt within brackets. Press Enter to accept this default location.

```
Where is the volume manager cdrom [default]?
```

- If the VxVM CD-ROM is not found, the prompt is displayed without a default location. Type the location of the CD-ROM or CD-ROM image.

```
Where is the volume manager cdrom?
```

8. When prompted, type your VxVM license key.

Please enter license key: *license*

The `scvxinstall` command automatically performs the following tasks:

- Installs the required VxVM software, licensing, and man-page packages, but does not install the GUI packages
- Selects a cluster-wide `vxio` driver major number
- Creates a root disk group by encapsulating the root disk
- Updates the `/global/.devices` entry in the `/etc/vfstab` file

See the `scvxinstall(1M)` man page for further details.

Note – Two automatic reboots occur during installation. After all installation tasks are completed, `scvxinstall` automatically reboots the node the second time unless you press Ctrl-C when prompted. If you press Ctrl-C to abort the second reboot, you must reboot the node later to complete VxVM installation.

9. If you intend to enable the VxVM cluster feature, supply the cluster feature license key.

See your VxVM documentation for information about how to add a license.

10. (Optional) Install the VxVM GUI.

See your VxVM documentation for information about installing the VxVM GUI.

11. Eject the CD-ROM.

12. Install any VxVM patches.

See “Patches and Required Firmware Levels” in *Sun Cluster 3.1 Release Notes* for the location of patches and installation instructions.

13. (Optional) If you prefer not to have VxVM man pages reside on the cluster node, remove the man-page package.

```
# pkgrm VRTSvmmman
```

14. Repeat Step 3 through Step 13 to install VxVM on additional nodes.

Note – If you intend to enable the VxVM cluster feature, you *must* install VxVM on all nodes of the cluster.

15. If you will *not* install VxVM on one or more nodes, modify the `/etc/name_to_major` file on each non-VxVM node.

a. On a node installed with VxVM, determine the `vxio` major number setting.

```
# grep vxio /etc/name_to_major
```

- b. Become superuser on a node that you do *not* intend to install with VxVM.
- c. Edit the `/etc/name_to_major` file and add an entry to set the `vxio` major number to `NNN`, the number derived in [Step a](#).

```
# vi /etc/name_to_major
vxio NNN
```

- d. Initialize the `vxio` entry.


```
# drvconfig -b -i vxio -m NNN
```
- e. Repeat [Step b](#) through [Step d](#) on all other nodes that you do *not* intend to install with VxVM.

When you finish, each node of the cluster should have the same `vxio` entry in its `/etc/name_to_major` file.

- 16. (Optional) To mirror the encapsulated root disk, go to [“SPARC: How to Mirror the Encapsulated Root Disk”](#) on page 168.
- 17. Go to [“SPARC: How to Create and Register a Disk Group”](#) on page 173.

▼ SPARC: How to Mirror the Encapsulated Root Disk

After you install VxVM and encapsulate the root disk, perform this procedure on each node on which you mirror the encapsulated root disk.

- 1. **Mirror the encapsulated root disk.**

Follow the procedures in your VxVM documentation. For maximum availability and simplified administration, use a local disk for the mirror. See [“Guidelines for Mirroring the Root Disk”](#) on page 39 for additional guidelines.



Caution – Do not use a quorum device to mirror a root disk. Using a quorum device to mirror a root disk might prevent the node from booting from the root-disk mirror under certain circumstances.

- 2. **Display the DID mappings.**

```
# sddidadm -L
```
- 3. **From the DID mappings, locate the disk that is used to mirror the root disk.**
- 4. **Extract the raw-disk device-group name from the device-ID name of the root-disk mirror.**

The name of the raw-disk device group follows the convention `dsk/dN`, where `N` is a number. In the following output, the portion of a `sddidadm` output line from which you extract the raw-disk device-group name is highlighted in bold.


```
N          node:/dev/rdisk/cNtXdY      /dev/did/rdisk/dN
```

5. View the node list of the raw-disk device group.

Output looks similar to the following.

```
# scconf -pvv | grep dsk/dN
Device group name:                dsk/dN
...
(dsk/dN) Device group node list:  phys-schost-1, phys-schost-3
...
```

6. If the node list contains more than one node name, remove from the node list all nodes except the node whose root disk you mirrored.

Only the node whose root disk you mirrored should remain in the node list for the raw-disk device group.

```
# scconf -r -D name=dsk/dN,nodelist=node

-D name=dsk/dN    Specifies the cluster-unique name of the raw-disk device
                  group
nodelist=node     Specifies the name of the node or nodes to remove from the
                  node list
```

7. Enable the `localonly` property of the raw-disk device group.

When the `localonly` property is enabled, the raw-disk device group is used exclusively by the node in its node list. This usage prevents unintentional fencing of the node from its boot device if the boot device is connected to multiple nodes.

```
# scconf -c -D name=dsk/dN,localonly=true
For more information about the localonly property, see the
scconf_dg_rawdisk(1M) man page.
```

8. Repeat this procedure for each node in the cluster whose encapsulated root disk you want to mirror.

9. Create disk groups.

Go to “[SPARC: How to Create and Register a Disk Group](#)” on page 173.

SPARC: Example—Mirroring the Encapsulated Root Disk

The following example shows a mirror created of the root disk for the node `phys-schost-1`. The mirror is created on the disk `c1t1d0`, whose raw-disk device-group name is `dsk/d2`. Disk `c1t1d0` is a multihost disk, so the node `phys-schost-3` is removed from the disk’s node list and the `localonly` property is enabled.

```
(Display the DID mappings)
# scdidadm -L
...
```

```

2         phys-schost-1: /dev/rdisk/c1t1d0 /dev/did/rdisk/d2
2         phys-schost-3: /dev/rdisk/c1t1d0 /dev/did/rdisk/d2
...

        (Display the node list of the mirror disk's raw-disk device group)
# scconf -pvv | grep dsk/d2
Device group name:                dsk/d2
...
        (dsk/d2) Device group node list:    phys-schost-1, phys-schost-3
...

        (Remove phys-schost-3 from the node list)
# scconf -r -D name=dsk/d2,nodelist=phys-schost-3

        (Enable the localonly property)
# scconf -c -D name=dsk/d2,localonly=true

```

▼ SPARC: How to Install VERITAS Volume Manager Software Only

This procedure uses the `scvxinstall` command to install VERITAS Volume Manager (VxVM) software only.

Note – To create the root disk group by encapsulating the root disk, do not use this procedure. Instead, go to [“SPARC: How to Install VERITAS Volume Manager Software and Encapsulate the Root Disk”](#) on page 165 to install VxVM software and encapsulate the root disk in one operation.

Perform this procedure on each node that you want to install with VxVM. You can install VxVM on all nodes of the cluster, or on only those nodes that are physically connected to the storage devices that VxVM will manage.

1. **Ensure that all nodes in the cluster are running in cluster mode.**
2. **Have available any VERITAS Volume Manager (VxVM) license keys you need to install.**
3. **Become superuser on a cluster node that you intend to install with VxVM.**
4. **Insert the VxVM CD-ROM into the CD-ROM drive on the node.**
5. **Start `scvxinstall` in noninteractive installation mode.**

```

# scvxinstall -i -L {license | none}
-i                Installs VxVM but does not encapsulate the root disk

```

-L {*license* | none} Installs the specified *license*. The none argument specifies that no additional license key is being added.

The `scvxinstall` command automatically performs the following tasks:

- Installs the required VxVM software, licensing, and man-page packages, but does not install the GUI packages
- Installs any specified license key
- Selects a cluster-wide `vxio` driver major number

See the `scvxinstall(1M)` man page for information.

6. (Optional) Install the VxVM GUI.

See your VxVM documentation for information about installing the VxVM GUI.

7. Eject the CD-ROM.

8. Install any VxVM patches.

See “Patches and Required Firmware Levels” in *Sun Cluster 3.1 Release Notes* for the location of patches and installation instructions.

9. (Optional) If you prefer not to have VxVM man pages reside on the cluster node, remove the man-page package.

```
# pkgrm VRTSvmmman
```

10. Repeat [Step 3](#) through [Step 9](#) to install VxVM on any additional nodes.

Note – If you intend to enable the VxVM cluster feature, you *must* install VxVM on all nodes of the cluster.

11. If you do *not* install one or more nodes with VxVM, modify the `/etc/name_to_major` file on each non-VxVM node.

a. On a node that is installed with VxVM, determine the `vxio` major number setting.

```
# grep vxio /etc/name_to_major
```

b. Become superuser on a node that you do *not* intend to install with VxVM.

c. Edit the `/etc/name_to_major` file and add an entry to set the `vxio` major number to *NNN*, the number derived in [Step a](#).

```
# vi /etc/name_to_major
vxio NNN
```

d. Initialize the `vxio` entry.

```
# drvconfig -b -i vxio -m NNN
```

- e. Repeat [Step a](#) through [Step c](#) on all other nodes that you do *not* intend to install with VxVM.

When you finish, each node of the cluster should have the same `vxio` entry in its `/etc/name_to_major` file.

12. (Optional) To create a root disk group, go to [“SPARC: How to Create a Root Disk Group on a Nonroot Disk”](#) on page 172.

13. Reboot each node.

```
# shutdown -g0 -y -i6
```

14. Create disk groups.

Go to [“SPARC: How to Create and Register a Disk Group”](#) on page 173.

▼ SPARC: How to Create a Root Disk Group on a Nonroot Disk

Use this procedure to create a root disk group by encapsulating or initializing local disks other than the root disk.

1. Become superuser on the node.
2. (Optional) If the disks are to be encapsulated, ensure that each disk has at least two slices with 0 cylinders.

If necessary, use the `format(1M)` command to assign 0 cylinders to each VxVM slice.

3. Start the `vxinstall` utility.

```
# vxinstall
```

When prompted, make the following choices or entries.

- If you intend to enable the VxVM cluster feature, supply the cluster feature license key.
 - Choose Custom Installation.
 - Do not encapsulate the boot disk.
 - Choose any disks to add to the root disk group.
 - Do not accept automatic reboot.
4. If the root disk group that you created contains one or more disks that connect to more than one node, enable the `localonly` property.

Use the following command to enable the `localonly` property of the `raw-disk` device group for each shared disk in the root disk group.

```
# scconf -c -D name=dsk/dN,localonly=true
```

When the `localonly` property is enabled, the raw-disk device group is used exclusively by the node in its node list. This usage prevents unintentional fencing of the node from the disk that is used by the root disk group if that disk is connected to multiple nodes.

For more information about the `localonly` property, see the `scconf_dg_rawdisk(1M)` man page.

5. Move any resource groups or device groups from the node.

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

-S Moves all resource groups and device groups

-h *from-node* Specifies the name of the node from which to move resource or device groups

6. Reboot the node.

```
# shutdown -g0 -y -i6
```

7. Use the `vxdiskadm` command to add multiple disks to the root disk group.

The root disk group becomes tolerant of a disk failure when it contains multiple disks. See VxVM documentation for procedures.

8. Create disk groups.

Go to “[SPARC: How to Create and Register a Disk Group](#)” on page 173.

▼ SPARC: How to Create and Register a Disk Group

Use this procedure to create your VxVM disk groups and volumes.

Note – After a disk group is registered with the cluster as a disk device group, you should never import or deport a VxVM disk group by using VxVM commands. The Sun Cluster software can handle all cases where disk groups need to be imported or deported. See “Administering Disk Device Groups” in *Sun Cluster System Administration Guide for Solaris OS* for procedures on how to manage Sun Cluster disk device groups.

Perform this procedure from a node that is physically connected to the disks that make the disk group you add.

1. Have available the following information.

- Mappings of your storage disk drives. See the appropriate manual in the *Sun Cluster Hardware Administration Collection* to perform an initial installation of your storage device.

- The following completed configuration planning worksheets.
 - “Local File System Layout Worksheet” on page 254
 - “Disk Device Group Configurations Worksheet” on page 260
 - “Volume-Manager Configurations Worksheet” on page 262

See “Planning Volume Management” on page 33 for planning guidelines.

2. **Become superuser on the node that will have ownership of the disk group.**
3. **Create a VxVM disk group and volume.**

If you are installing Oracle Real Application Clusters, create shared VxVM disk groups by using the cluster feature of VxVM as described in the *VERITAS Volume Manager Administrator’s Reference Guide*. Otherwise, create VxVM disk groups by using the standard procedures that are documented in the VxVM documentation.

Note – You can use Dirty Region Logging (DRL) to decrease volume recovery time if a node failure occurs. However, DRL might decrease I/O throughput.

4. **If the VxVM cluster feature is *not* enabled, register the disk group as a Sun Cluster disk device group.**

If the VxVM cluster feature is enabled, do *not* register a shared disk group as a Sun Cluster disk device group. Instead, go to “SPARC: How to Verify the Disk Group Configuration” on page 176.

- a. **Start the `scsetup(1M)` utility.**

```
# scsetup
```

- b. **Choose Device groups and volumes.**

- c. **Choose Register a VxVM disk group.**

- d. **Follow the instructions to specify the VxVM disk device group that you want to register as a Sun Cluster disk device group.**

- e. **If you encounter the following error message when you attempt to register the disk device group, reinit the disk device group.**

```
scconf: Failed to add device group - in use
```

To reinit the disk device group, use the procedure “SPARC: How to Assign a New Minor Number to a Disk Device Group” on page 175. This procedure enables you to assign a new minor number that does not conflict with a minor number used by existing disk device groups.

- f. **When finished, quit the `scsetup` utility.**

- g. **Verify that the disk device group is registered.**

Look for the disk device information for the new disk that is displayed by the following command.

```
# scstat -D
```

Tip – If you experience a stack overflow when the disk device group is brought online, the default value of the thread stack size might be insufficient. Add the following entry to the `/etc/system` file on each node, where *size* is a number greater than 8000, the default setting:

```
set cl_comm:rm_thread_stacksize=0xsize
```

Note – If you change any configuration information for a VxVM disk group or volume, you must register the configuration changes by using the `scsetup` utility. Configuration changes you must register include adding or removing volumes and changing the group, owner, or permissions of existing volumes. See “Administering Disk Device Groups” in *Sun Cluster System Administration Guide for Solaris OS* for procedures to register configuration changes to a disk device group.

5. Go to [“SPARC: How to Verify the Disk Group Configuration”](#) on page 176.

▼ SPARC: How to Assign a New Minor Number to a Disk Device Group

If disk device group registration fails because of a minor-number conflict with another disk group, you must assign the new disk group a new, unused minor number. Perform this procedure to renumber a disk group.

1. Become superuser on a node of the cluster.
2. Determine the minor numbers in use.
3. Choose any other multiple of 1000 that is not in use to become the base minor number for the new disk group.
4. Assign the new base minor number to the disk group.

```
# vxdg renumber diskgroup base-minor-number
```

5. Go to [“SPARC: How to Create and Register a Disk Group” on page 173](#) of [“SPARC: How to Create and Register a Disk Group” on page 173](#) to register the disk group as a Sun Cluster disk device group.

SPARC: Example—How to Assign a New Minor Number to a Disk Device Group

This example uses the minor numbers 16000-16002 and 4000-4001. The `vx dg` `reminor` command renumbers the new disk device group to use the base minor number 5000.

```
# ls -l /global/.devices/node@1/dev/vx/dsk/*
/global/.devices/node@1/dev/vx/dsk/dg1
brw----- 1 root    root      56,16000 Oct  7 11:32 dg1v1
brw----- 1 root    root      56,16001 Oct  7 11:32 dg1v2
brw----- 1 root    root      56,16002 Oct  7 11:32 dg1v3

/global/.devices/node@1/dev/vx/dsk/dg2
brw----- 1 root    root      56,4000 Oct  7 11:32 dg2v1
brw----- 1 root    root      56,4001 Oct  7 11:32 dg2v2
# vx dg reminor dg3 5000
```

▼ SPARC: How to Verify the Disk Group Configuration

Perform this procedure on each node of the cluster.

1. Verify that only the local disks are included in the root disk group, and disk groups are imported on the current primary node only.

```
# vx disk list
```

2. Verify that all volumes have been started.

```
# vx print
```

3. Verify that all disk groups have been registered as Sun Cluster disk device groups and are online.

```
# scstat -D
```

4. Configure the cluster.

Go to [“Configuring the Cluster” on page 108](#).

▼ SPARC: How to Unencapsulate the Root Disk

Perform this procedure to unencapsulate the root disk.

1. Ensure that only Solaris root file systems are present on the root disk.

The Solaris root file systems are root (/), swap, the global devices namespace, /usr, /var, /opt, and /home. If any other file systems reside on the root disk, back them up and remove them from the root disk.

2. Become superuser on the node that you intend to unencapsulate.

3. Move all resource groups and device groups from the node.

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

-S Moves all resource groups and device groups

-h *from-node* Specifies the name of the node from which to move resource or device groups

4. Determine the node-ID number of the node.

```
# clinfo -nN
```

5. Unmount the global-devices file system for this node, where *N* is the node ID number that is returned in [Step 4](#).

```
# umount /global/.devices/node@N
```

6. View the /etc/vfstab file and determine which VxVM volume corresponds to the global-devices file system.

```
# vi /etc/vfstab
```

#device	device	mount	FS	fck	mount	mount
#to	to	point	type	pass	at	options
#						
#NOTE:	volume	rootdiskxNvol	(/global/.devices/node@N)	encapsulated		
#partition	cNtXdYsZ					

7. Remove the VxVM volume that corresponds to the global-devices file system from the root disk group.

- For VxVM 3.5 and earlier, use the following command:

```
# vxedit -rf rm rootdiskxNvol
```

- For VxVM 4.0, use the following command:

```
# vxedit -g rootdiskgroup
```



Caution – Do not store data other than device entries for global devices in the global-devices file system. All data in the global-devices file system is destroyed when you remove the VxVM volume. Only data that is related to global devices entries is restored after the root disk is unencapsulated.

8. Unencapsulate the root disk.

Note – Do **not** accept the shutdown request from the command.

```
# /etc/vx/bin/vxunroot
```

See your VxVM documentation for details.

9. Use the `format(1M)` command to add a 512-Mbyte partition to the root disk to use for the global-devices file system.

Tip – Use the same slice that was allocated to the global-devices file system before the root disk was encapsulated, as specified in the `/etc/vfstab` file.

10. Set up a file system on the partition that you created in [Step 9](#).

```
# newfs /dev/rdisk/cNtXdYsZ
```

11. Determine the DID name of the root disk.

```
# scdidadm -l cNtXdY
1      phys-schost-1:/dev/rdisk/cNtXdY    /dev/did/rdisk/dN
```

12. In the `/etc/vfstab` file, replace the path names in the global-devices file system entry with the DID path that you identified in [Step 11](#).

The original entry would look similar to the following.

```
# vi /etc/vfstab
/dev/vx/dsk/rootdiskxNvol /dev/vx/rdisk/rootdiskxNvol /global/.devices/node@N ufs 2 no global
The revised entry that uses the DID path would look similar to the following.
/dev/did/dsk/dNsX /dev/did/rdisk/dNsX /global/.devices/node@N ufs 2 no global
```

13. Mount the global-devices file system.

```
# mount /global/.devices/node@N
```

14. From one node of the cluster, repopulate the global-devices file system with device nodes for any raw-disk devices and Solstice DiskSuite or Solaris Volume Manager devices.

```
# scgdevs
```

VxVM devices are recreated during the next reboot.

15. Reboot the node.

```
# reboot
```

16. Repeat this procedure on each node of the cluster to unencapsulate the root disk on those nodes.

Upgrading Sun Cluster Software

This chapter provides the following information and procedures to upgrade a Sun Cluster 3.x configuration to Sun Cluster 3.1 9/04 software:

- “Upgrade Requirements and Support Guidelines” on page 182
- “Choosing a Sun Cluster Upgrade Method” on page 182
- “Task Map: Upgrading to Sun Cluster 3.1 9/04 Software (Nonrolling)” on page 184
- “How to Prepare the Cluster for a Nonrolling Upgrade” on page 184
- “How to Perform a Nonrolling Upgrade of the Solaris OS” on page 189
- “How to Perform a Nonrolling Upgrade of Sun Cluster 3.1 9/04 Software” on page 192
- “How to Finish a Nonrolling Upgrade to Sun Cluster 3.1 9/04 Software” on page 198
- “Task Map: Upgrading to Sun Cluster 3.1 9/04 Software (Rolling)” on page 201
- “How to Prepare a Cluster Node for a Rolling Upgrade” on page 202
- “How to Perform a Rolling Upgrade of a Solaris Maintenance Update” on page 206
- “How to Perform a Rolling Upgrade of Sun Cluster 3.1 9/04 Software” on page 207
- “How to Finish a Rolling Upgrade to Sun Cluster 3.1 9/04 Software” on page 212
- “How to Handle Storage Reconfiguration During an Upgrade” on page 215
- “How to Resolve Mistaken Storage Changes During an Upgrade” on page 216
- “SPARC: How to Upgrade Sun Cluster-Module Software for Sun Management Center” on page 217
- “SPARC: How to Upgrade Sun Management Center Software” on page 218

Overview of Upgrading a Sun Cluster Configuration

This section provides the following guidelines to upgrade a Sun Cluster configuration:

- “Upgrade Requirements and Support Guidelines” on page 182
- “Choosing a Sun Cluster Upgrade Method” on page 182

Upgrade Requirements and Support Guidelines

Observe the following requirements and support guidelines when you upgrade to Sun Cluster 3.1 9/04 software:

- The cluster must run on or be upgraded to at least Solaris 8 2/02 software, including the most current required patches.
- The cluster hardware must be a supported configuration for Sun Cluster 3.1 9/04 software. Contact your Sun representative for information about current supported Sun Cluster configurations.
- You must upgrade all software to a version that is supported by Sun Cluster 3.1 9/04 software. For example, if a data service is supported on Sun Cluster 3.0 software but is not supported on Sun Cluster 3.1 9/04 software, you must upgrade that data service to the version of that data service that is supported on Sun Cluster 3.1 9/04 software. See “Supported Products” in *Sun Cluster 3.1 9/04 Release Notes for Solaris OS* for support information about specific data services.
- If the related application of a data service is not supported on Sun Cluster 3.1 9/04 software, you must upgrade that application to a supported release.
- The `scinstall` upgrade utility only upgrades those data services that are provided with Sun Cluster 3.1 9/04 software. You must manually upgrade any custom or third-party data services.
- For upgrade from a Sun Cluster 3.0 release, have available the test IP addresses to use with your public-network adapters when NAFO groups are converted to Internet Protocol (IP) Network Multipathing groups. The `scinstall` upgrade utility prompts you for a test IP address for each public-network adapter in the cluster. A test IP address must be on the same subnet as the primary IP address for the adapter.
See the IP Network Multipathing Administration Guide (Solaris 8) or System Administration Guide: IP Services (Solaris 9) for information about test IP addresses for IP Network Multipathing groups.
- Sun Cluster 3.1 9/04 software supports only nonrolling upgrade from Solaris 8 software to Solaris 9 software.
- Sun Cluster 3.1 9/04 software supports direct upgrade only from Sun Cluster 3.x software.
- Sun Cluster 3.1 9/04 software does not support any downgrade of Sun Cluster software.
- Sun Cluster 3.1 9/04 software does not support upgrade between architectures.
- Sun Cluster 3.1 9/04 software does not support the Live Upgrade method to upgrade Solaris software in a Sun Cluster configuration.

Choosing a Sun Cluster Upgrade Method

Choose one of the following methods to upgrade your cluster to Sun Cluster 3.1 9/04 software:

- **Nonrolling upgrade** – In a *nonrolling upgrade*, you shut down the cluster before you upgrade the cluster nodes. You return the cluster to production after all nodes are fully upgraded. You **must** use the nonrolling-upgrade method if one or more of the following conditions apply:
 - You are upgrading from Sun Cluster 3.0 software.
 - You are upgrading from Solaris 8 software to Solaris 9 software.
 - Any software products that you are upgrading, such as applications or databases, require that the same version of the software is running on all cluster nodes at the same time.
 - You are upgrading the Sun Cluster-module software for Sun Management Center.
 - You are also upgrading VxVM or VxFS.
- **Rolling upgrade** – In a *rolling upgrade*, you upgrade one node of the cluster at a time. The cluster remains in production with services running on the other nodes. You can use the rolling-upgrade method only if **all** of the following conditions apply:
 - You are upgrading from Sun Cluster 3.1 software.
 - You are upgrading the Solaris operating system only to a Solaris update, if at all.
 - For any applications or databases you must upgrade, the current version of the software can coexist in a running cluster with the upgrade version of that software.

If your cluster configuration meets the requirements to perform a rolling upgrade, you can still choose to perform a nonrolling upgrade instead. A nonrolling upgrade might be preferable to a rolling upgrade if you wanted to use the Cluster Control Panel to issue commands to all cluster nodes at the same time and you could tolerate the cluster downtime.

For overview information about planning your Sun Cluster 3.1 9/04 configuration, see [Chapter 1](#).

Upgrading to Sun Cluster 3.1 9/04 Software (Nonrolling)

Follow the tasks in this section to perform a nonrolling upgrade from Sun Cluster 3.x software to Sun Cluster 3.1 9/04 software. In a nonrolling upgrade, you shut down the entire cluster before you upgrade the cluster nodes. This procedure also enables you to upgrade the cluster from Solaris 8 software to Solaris 9 software.

Note – To perform a rolling upgrade to Sun Cluster 3.1 9/04 software, instead follow the procedures in [“Upgrading to Sun Cluster 3.1 9/04 Software \(Rolling\)”](#) on page 201.

Task Map: Upgrading to Sun Cluster 3.1 9/04 Software (Nonrolling)

TABLE 5–1 Task Map: Upgrading to Sun Cluster 3.1 9/04 Software

Task	Instructions
1. Read the upgrade requirements and restrictions.	“Upgrade Requirements and Support Guidelines” on page 182
2. Remove the cluster from production, disable resources, and back up shared data and system disks. If the cluster uses dual-string mediators for Solstice DiskSuite or Solaris Volume Manager software, unconfigure the mediators.	“How to Prepare the Cluster for a Nonrolling Upgrade” on page 184
3. Upgrade the Solaris software, if necessary, to a supported Solaris update. Optionally, upgrade VERITAS Volume Manager (VxVM).	“How to Perform a Nonrolling Upgrade of the Solaris OS” on page 189
4. Upgrade to Sun Cluster 3.1 9/04 framework and data-service software. If necessary, upgrade applications. If the cluster uses dual-string mediators, reconfigure the mediators. SPARC: If you upgraded VxVM, upgrade disk groups.	“How to Perform a Nonrolling Upgrade of Sun Cluster 3.1 9/04 Software” on page 192
5. Enable resources and bring resource groups online. Optionally, migrate existing resources to new resource types.	“How to Finish a Nonrolling Upgrade to Sun Cluster 3.1 9/04 Software” on page 198
6. <i>(Optional)</i> SPARC: Upgrade the Sun Cluster module for Sun Management Center, if needed.	“SPARC: How to Upgrade Sun Cluster-Module Software for Sun Management Center” on page 217

▼ How to Prepare the Cluster for a Nonrolling Upgrade

Before you upgrade the software, perform the following steps to remove the cluster from production:

- 1. Ensure that the configuration meets requirements for upgrade.**
See [“Upgrade Requirements and Support Guidelines”](#) on page 182.
- 2. Have available the CD-ROMs, documentation, and patches for all software products you are upgrading.**

- Solaris 8 or Solaris 9 OS
- Sun Cluster 3.1 9/04 framework
- Sun Cluster 3.1 9/04 data services (agents)
- Applications that are managed by Sun Cluster 3.1 9/04 data-service agents
- SPARC: VERITAS Volume Manager

See “Patches and Required Firmware Levels” in *Sun Cluster 3.1 Release Notes* for the location of patches and installation instructions.

3. (Optional) Install Sun Cluster 3.1 9/04 documentation.

Install the documentation packages on your preferred location, such as an administrative console or a documentation server. See the `index.html` file at the top level of the Sun Cluster 3.1 9/04 CD-ROM to access installation instructions.

4. If you are upgrading from Sun Cluster 3.0 software, have available your list of test IP addresses.

Each public-network adapter in the cluster must have at least one test IP address. This requirement applies regardless of whether the adapter is the active adapter or the backup adapter in the group. The test IP addresses are used to reconfigure the adapters to use IP Network Multipathing.

Note – Each test IP address must be on the same subnet as the existing IP address that is used by the public-network adapter.

To list the public-network adapters on a node, run the following command:

```
% pnmstat
```

See the *IP Network Multipathing Administration Guide* (Solaris 8) or *System Administration Guide: IP Services* (Solaris 9) for more information about test IP addresses for IP Network Multipathing.

5. Notify users that cluster services will be unavailable during the upgrade.

6. Ensure that the cluster is functioning normally.

- To view the current status of the cluster, run the following command from any node:

```
% scstat
```

See the `scstat(1M)` man page for more information.

- Search the `/var/adm/messages` log on the same node for unresolved error messages or warning messages.
- Check the volume-manager status.

7. Become superuser on a node of the cluster.

8. Start the `scsetup(1m)` utility.

```
# scsetup
```

The Main Menu displays.

9. Switch each resource group offline.

- a. From the `scsetup` Main Menu, choose Resource groups.
- b. From the Resource Group Menu, choose Online/Offline or Switchover a resource group.
- c. Follow the prompts to take offline all resource groups and to put them in the unmanaged state.
- d. When all resource groups are offline, type `q` to return to the Resource Group Menu.

10. Disable all resources in the cluster.

The disabling of resources before upgrade prevents the cluster from bringing the resources online automatically if a node is mistakenly rebooted into cluster mode.

- a. From the Resource Group Menu, choose Enable/Disable a resource.
- b. Choose a resource to disable and follow the prompts.
- c. Repeat [Step b](#) for each resource.
- d. When all resources are disabled, type `q` to return to the Resource Group Menu.

11. Exit the `scsetup` utility.

Type `q` to back out of each submenu or press `Ctrl-C`.

12. Verify that all resources on all nodes are `Offline` and that all resource groups are in the `Unmanaged` state.

```
# scstat -g
```

13. If your cluster uses dual-string mediators for Solstice DiskSuite or Solaris Volume Manager software, unconfigure your mediators.

See [“Configuring Dual-String Mediators”](#) on page 159 for more information.

a. Run the following command to verify that no mediator data problems exist.

```
# medstat -s setname
```

`-s setname` Specifies the disk set name

If the value in the Status field is `Bad`, repair the affected mediator host. Follow the procedure [“How to Fix Bad Mediator Data”](#) on page 161.

b. List all mediators.

Save this information for when you restore the mediators during the procedure [“How to Finish a Nonrolling Upgrade to Sun Cluster 3.1 9/04 Software”](#) on page 198.

- c. For a disk set that uses mediators, take ownership of the disk set if no node already has ownership.

```
# metaset -s setname -t
-t          Takes ownership of the disk set
```

- d. Unconfigure all mediators for the disk set.

```
# metaset -s setname -d -m mediator-host-list
-s setname          Specifies the disk set name
-d                  Deletes from the disk set
-m mediator-host-list Specifies the name of the node to remove as a
                    mediator host for the disk set
```

See the `mediator(7D)` man page for further information about mediator-specific options to the `metaset` command.

- e. Repeat **Step c** through **Step d** for each remaining disk set that uses mediators.

14. If not already installed, install Sun Web Console packages.

Perform this step on each node of the cluster. These packages are required by Sun Cluster software, even if you do not use Sun Web Console.

- a. Insert the Sun Cluster 3.1 9/04 CD-ROM in the CD-ROM drive.
- b. Change to the
`/cdrom/cdrom0/Solaris_arch/Product/sun_web_console/2.1/`
directory, where *arch* is `sparc` or `x86`.
- c. Run the `setup` command.

```
# ./setup
The setup command installs all packages to support Sun Web Console.
```

15. For a two-node cluster, if the cluster uses Sun StorEdge Availability Suite software, ensure that the configuration data for availability services resides on the quorum disk.

The configuration data must reside on a quorum disk to ensure the proper functioning of Sun StorEdge Availability Suite after you upgrade the cluster software.

- a. Become superuser on a node of the cluster that runs Sun StorEdge Availability Suite software.
- b. Identify the device ID and the slice that is used by the Sun StorEdge Availability Suite configuration file.

```
# /usr/opt/SUNWscm/sbin/dscfg
/dev/did/rdisk/dNsS
```

In this example output, *N* is the device ID and *S* the slice of device *N*.

c. Identify the existing quorum device.

```
# scstat -q
-- Quorum Votes by Device --
                Device Name                Present Possible Status
                -----                -
Device votes:  /dev/did/rdisk/dQsS  1          1          Online
```

In this example output, *dQsS* is the existing quorum device.

d. If the quorum device is not the same as the Sun StorEdge Availability Suite configuration-data device, move the configuration data to an available slice on the quorum device.

```
# dd if='/usr/opt/SUNWesm/sbin/dscfg' of=/dev/did/rdisk/dQsS
```

Note – You must use the name of the raw DID device, */dev/did/rdisk/*, not the block DID device, */dev/did/dsk/*.

e. If you moved the configuration data, configure Sun StorEdge Availability Suite software to use the new location.

As superuser, issue the following command on each node that runs Sun StorEdge Availability Suite software.

```
# /usr/opt/SUNWesm/sbin/dscfg -s /dev/did/rdisk/dQsS
```

16. Stop all applications that are running on each node of the cluster.

17. Ensure that all shared data is backed up.

18. From one node, shut down the cluster.

```
# scshutdown -g0 -y
```

See the *scshutdown(1M)* man page for more information.

19. Boot each node into noncluster mode.

On SPARC based systems, perform the following command:

```
ok boot -x
```

On x86 based systems, perform the following commands:

```
...
                <<< Current Boot Parameters >>>
Boot path: /pci@0,0/pci-ide@7,1/ata@1/cmdk@0,0:b
Boot args:

Type  b [file-name] [boot-flags] <ENTER>  to boot with options
or    i <ENTER>                             to enter boot interpreter
or    <ENTER>                               to boot with defaults
```

```
<<< timeout in 5 seconds >>>
Select (b)oot or (i)nterpreter: b -x
```

20. Ensure that each system disk is backed up.

21. Upgrade the Sun Cluster software or the Solaris operating system.

- To upgrade Solaris software before you perform Sun Cluster upgrade, go to [“How to Perform a Nonrolling Upgrade of the Solaris OS”](#) on page 189.

If Sun Cluster 3.1 9/04 software does not support the release of the Solaris OS that you currently run on your cluster, you must upgrade the Solaris software to a supported release. If Sun Cluster 3.1 9/04 software supports the release of the Solaris OS that you currently run on your cluster, further Solaris software upgrade is optional. See “Supported Products” in *Sun Cluster Release Notes for Solaris OS* for more information.

- To upgrade Sun Cluster software, go to [“How to Perform a Nonrolling Upgrade of Sun Cluster 3.1 9/04 Software”](#) on page 192.

▼ How to Perform a Nonrolling Upgrade of the Solaris OS

Perform this procedure on each node in the cluster to upgrade the Solaris OS. If the cluster already runs on a version of the Solaris OS that supports Sun Cluster 3.1 9/04 software, further upgrade of the Solaris OS is optional. If you do not intend to upgrade the Solaris OS, go to [“How to Perform a Nonrolling Upgrade of Sun Cluster 3.1 9/04 Software”](#) on page 192.

Note – The cluster must already run on, or be upgraded to, at least the minimum required level of the Solaris 8 or Solaris 9 OS to support Sun Cluster 3.1 9/04 software. See “Supported Products” in *Sun Cluster Release Notes for Solaris OS* for more information.

1. **Ensure that all steps in [“How to Prepare the Cluster for a Nonrolling Upgrade”](#) on page 184 are completed.**
2. **Become superuser on the cluster node to upgrade.**
3. **(Optional) Upgrade VxFS.**
Follow procedures that are provided in your VxFS documentation.
4. **Determine whether the following Apache links already exist, and if so, whether the file names contain an uppercase K or S:**

```

/etc/rc0.d/K16apache
/etc/rc1.d/K16apache
/etc/rc2.d/K16apache
/etc/rc3.d/S50apache
/etc/rcS.d/K16apache

```

- If these links already exist and do contain an uppercase K or S in the file name, no further action is necessary for these links.
- If these links do not exist, or if these links exist but instead contain a lowercase k or s in the file name, you move aside these links in [Step 9](#).

5. Comment out all entries for globally mounted file systems in the node's `/etc/vfstab` file.

- a. For later reference, make a record of all entries that are already commented out.
- b. Temporarily comment out all entries for globally mounted file systems in the `/etc/vfstab` file.

Entries for globally mounted file systems contain the `global` mount option. Comment out these entries to prevent the Solaris upgrade from attempting to mount the global devices.

6. Determine which procedure to follow to upgrade the Solaris OS.

Volume Manager	Procedure to Use	Location of Instructions
Solstice DiskSuite or Solaris Volume Manager	Any Solaris upgrade method <i>except</i> the Live Upgrade method	Solaris 8 or Solaris 9 installation documentation
SPARC: VERITAS Volume Manager	"Upgrading VxVM and Solaris"	VERITAS Volume Manager installation documentation

Note – If your cluster has VxVM installed, you must reinstall the existing VxVM software or upgrade to the Solaris 9 version of VxVM software as part of the Solaris upgrade process.

7. Upgrade the Solaris software, following the procedure that you selected in [Step 6](#).

- a. When you are instructed to reboot a node during the upgrade process, always add the `-x` option to the command. Or, if the instruction says to run the `init S` command, use the `reboot -- -xs` command instead.

The `-x` option ensures that the node reboots into noncluster mode. For example, either of the following two commands boot a node into single-user noncluster mode:

- On SPARC based systems, perform the following commands:

```
# reboot -- -xs
ok boot -xs
```

- On x86 based systems, perform the following commands:

```
# reboot -- -xs
...
<<< Current Boot Parameters >>>
Boot path: /pci@0,0/pci-ide@7,1/ata@1/cmdk@0,0:b
Boot args:

Type  b [file-name] [boot-flags] <ENTER>  to boot with options
or    i <ENTER>                          to enter boot interpreter
or    <ENTER>                             to boot with defaults

<<< timeout in 5 seconds >>>
Select (b)oot or (i)nterpreter: b -xs
```

- Do not perform the final reboot instruction in the Solaris software upgrade. Instead, return to this procedure to perform [Step 8](#) and [Step 9](#), then reboot into noncluster mode in [Step 10](#) to complete Solaris software upgrade.**

- In the `/a/etc/vfstab` file, uncomment those entries for globally mounted file systems that you commented out in [Step 5](#).**

- Move aside restored Apache links if either of the following conditions was true *before* you upgraded the Solaris software:**

- The Apache links listed in [Step 4](#) did not exist.
- The Apache links listed in [Step 4](#) existed *and* contained a lowercase k or s in the file names.

To move aside restored Apache links, which contain an uppercase K or S in the name, use the following commands to rename the files with a lowercase k or s.

```
# mv /a/etc/rc0.d/K16apache /a/etc/rc0.d/k16apache
# mv /a/etc/rc1.d/K16apache /a/etc/rc1.d/k16apache
# mv /a/etc/rc2.d/K16apache /a/etc/rc2.d/k16apache
# mv /a/etc/rc3.d/S50apache /a/etc/rc3.d/s50apache
# mv /a/etc/rcS.d/K16apache /a/etc/rcS.d/k16apache
```

- Reboot the node into noncluster mode.**

Include the double dashes (--) in the following command:

```
# reboot -- -x
```

- SPARC: If your cluster runs VxVM, perform the remaining steps in the procedure “Upgrading VxVM and Solaris” to reinstall or upgrade VxVM.**

Note the following special instructions:

- After VxVM upgrade is complete but before you reboot, verify the entries in the `/etc/vfstab` file. If any of the entries that you uncommented in [Step 8](#)**

were commented out, make those entries uncommented again.

- b. When the VxVM procedures instruct you to perform a final reconfiguration reboot by using the `-r` option, reboot into noncluster mode by using the `-x` option instead.

```
# reboot -- -x
```

Note – If you see a message similar to the following, type the root password to continue upgrade processing. Do *not* run the `fsck` command nor type Ctrl-D.

```
WARNING - Unable to repair the /global/.devices/node@1 filesystem.  
Run fsck manually (fsck -F ufs /dev/vx/rdisk/rootdisk_13vol). Exit the  
shell when done to continue the boot process.
```

```
Type control-d to proceed with normal startup,  
(or give root password for system maintenance):
```

Type the root password

12. Install any required Solaris software patches and hardware-related patches, and download any needed firmware that is contained in the hardware patches.

For Solstice DiskSuite software (Solaris 8), also install any Solstice DiskSuite software patches.

Note – Do not reboot after you add patches. Wait to reboot the node until after you upgrade the Sun Cluster software.

See “Patches and Required Firmware Levels” in *Sun Cluster 3.1 Release Notes* for the location of patches and installation instructions.

13. Upgrade to Sun Cluster 3.1 9/04 software.

Go to “How to Perform a Nonrolling Upgrade of Sun Cluster 3.1 9/04 Software” on page 192.

Note – To complete the upgrade from Solaris 8 to Solaris 9 software, you must also upgrade to the Solaris 9 version of Sun Cluster 3.1 9/04 software, even if the cluster already runs on the Solaris 8 version of Sun Cluster 3.1 9/04 software.

▼ How to Perform a Nonrolling Upgrade of Sun Cluster 3.1 9/04 Software

Perform this procedure to upgrade each node of the cluster to Sun Cluster 3.1 9/04 software. You must also perform this procedure to complete cluster upgrade from Solaris 8 to Solaris 9 software.

Tip – You can perform this procedure on more than one node at the same time.

1. Ensure that all steps in “How to Prepare the Cluster for a Nonrolling Upgrade” on page 184 are completed.

2. If you upgraded from Solaris 8 to Solaris 9 software, ensure that all steps in “How to Perform a Nonrolling Upgrade of the Solaris OS” on page 189 are completed.

3. Ensure that you have installed all required Solaris software patches and hardware-related patches.

For Solstice DiskSuite software (Solaris 8), also ensure that you have installed all required Solstice DiskSuite software patches.

4. Become superuser on a node of the cluster.

5. Insert the Sun Java Enterprise System 1/05 2 of 2 CD-ROM into the CD-ROM drive on the node.

If the volume management daemon `vold(1M)` is running and is configured to manage CD-ROM devices, the daemon automatically mounts the CD-ROM on the `/cdrom/cdrom0/` directory.

6. On the Sun Cluster 3.1 9/04 CD-ROM, change to the `Solaris_arch/Product/sun_cluster/Solaris_ver/Tools/` directory, where `arch` is `sparc` or `x86` and where `ver` is 8 (for Solaris 8) or 9 (for Solaris 9) .

```
# cd /cdrom/cdrom0/Solaris_arch/Product/sun_cluster/Solaris_ver/Tools
```

7. Upgrade the cluster framework software.

Note – Do not use the `/usr/cluster/bin/scinstall` command that is already installed on the node. You must use the `scinstall` command that is on the Sun Cluster 3.1 9/04 CD-ROM.

■ To upgrade from Sun Cluster 3.0 software, run the following command:

```
# ./scinstall -u update -S interact [-M patchdir=dirname]
```

`-S`

Specifies the test IP addresses to use to convert NAFO groups to IP Network Multipathing groups

`interact`

Specifies that `scinstall` prompts the user for each test IP address needed

`-M patchdir=dirname [, patchlistfile=filename]`

Specifies the path to patch information so that the specified patches can be installed using the `scinstall` command. If you do not specify a patch-list

file, the `scinstall` command installs all the patches in the directory `dirname`, including tarred, jarred, and zipped patches.

The `-M` option is not required. You can use any method you prefer for installing patches.

- To upgrade from Sun Cluster 3.1 software, run the following command:

```
# ./scinstall -u update [-M patchdir=dirname]
```

```
-M patchdir=dirname [, patchlistfile=filename]
```

Specifies the path to patch information so that the specified patches can be installed by the `scinstall` command. If you do not specify a patch-list file, the `scinstall` command installs all the patches in the directory `dirname`, including tarred, jarred, and zipped patches.

The `-M` option is not required. You can use any method you prefer for installing patches.

See “Patches and Required Firmware Levels” in *Sun Cluster 3.1 Release Notes* for the location of patches and installation instructions.

Note – Sun Cluster 3.1 9/04 software requires at least version 3.5.1 of Sun Explorer software. Upgrading to Sun Cluster software includes installing Sun Explorer data collector software, to be used in conjunction with the `sccheck` utility. If another version of Sun Explorer software was already installed before the Sun Cluster upgrade, it is replaced by the version that is provided with Sun Cluster software. Options such as user identity and data delivery are preserved, but `crontab` entries must be manually re-created.

During Sun Cluster upgrade, `scinstall` might make one or more of the following configuration changes:

- Convert NAFO groups to IP Network Multipathing groups but keep the original NAFO-group name.
See the `scinstall(1M)` man page for more information. See the *IP Network Multipathing Administration Guide* (Solaris 8) or *System Administration Guide: IP Services* (Solaris 9) for information about test addresses for IP Network Multipathing.
- Rename the `ntp.conf` file to `ntp.conf.cluster`, if `ntp.conf.cluster` does not already exist on the node.
- Set the `local-mac-address?` variable to `true`, if the variable is not already set to that value.

Upgrade processing is finished when the system displays the message `Completed Sun Cluster framework upgrade` and the path to the upgrade log.

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

```
# eject cdrom
```

9. (Optional) Upgrade Sun Cluster data services.

Note – If you are using the Sun Cluster HA for Oracle 3.0 64-bit for Solaris 9 data service, you must upgrade to the Sun Cluster 3.1 9/04 version.

You can continue to use any other Sun Cluster 3.0 data services after you upgrade to Sun Cluster 3.1 9/04 software.

a. Insert the Sun Cluster 3.1 9/04 Agents CD-ROM into the CD-ROM drive on the node.

b. Upgrade the data-service software.

Use one of the following methods:

- To upgrade one or more specified data services, type the following command.

```
# scinstall -u update -s svc [,svc, ...] -d /cdrom/cdrom0
```

-u update	Upgrades a cluster node to a later Sun Cluster software release
-s <i>svc</i>	Upgrades the specified data service
-d	Specifies an alternate directory location for the CD-ROM image

- To upgrade all data services present on the node, type the following command.

```
# scinstall -u update -s all -d /cdrom/cdrom0
```

-s all	Upgrades all data services
--------	----------------------------

The `scinstall` command assumes that updates for all installed data services exist on the update release. If an update for a particular data service does not exist in the update release, that data service is not upgraded.

Upgrade processing is finished when the system displays the message `Completed upgrade of Sun Cluster data services agents` and displays the path to the upgrade log.

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

```
# eject cdrom
```

10. As needed, manually upgrade any custom data services that are not supplied on the Sun Cluster 3.1 9/04 Agents CD-ROM.

11. Verify that each data-service update is installed successfully.

View the upgrade log file that is referenced at the end of the upgrade output messages.

12. Install any Sun Cluster 3.1 9/04 software patches, if you did not already install them by using the `scinstall` command.

13. Install any Sun Cluster 3.1 9/04 data-service software patches.

See “Patches and Required Firmware Levels” in *Sun Cluster 3.1 Release Notes* for the location of patches and installation instructions.

14. Upgrade software applications that are installed on the cluster.

Ensure that application levels are compatible with the current versions of Sun Cluster and Solaris software. See your application documentation for installation instructions.

15. After all nodes are upgraded, reboot each node into the cluster.

```
# reboot
```

16. Verify that all upgraded software is at the same version on all upgraded nodes.

a. On each upgraded node, view the installed levels of Sun Cluster software.

```
# scinstall -pv
```

The first line of output states which version of Sun Cluster software the node is running. This version should match the version that you just upgraded to.

b. From any node, verify that all upgraded cluster nodes are running in cluster mode (Online).

```
# scstat -n
```

See the `scstat(1M)` man page for more information about displaying cluster status.

17. If you upgraded from Solaris 8 to Solaris 9 software, verify the consistency of the storage configuration.

a. On each node, run the following command to verify the consistency of the storage configuration.

```
# scdidadm -c
```

```
-c          Perform a consistency check
```



Caution – Do not proceed to [Step b](#) until your configuration passes this consistency check. Failure to pass this check might result in errors in device identification and cause data corruption.

The following table lists the possible output from the `scdidadm -c` command and the action you must take, if any.

Example Message	Action
device id for 'phys-schost-1:/dev/rdisk/c1t3d0' does not match physical device's id, device may have been replaced	Go to “Recovering From Storage Configuration Changes During Upgrade” on page 214 and perform the appropriate repair procedure.
device id for 'phys-schost-1:/dev/rdisk/c0t0d0' needs to be updated, run <code>scdidadm -R</code> to update	None. You update this device ID in Step b .
No output message	None.

See the `scdidadm(1M)` man page for more information.

b. On each node, migrate the Sun Cluster storage database to Solaris 9 device IDs.

```
# scdidadm -R all
-R      Perform repair procedures
all     Specify all devices
```

c. On each node, run the following command to verify that storage database migration to Solaris 9 device IDs is successful.

```
# scdidadm -c
```

- If the `scdidadm` command displays a message, return to [Step a](#) to make further corrections to the storage configuration or the storage database.
- If the `scdidadm` command displays no messages, the device-ID migration is successful. When device-ID migration is verified on all cluster nodes, proceed to [Step 4](#).

18. Go to [“How to Finish a Nonrolling Upgrade to Sun Cluster 3.1 9/04 Software”](#) on page 198.

Example—Nonrolling Upgrade From Sun Cluster 3.0 to Sun Cluster 3.1 9/04 Software

The following example shows the process of a nonrolling upgrade of a two-node cluster from Sun Cluster 3.0 to Sun Cluster 3.1 9/04 software on the Solaris 8 OS. The example includes the installation of Sun Web Console software and the upgrade of all installed data services that have new versions on the Sun Cluster 3.1 9/04 Agents CD-ROM. The cluster node names are `phys-schost-1` and `phys-schost-2`.

(On the first node, install Sun Web Console software from the Sun Cluster 3.1 9/04 CD-ROM)

```
phys-schost-1# cd /cdrom/cdrom0/Solaris_sparc/Product/sun_cluster/ \
Solaris_8/Misc
phys-schost-1# ./setup
```

(On the first node, upgrade framework software from the Sun Cluster 3.1 9/04 CD-ROM)

```
phys-schost-1# cd /cdrom/cdrom0/Solaris_sparc/Product/sun_cluster/Solaris_8/Tools
phys-schost-1# ./scinstall -u update -S interact
```

(On the first node, upgrade data services from the Sun Cluster 3.1 9/04 Agents CD-ROM)

```
phys-schost-1# scinstall -u update -s all -d /cdrom/cdrom0
```

(On the second node, install Sun Web Console software from the Sun Cluster 3.1 9/04 CD-ROM)

```
phys-schost-2# cd /cdrom/cdrom0/Solaris_sparc/Product/sun_cluster/ \
Solaris_8/Misc
phys-schost-2# ./setup
```

(On the second node, upgrade framework software from the Sun Cluster 3.1 9/04 CD-ROM)

```
phys-schost-2# cd /cdrom/cdrom0/Solaris_sparc/Product/sun_cluster/Solaris_8/Tools
phys-schost-2# ./scinstall -u update -S interact
```

(On the second node, upgrade data services from the Sun Cluster 3.1 9/04 Agents CD-ROM)

```
phys-schost-2# scinstall -u update -s all -d /cdrom/cdrom0
```

(Reboot each node into the cluster)

```
phys-schost-1# reboot
phys-schost-2# reboot
```

(Verify that software versions are the same on all nodes)

```
# scinstall -pv
```

(Verify cluster membership)

```
# scstat -n
-- Cluster Nodes --
      Node name      Status
-----
Cluster node:      phys-schost-1 Online
Cluster node:      phys-schost-2 Online
```

▼ How to Finish a Nonrolling Upgrade to Sun Cluster 3.1 9/04 Software

Perform this procedure to finish Sun Cluster upgrade. First, reregister all resource types that received a new version from the upgrade. Second, modify eligible resources to use the new version of the resource type that the resource uses. Third, re-enable resources. Finally, bring resource groups back online.

Note – To upgrade future versions of resource types, see “Upgrading a Resource Type” in *Sun Cluster Data Service Planning and Administration Guide for Solaris OS*.

1. Ensure that all steps in “How to Perform a Nonrolling Upgrade of Sun Cluster 3.1 9/04 Software” on page 192 are completed.
2. If you upgraded any data services that are not supplied on the Sun Cluster 3.1 9/04 Agents CD-ROM, register the new resource types for those data services. Follow the documentation that accompanies the data services.
3. If you upgraded Sun Cluster HA for SAP liveCache from the version for Sun Cluster 3.0 to the version for Sun Cluster 3.1, modify the `/opt/SUNWsc1c/livecache/bin/lccluster` configuration file.
In the `lccluster` file, specify the value of `put-Confdir_list-here` in the `CONFDIR_LIST="put-Confdir_list-here"` entry. This entry did not exist in the Sun Cluster 3.0 version of the `lccluster` file. Follow instructions in “Registering and Configuring the Sun Cluster HA for SAP liveCache” in *Sun Cluster Data Service for SAP liveCache Guide for Solaris OS*.
4. If your configuration uses dual-string mediators for Solstice DiskSuite or Solaris Volume Manager software, restore the mediator configurations.
 - a. Determine which node has ownership of a disk set to which you will add the mediator hosts.

```
# metaset -s setname
```

`-s setname` Specifies the disk set name
 - b. If no node has ownership, take ownership of the disk set.

```
# metaset -s setname -t
```

`-t` Takes ownership of the disk set
 - c. Re-create the mediators.

```
# metaset -s setname -a -m mediator-host-list
```

`-a` Adds to the disk set
`-m mediator-host-list` Specifies the names of the nodes to add as mediator hosts for the disk set
 - d. Repeat [Step a](#) through [Step c](#) for each disk set in the cluster that uses mediators.
5. SPARC: If you upgraded VxVM, upgrade all disk groups.

To upgrade a disk group to the highest version supported by the VxVM release you installed, run the following command from the primary node of the disk group:

```
# vxvg upgrade dgname
```

See your VxVM administration documentation for more information about upgrading disk groups.

6. From any node, start the `scsetup(1M)` utility.

```
# scsetup
```

7. Re-enable all disabled resources.

- a. From the Resource Group Menu, choose Enable/Disable a resource.
- b. Choose a resource to enable and follow the prompts.
- c. Repeat [Step b](#) for each disabled resource.
- d. When all resources are re-enabled, type `q` to return to the Resource Group Menu.

8. Bring each resource group back online.

- a. From the Resource Group Menu, choose Online/Offline or Switchover a resource group.
- b. Follow the prompts to put each resource group into the managed state and then bring the resource group online.

9. When all resource groups are back online, exit the `scsetup` utility.

Type `q` to back out of each submenu, or press Ctrl-C.

10. (Optional) Migrate resources to new resource type versions.

See “Upgrading a Resource Type” in *Sun Cluster Data Services Planning and Administration Guide for Solaris OS*, which contains procedures which use the command line. Alternatively, you can perform the same tasks by using the Resource Group menu of the `scsetup` utility. The process involves performing the following tasks:

- Register the new resource type.
- Migrate the eligible resource to the new version of its resource type.
- Modify the extension properties of the resource type as specified in the manual for the related data service.

11. If you have a SPARC based system and use Sun Management Center to monitor the cluster, go to [“SPARC: How to Upgrade Sun Cluster-Module Software for Sun Management Center”](#) on page 217.

The cluster upgrade is complete.

Upgrading to Sun Cluster 3.1 9/04 Software (Rolling)

This section provides procedures to perform a rolling upgrade from Sun Cluster 3.1 software to Sun Cluster 3.1 9/04 software. In a rolling upgrade, you upgrade one cluster node at a time, while the other cluster nodes remain in production. After all nodes are upgraded and have rejoined the cluster, you must commit the cluster to the new software version before you can use any new features.

To upgrade from Sun Cluster 3.0 software, follow the procedures in [“Upgrading to Sun Cluster 3.1 9/04 Software \(Nonrolling\)”](#) on page 183.

Note – Sun Cluster 3.1 9/04 software does not support rolling upgrade from Solaris 8 software to Solaris 9 software. You can upgrade Solaris software to an update release during Sun Cluster rolling upgrade. To upgrade a Sun Cluster configuration from Solaris 8 software to Solaris 9 software, perform the procedures in [“Upgrading to Sun Cluster 3.1 9/04 Software \(Nonrolling\)”](#) on page 183.

Task Map: Upgrading to Sun Cluster 3.1 9/04 Software (Rolling)

To perform a rolling upgrade, follow the tasks that are listed in [Table 5–2](#).

TABLE 5–2 Task Map: Upgrading to Sun Cluster 3.1 9/04 Software

Task	Instructions
1. Read the upgrade requirements and restrictions.	“Upgrade Requirements and Support Guidelines” on page 182
2. On one node of the cluster, move resource groups and device groups to another cluster node, and ensure that shared data and system disks are backed up. If the cluster uses dual-string mediators for Solstice DiskSuite or Solaris Volume Manager software, unconfigure the mediators. Then reboot the node into noncluster mode.	“How to Prepare a Cluster Node for a Rolling Upgrade” on page 202

TABLE 5-2 Task Map: Upgrading to Sun Cluster 3.1 9/04 Software (Continued)

Task	Instructions
3. Upgrade the Solaris OS on the cluster node, if necessary, to a supported Solaris update release. SPARC: Optionally, upgrade VERITAS File System (VxFS) and VERITAS Volume Manager (VxVM).	“How to Perform a Rolling Upgrade of a Solaris Maintenance Update” on page 206
4. Upgrade the cluster node to Sun Cluster 3.1 9/04 framework and data-service software. If necessary, upgrade applications. SPARC: If you upgraded VxVM, upgrade disk groups. Then reboot the node back into the cluster.	“How to Perform a Rolling Upgrade of Sun Cluster 3.1 9/04 Software” on page 207
5. Repeat Tasks 2 through 4 on each remaining node to upgrade.	
6. Use the <code>scversions</code> command to commit the cluster to the upgrade. If the cluster uses dual-string mediators, reconfigure the mediators. Optionally, migrate existing resources to new resource types.	“How to Finish a Rolling Upgrade to Sun Cluster 3.1 9/04 Software” on page 212
7. (Optional) SPARC: Upgrade the Sun Cluster module to Sun Management Center.	“SPARC: How to Upgrade Sun Cluster-Module Software for Sun Management Center” on page 217

▼ How to Prepare a Cluster Node for a Rolling Upgrade

Perform this procedure on one node at a time. You will take the upgraded node out of the cluster while the remaining nodes continue to function as active cluster members.

Note – Observe the following guidelines when you perform a rolling upgrade:

- Limit the amount of time that you take to complete a rolling upgrade of all cluster nodes. After a node is upgraded, begin the upgrade of the next cluster node as soon as possible. You can experience performance penalties and other penalties when you run a mixed-version cluster for an extended period of time.
 - Avoid installing new data services or issuing any administrative configuration commands during the upgrade.
 - Until all nodes of the cluster are successfully upgraded and the upgrade is committed, new features that are introduced by the new release might not be available.
-

1. Ensure that the configuration meets requirements for upgrade.

See [“Upgrade Requirements and Support Guidelines” on page 182](#).

2. **Have available the CD-ROMs, documentation, and patches for all the software products you are upgrading before you begin to upgrade the cluster.**
 - Solaris 8 or Solaris 9 OS
 - Sun Cluster 3.1 9/04 framework
 - Sun Cluster 3.1 9/04 data services (agents)
 - Applications that are managed by Sun Cluster 3.1 9/04 data-service agents

See “Patches and Required Firmware Levels” in *Sun Cluster 3.1 Release Notes* for the location of patches and installation instructions.
3. **(Optional) Install Sun Cluster 3.1 9/04 documentation.**

Install the documentation packages on your preferred location, such as an administrative console or a documentation server. See the `index.html` file at the top level of the Sun Cluster 3.1 9/04 CD-ROM to access installation instructions.
4. **Become superuser on one node of the cluster to upgrade.**
5. **If not already installed, install Sun Web Console packages.**

These packages are required by Sun Cluster software, even if you do not use Sun Web Console.

 - a. **Insert the Sun Cluster 3.1 9/04 CD-ROM in the CD-ROM drive.**
 - b. **Change to the**
`/cdrom/cdrom0/Solaris_arch/Product/sun_web_console/2.1/`
directory, where *arch* is `sparc` or `x86`.
 - c. **Run the `setup` command.**

```
# ./setup
```

The `setup` command installs all packages to support Sun Web Console.
6. **For a two-node cluster, if the cluster uses Sun StorEdge Availability Suite software, ensure that the configuration data for availability services resides on the quorum disk.**

The configuration data must reside on a quorum disk to ensure the proper functioning of Sun StorEdge Availability Suite after you upgrade the cluster software.

 - a. **Become superuser on a node of the cluster that runs Sun StorEdge Availability Suite software.**
 - b. **Identify the device ID and the slice that is used by the Sun StorEdge Availability Suite configuration file.**

```
# /usr/opt/SUNWscm/sbin/dscfg  
/dev/did/rdisk/dNsS
```

In this example output, *N* is the device ID and *S* the slice of device *N*.
 - c. **Identify the existing quorum device.**

```
# scstat -g
-- Quorum Votes by Device --
      Device Name          Present Possible Status
-----
Device votes:  /dev/did/rdisk/dQsS  1      1      Online
```

In this example output, dQsS is the existing quorum device.

- d. **If the quorum device is not the same as the Sun StorEdge Availability Suite configuration-data device, move the configuration data to an available slice on the quorum device.**

```
# dd if='/usr/opt/SUNWesm/sbin/dscfg' of=/dev/did/rdisk/dQsS
```

Note – You must use the name of the raw DID device, /dev/did/rdisk/, not the block DID device, /dev/did/dsk/.

- e. **If you moved the configuration data, configure Sun StorEdge Availability Suite software to use the new location.**

As superuser, issue the following command on each node that runs Sun StorEdge Availability Suite software.

```
# /usr/opt/SUNWesm/sbin/dscfg -s /dev/did/rdisk/dQsS
```

7. **From any node, view the current status of the cluster.**

Save the output as a baseline for later comparison.

```
% scstat
% scrgadm -pv[v]
```

See the scstat(1M) and scrgadm(1M) man pages for more information.

8. **Move all resource groups and device groups that are running on the node to upgrade.**

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

-S Moves all resource groups and device groups

-h *from-node* Specifies the name of the node from which to move resource groups and device groups

See the scswitch(1M) man page for more information.

9. **Verify that the move was completed successfully.**

```
# scstat -g -D
```

-g Shows status for all resource groups

-D Shows status for all disk device groups

10. **Ensure that the system disk, applications, and all data are backed up.**

11. If your cluster uses dual-string mediators for Solstice DiskSuite or Solaris Volume Manager software, unconfigure your mediators.

See “Configuring Dual-String Mediators” on page 159 for more information.

a. Run the following command to verify that no mediator data problems exist.

```
# medstat -s setname
```

-s *setname* Specifies the disk set name

If the value in the Status field is Bad, repair the affected mediator host. Follow the procedure “How to Fix Bad Mediator Data” on page 161.

b. List all mediators.

Save this information for when you restore the mediators during the procedure “How to Finish a Rolling Upgrade to Sun Cluster 3.1 9/04 Software” on page 212.

c. For a disk set that uses mediators, take ownership of the disk set if no node already has ownership.

```
# metaset -s setname -t
```

-t Takes ownership of the disk set

d. Unconfigure all mediators for the disk set.

```
# metaset -s setname -d -m mediator-host-list
```

-s *setname* Specifies the disk-set name

-d Deletes from the disk set

-m *mediator-host-list* Specifies the name of the node to remove as a mediator host for the disk set

See the mediator(7D) man page for further information about mediator-specific options to the metaset command.

e. Repeat Step c through Step d for each remaining disk set that uses mediators.

12. Shut down the node that you want to upgrade and boot it into noncluster mode.

On SPARC based systems, perform the following commands:

```
# shutdown -y -g0
```

```
ok boot -x
```

On x86 based systems, perform the following commands:

```
# shutdown -y -g0
```

```
...
```

```
<<< Current Boot Parameters >>>
```

```
Boot path: /pci@0,0/pci-ide@7,1/ata@1/cmdk@0,0:b
```

```
Boot args:
```

```
Type b [file-name] [boot-flags] <ENTER> to boot with options
```

```
or i <ENTER> to enter boot interpreter
```

or <ENTER> to boot with defaults

<<< timeout in 5 seconds >>>

Select (b)oot or (i)nterpreter: **b -x**

The other nodes of the cluster continue to function as active cluster members.

13. To upgrade the Solaris software to a Maintenance Update release, go to [“How to Perform a Rolling Upgrade of a Solaris Maintenance Update”](#) on page 206.

Note – The cluster must already run on, or be upgraded to, at least the minimum required level of the Solaris OS to support Sun Cluster 3.1 9/04 software. See the *Sun Cluster Release Notes for Solaris OS* for information about supported releases of the Solaris OS.

14. Go to [“How to Perform a Rolling Upgrade of Sun Cluster 3.1 9/04 Software”](#) on page 207.

▼ How to Perform a Rolling Upgrade of a Solaris Maintenance Update

Perform this procedure to upgrade the Solaris 8 or Solaris 9 OS to a supported Maintenance Update release.

Note – To upgrade a cluster from Solaris 8 to Solaris 9 software, with or without upgrading Sun Cluster software as well, you must perform a nonrolling upgrade. Go to [“Upgrading to Sun Cluster 3.1 9/04 Software \(Nonrolling\)”](#) on page 183.

1. Ensure that all steps in [“How to Prepare a Cluster Node for a Rolling Upgrade”](#) on page 202 are completed.
2. Temporarily comment out all entries for globally mounted file systems in the node's `/etc/vfstab` file.
Perform this step to prevent the Solaris upgrade from attempting to mount the global devices.
3. Follow the instructions in the Solaris maintenance update installation guide to install the Maintenance Update release.

Note – Do not reboot the node when prompted to reboot at the end of installation processing.

4. Uncomment all entries in the `/a/etc/vfstab` file for globally mounted file systems that you commented out in [Step 2](#).
5. Install any required Solaris software patches and hardware-related patches, and download any needed firmware that is contained in the hardware patches.

Note – Do not reboot the node until [Step 6](#).

6. Reboot the node into noncluster mode.
Include the double dashes (--) in the following command:

```
# reboot -- -x
```

7. Upgrade the Sun Cluster software.
Go to “[How to Perform a Rolling Upgrade of Sun Cluster 3.1 9/04 Software](#)” on page 207.

▼ How to Perform a Rolling Upgrade of Sun Cluster 3.1 9/04 Software

Perform this procedure to upgrade a node to Sun Cluster 3.1 9/04 software while the remaining cluster nodes are in cluster mode.

Note – Until all nodes of the cluster are upgraded and the upgrade is committed, new features that are introduced by the new release might not be available.

1. Ensure that all steps in “[How to Prepare a Cluster Node for a Rolling Upgrade](#)” on page 202 are completed.
2. If you upgraded the Solaris OS to a Maintenance Update release, ensure that all steps in “[How to Perform a Rolling Upgrade of a Solaris Maintenance Update](#)” on page 206 are completed.
3. Ensure that you have installed all required Solaris software patches and hardware-related patches.

For Solstice DiskSuite software (Solaris 8), also ensure that you have installed all required Solstice DiskSuite software patches.

4. Become superuser on the node of the cluster.

5. Install Sun Web Console packages.

Perform this step on each node of the cluster. These packages are required by Sun Cluster software, even if you do not use Sun Web Console.

a. Insert the Sun Cluster 3.1 9/04 CD-ROM in the CD-ROM drive.

b. Change to the

`/cdrom/cdrom0/Solaris_arch/Product/sun_web_console/2.1/`
directory, where *arch* is *sparc* or *x86*.

c. Run the `setup` command.

```
# ./setup
```

The `setup` command installs all packages to support Sun Web Console.

6. On the Sun Cluster 3.1 9/04 CD-ROM, change to the

`Solaris_arch/Product/sun_cluster/Solaris_ver/Tools/` directory, where *arch* is *sparc* or *x86* and where *ver* is 8 (for Solaris 8) or 9 (for Solaris 9) .

```
# cd /cdrom/cdrom0/Solaris_arch/Product/sun_cluster/Solaris_ver/Tools
```

7. Upgrade the cluster framework software.

Note – Do not use the `/usr/cluster/bin/scinstall` command that is already installed on the node. You must use the `scinstall` command that is on the Sun Cluster 3.1 9/04 CD-ROM.

```
# ./scinstall -u update [-M patchdir=dirname]
```

```
-M patchdir=dirname [, patchlistfile=filename]
```

Specifies the path to patch information so that the specified patches can be installed by the `scinstall` command. If you do not specify a patch-list file, the `scinstall` command installs all the patches in the directory *dirname*, including tarred, jarred, and zipped patches.

The `-M` option is not required. You can use any method you prefer for installing patches.

See “Patches and Required Firmware Levels” in *Sun Cluster 3.1 Release Notes* for the location of patches and installation instructions.

Note – Sun Cluster 3.1 9/04 software requires at least version 3.5.1 of Sun Explorer software. Upgrading to Sun Cluster software includes installing Sun Explorer data collector software, to be used in conjunction with the `sccheck` utility. If another version of Sun Explorer software was already installed before the Sun Cluster upgrade, it is replaced by the version that is provided with Sun Cluster software. Options such as user identity and data delivery are preserved, but `crontab` entries must be manually re-created.

Upgrade processing is finished when the system displays the message `Completed Sun Cluster framework upgrade` and the path to the upgrade log.

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

```
# eject cdrom
```

9. (Optional) Upgrade Sun Cluster data services.

Note – If you are using the Sun Cluster HA for Oracle 3.0 64-bit for Solaris 9 data service, you must upgrade to the Sun Cluster 3.1 9/04 version.

You can continue to use any other Sun Cluster 3.0 data services after you upgrade to Sun Cluster 3.1 9/04 software.

a. Insert the Sun Cluster 3.1 9/04 Agents CD-ROM into the CD-ROM drive on the node.

b. Upgrade the data-service software.

Use one of the following methods:

- To upgrade one or more specified data services, type the following command.

```
# scinstall -u update -s svc[,svc,...] -d /cdrom/cdrom0
```

<code>-u update</code>	Upgrades a cluster node to a later Sun Cluster software release
<code>-s <i>svc</i></code>	Upgrades the specified data service
<code>-d</code>	Specifies an alternate directory location for the CD-ROM image

- To upgrade all data services present on the node, type the following command.

```
# scinstall -u update -s all -d /cdrom/cdrom0
```

<code>-s all</code>	Upgrades all data services
---------------------	----------------------------

The `scinstall` command assumes that updates for all installed data services exist on the update release. If an update for a particular data service does not exist in the update release, that data service is not upgraded.

Upgrade processing is finished when the system displays the message `Completed upgrade of Sun Cluster data services agents` and displays the path to the upgrade log.

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

```
# eject cdrom
```

10. **As needed, manually upgrade any custom data services that are not supplied on the Sun Cluster 3.1 9/04 Agents CD-ROM.**

11. **Verify that each data-service update is installed successfully.**

View the upgrade log file that is referenced at the end of the upgrade output messages.

12. **Install any Sun Cluster 3.1 9/04 software patches, if you did not already install them by using the `scinstall` command.**

13. **Install any Sun Cluster 3.1 9/04 data-service software patches.**

See "Patches and Required Firmware Levels" in *Sun Cluster 3.1 Release Notes* for the location of patches and installation instructions.

14. **Upgrade software applications that are installed on the cluster.**

Ensure that application levels are compatible with the current versions of Sun Cluster and Solaris software. See your application documentation for installation instructions. In addition, follow these guidelines to upgrade applications in a Sun Cluster 3.1 9/04 configuration:

- If the applications are stored on shared disks, you must master the relevant disk groups and manually mount the relevant file systems before you upgrade the application.
- If you are instructed to reboot a node during the upgrade process, always add the `-x` option to the command.

The `-x` option ensures that the node reboots into noncluster mode. For example, either of the following two commands boot a node into single-user noncluster mode:

On SPARC based systems, perform the following commands:

```
# reboot -- -xs
ok boot -xs
```

On x86 based systems, perform the following commands:

```

# reboot -- -xs
...
                                <<< Current Boot Parameters >>>
Boot path: /pci@0,0/pci-ide@7,1/ata@1/cmdk@0,0:b
Boot args:

Type   b [file-name] [boot-flags] <ENTER>   to boot with options
or     i <ENTER>                             to enter boot interpreter
or     <ENTER>                               to boot with defaults

                                <<< timeout in 5 seconds >>>
Select (b)oot or (i)nterpreter: b -xs

```

Note – Do not upgrade an application if the newer version of the application cannot coexist in the cluster with the older version of the application.

15. Reboot the node into the cluster.

```
# reboot
```

16. Run the following command on the upgraded node to verify that Sun Cluster 3.1 9/04 software was installed successfully.

```
# scinstall -pv
```

The first line of output states which version of Sun Cluster software the node is running. This version should match the version you just upgraded to.

17. From any node, verify the status of the cluster configuration.

```
% scstat
% scrgadm -pv [v]
```

Output should be the same as for Step 7 in “How to Prepare a Cluster Node for a Rolling Upgrade” on page 202.

18. If you have another node to upgrade, return to “How to Prepare a Cluster Node for a Rolling Upgrade” on page 202 and repeat all upgrade procedures on the next node to upgrade.

19. When all nodes in the cluster are upgraded, go to “How to Finish a Rolling Upgrade to Sun Cluster 3.1 9/04 Software” on page 212.

Example—Rolling Upgrade From Sun Cluster 3.1 to Sun Cluster 3.1 9/04 Software

The following example shows the process of a rolling upgrade of a cluster node from Sun Cluster 3.1 to Sun Cluster 3.1 9/04 software on the Solaris 8 OS. The example includes the installation of Sun Web Console software and the upgrade of all installed data services that have new versions on the Sun Cluster 3.1 9/04 Agents CD-ROM.

The cluster node name is `phys-schost-1`.

(Install Sun Web Console software from the Sun Cluster 3.1 9/04 CD-ROM)

```
phys-schost-1# cd /cdrom/cdrom0/Solaris_sparc/Product/sun_cluster/ \
Solaris_8/Misc
phys-schost-1# ./setup
```

(Upgrade framework software from the Sun Cluster 3.1 9/04 CD-ROM)

```
phys-schost-1# cd /cdrom/cdrom0/Solaris_sparc/Product/sun_cluster/Solaris_8/Tools
phys-schost-1# ./scinstall -u update -S interact
```

(Upgrade data services from the Sun Cluster 3.1 9/04 Agents CD-ROM)

```
phys-schost-1# scinstall -u update -s all -d /cdrom/cdrom0
```

(Reboot the node into the cluster)

```
phys-schost-1# reboot
```

(Verify that software upgrade succeeded)

```
# scinstall -pv
```

(Verify cluster status)

```
# scstat
# scrgadm -pv
```

▼ How to Finish a Rolling Upgrade to Sun Cluster 3.1 9/04 Software

1. Ensure that all upgrade procedures are completed for all cluster nodes that you are upgrading.
2. From one node, check the upgrade status of the cluster.

```
# scversions
```
3. From the following table, perform the action that is listed for the output message from [Step 2](#).

Output Message	Action
Upgrade commit is needed.	Go to Step 4 .

Output Message	Action
Upgrade commit is NOT needed. All versions match.	Skip to Step 6 .
Upgrade commit cannot be performed until all cluster nodes are upgraded. Please run <code>scinstall(1m)</code> on cluster nodes to identify older versions.	Return to “How to Perform a Rolling Upgrade of Sun Cluster 3.1 9/04 Software” on page 207 to upgrade the remaining cluster nodes.
Check upgrade cannot be performed until all cluster nodes are upgraded. Please run <code>scinstall(1m)</code> on cluster nodes to identify older versions.	Return to “How to Perform a Rolling Upgrade of Sun Cluster 3.1 9/04 Software” on page 207 to upgrade the remaining cluster nodes.

4. After all nodes have rejoined the cluster, from one node commit the cluster to the upgrade.

```
# scversions -c
```

Committing the upgrade enables the cluster to utilize all features in the newer software. New features are available only after you perform the upgrade commitment.

5. From one node, verify that the cluster upgrade commitment has succeeded.

```
# scversions
```

Upgrade commit is NOT needed. All versions match.

6. If your configuration uses dual-string mediators for Solstice DiskSuite or Solaris Volume Manager software, restore the mediator configurations.

a. Determine which node has ownership of a disk set to which you are adding the mediator hosts.

```
# metaset -s setname
```

`-s setname` Specifies the disk-set name

b. If no node has ownership, take ownership of the disk set.

```
# metaset -s setname -t
```

`-t` Takes ownership of the disk set

c. Re-create the mediators.

```
# metaset -s setname -a -m mediator-host-list
```

`-a` Adds to the disk set

`-m mediator-host-list` Specifies the names of the nodes to add as mediator hosts for the disk set

d. Repeat [Step a](#) through [Step c](#) for each disk set in the cluster that uses mediators.

7. If you upgraded any data services that are not supplied on the Sun Cluster 3.1 9/04 Agents CD-ROM, register the new resource types for those data services.

Follow the documentation that accompanies the data services.

8. (Optional) Switch each resource group and device group back its original node.

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

-z	Performs the switch
-g <i>resource-group</i>	Specifies the resource group to switch
-h <i>node</i>	Specifies the name of the node to switch to
-D <i>disk-device-group</i>	Specifies the device group to switch

9. Restart any applications.

Follow the instructions that are provided in your vendor documentation.

10. (Optional) Migrate resources to new resource type versions.

See "Upgrading a Resource Type" in *Sun Cluster Data Services Planning and Administration Guide for Solaris OS*, which contains procedures which use the command line. Alternatively, you can perform the same tasks by using the Resource Group menu of the `scsetup` utility. The process involves performing the following tasks:

- Register the new resource type.
- Migrate the eligible resource to the new version of its resource type.
- Modify the extension properties of the resource type as specified in the manual for the related data service.

11. If you have a SPARC based system and use Sun Management Center to monitor the cluster, go to ["SPARC: How to Upgrade Sun Cluster-Module Software for Sun Management Center"](#) on page 217.

The cluster upgrade is complete.

Recovering From Storage Configuration Changes During Upgrade

This section provides the following repair procedures to follow if changes were inadvertently made to the storage configuration during upgrade:

- ["How to Handle Storage Reconfiguration During an Upgrade"](#) on page 215

- [“How to Resolve Mistaken Storage Changes During an Upgrade”](#) on page 216

▼ How to Handle Storage Reconfiguration During an Upgrade

Any changes to the storage topology, including running Sun Cluster commands, should be completed before you upgrade the cluster to Solaris 9 software. If, however, changes were made to the storage topology during the upgrade, perform the following procedure. This procedure ensures that the new storage configuration is correct and that existing storage that was not reconfigured is not mistakenly altered.

1. Ensure that the storage topology is correct.

Check whether the devices that were flagged as possibly being replaced map to devices that actually were replaced. If the devices were not replaced, check for and correct possible accidental configuration changes, such as incorrect cabling.

2. Become superuser on a node that is attached to the unverified device.

3. Manually update the unverified device.

```
# sddidadm -R device
-R device          Performs repair procedures on the specified device
See the sddidadm(1M) man page for more information.
```

4. Update the DID driver.

```
# sddidadm -ui
# sddidadm -r
-u          Loads the device ID configuration table into the kernel
-i          Initializes the DID driver
-r          Reconfigures the database
```

5. Repeat [Step 2](#) through [Step 4](#) on all other nodes that are attached to the unverified device.

6. Return to the remaining upgrade tasks.

- For a nonrolling upgrade, go to [Step a](#) in [“How to Perform a Nonrolling Upgrade of Sun Cluster 3.1 9/04 Software”](#) on page 192.
- For a rolling upgrade, go to [Step 4](#) in [“How to Perform a Rolling Upgrade of Sun Cluster 3.1 9/04 Software”](#) on page 207.

▼ How to Resolve Mistaken Storage Changes During an Upgrade

If accidental changes are made to the storage cabling during the upgrade, perform the following procedure to return the storage configuration to the correct state.

Note – This procedure assumes that no physical storage was actually changed. If physical or logical storage devices were changed or replaced, instead follow the procedures in “How to Handle Storage Reconfiguration During an Upgrade” on page 215.

1. Return the storage topology to its original configuration.

Check the configuration of the devices that were flagged as possibly being replaced, including the cabling.

2. As superuser, update the DID driver on each node of the cluster.

```
# sddidadm -ui# sddidadm -r
-u          Loads the device-ID configuration table into the kernel
-i          Initializes the DID driver
-r          Reconfigures the database
```

See the `sddidadm(1M)` man page for more information.

3. If the `sddidadm` command returned any error messages in [Step 2](#), return to [Step 1](#) to make further modifications to correct the storage configuration, then repeat [Step 2](#).

4. Return to the remaining upgrade tasks.

- For a nonrolling upgrade, go to [Step a](#) in “How to Perform a Nonrolling Upgrade of Sun Cluster 3.1 9/04 Software” on page 192.
- For a rolling upgrade, go to [Step 4](#) in “How to Perform a Rolling Upgrade of Sun Cluster 3.1 9/04 Software” on page 207.

SPARC: Sun Management Center Software Upgrade

This section provides procedures to upgrade the Sun Cluster-module for Sun Management Center software and procedures to upgrade both Sun Management Center software and Sun Cluster-module software.

▼ SPARC: How to Upgrade Sun Cluster-Module Software for Sun Management Center

Perform the following steps to upgrade Sun Cluster-module software on the Sun Management Center server machine, help-server machine, and console machine.

Note – If you intend to upgrade the Sun Management Center software itself, do not perform this procedure. Instead, go to [“SPARC: How to Upgrade Sun Management Center Software” on page 218](#) to upgrade the Sun Management Center software and the Sun Cluster module.

1. As superuser, remove the existing Sun Cluster-module packages.

Use the `pkgrm(1M)` command to remove all Sun Cluster-module packages from all locations that are listed in the following table.

```
# pkgrm module-package
```

Location	Module Package to Remove
Sun Management Center console machine	SUNWscscn
Sun Management Center server machine	SUNWscssv
Sun Management Center help-server machine	SUNWscsh1

Note – Sun Cluster-module software on the cluster nodes was already upgraded during the cluster-framework upgrade.

2. As superuser, reinstall Sun Cluster–module packages from the Sun Cluster 3.1 9/04 CD-ROM to the locations that are listed in the following table.

In the CD-ROM path, the value of *arch* is **sparc** or **x86**, and the value of *ver* is **8** (for Solaris 8) or **9** (for Solaris 9).

```
# cd /cdrom/cdrom0/Solaris_arch/Product/sun_cluster/Solaris_ver/Packages/
# pkgadd -d . module-package
```

Location	Module Package to Install
Sun Management Center console machine	SUNWscsh1
Sun Management Center server machine	SUNWscssv
Sun Management Center help-server machine	SUNWscsh1

Note that you install the help-server package SUNWscsh1 on both the console machine and the help-server machine. Also, you do not upgrade to a new SUNWscscn package on the console machine.

▼ SPARC: How to Upgrade Sun Management Center Software

Perform the following steps to upgrade from Sun Management Center 2.1.1 to either Sun Management Center 3.0 software or Sun Management Center 3.5 software.

1. Have available the following items:

- **Sun Cluster 3.1 9/04 CD-ROM or the path to the CD-ROM image.**
You use the CD-ROM to reinstall the Sun Cluster 3.1 9/04 version of the Sun Cluster–module packages after you upgrade Sun Management Center software.
- **Sun Management Center documentation.**
- **Sun Management Center patches and Sun Cluster–module patches, if any.**
See “Patches and Required Firmware Levels” in *Sun Cluster 3.1 Release Notes* for the location of patches and installation instructions.

2. Stop any Sun Management Center processes.

- a. **If the Sun Management Center console is running, exit the console.**
In the console window, choose Exit from the File menu.
- b. **On each Sun Management Center agent machine (cluster node), stop the Sun Management Center agent process.**

```
# /opt/SUNWsymon/sbin/es-stop -a
```

- c. **On the Sun Management Center server machine, stop the Sun Management Center server process.**

```
# /opt/SUNWsymon/sbin/es-stop -S
```

- 3. **As superuser, remove Sun Cluster–module packages.**

Use the `pkgrm(1M)` command to remove all Sun Cluster–module packages from all locations that are listed in the following table.

```
# pkgrm module-package
```

Location	Module Package to Remove
Each cluster node	SUNWscsam, SUNWscsal
Sun Management Center console machine	SUNWscscn
Sun Management Center server machine	SUNWscssv
Sun Management Center help-server machine	SUNWscsh1

If you do not remove the listed packages, the Sun Management Center software upgrade might fail because of package dependency problems. You reinstall these packages in [Step 5](#), after you upgrade Sun Management Center software.

- 4. **Upgrade the Sun Management Center software.**

Follow the upgrade procedures in your Sun Management Center documentation.

- 5. **As superuser, reinstall Sun Cluster–module packages from the Sun Cluster 3.1 9/04 CD-ROM to the locations that are listed in the following table.**

In the CD-ROM path, the value of *arch* is `sparc` or `x86`, and the value of *ver* is `8` (for Solaris 8) or `9` (for Solaris 9).

```
# cd /cdrom/cdrom0/Solaris_arch/Product/sun_cluster/Solaris_ver/Packages/  
# pkgadd -d . module-package
```

Location	Module Package to Install
Each cluster node	SUNWscsam, SUNWscsal
Sun Management Center server machine	SUNWscssv
Sun Management Center console machine	SUNWscsh1
Sun Management Center help-server machine	SUNWscsh1

You install the help-server package `SUNWscsh1` on both the console machine and the help-server machine.

6. **Apply any Sun Management Center patches and any Sun Cluster-module patches to each node of the cluster.**
7. **Restart Sun Management Center agent, server, and console processes.**
Follow procedures in [“SPARC: How to Start Sun Management Center”](#) on page 120.
8. **Load the Sun Cluster module.**
Follow procedures in [“SPARC: How to Load the Sun Cluster Module”](#) on page 121.
If the Sun Cluster module was previously loaded, unload the module and then reload it to clear all cached alarm definitions on the server. To unload the module, choose Unload Module from the Module menu on the console’s Details window.

Configuring Data Replication With Sun StorEdge Availability Suite 3.1 Software

This chapter provides guidelines for configuring data replication between clusters by using Sun StorEdge Availability Suite 3.1 software.

This chapter also contains an example of how data replication was configured for an NFS application by using Sun StorEdge Availability Suite 3.1 software. This example uses a specific cluster configuration and provides detailed information about how individual tasks can be performed. It does not include all of the steps that are required by other applications or other cluster configurations.

This chapter contains the following sections:

- [“Introduction to Data Replication” on page 222](#)
- [“Guidelines for Configuring Data Replication” on page 225](#)
- [“Example Configuration” on page 230](#)

The following procedures are in this chapter:

- [“How to Configure a Disk Device Group on the Primary Cluster” on page 234](#)
- [“How to Configure a Disk Device Group on the Secondary Cluster” on page 235](#)
- [“How to Configure the File System on the Primary Cluster for the NFS Application” on page 235](#)
- [“How to Configure the File System on the Secondary Cluster for the NFS Application” on page 237](#)
- [“How to Create a Replication Resource Group on the Primary Cluster” on page 237](#)
- [“How to Create a Replication Resource Group on the Secondary Cluster” on page 238](#)
- [“How to Create an Application Resource Group on the Primary Cluster” on page 238](#)
- [“How to Create an Application Resource Group on the Secondary Cluster” on page 240](#)
- [“How to Enable Replication on the Primary Cluster” on page 241](#)
- [“How to Enable Replication on the Secondary Cluster” on page 242](#)
- [“How to Perform a Remote Mirror Replication” on page 243](#)
- [“How to Perform a Point-in-Time Snapshot” on page 244](#)
- [“How to Verify That Replication Is Configured Correctly” on page 245](#)
- [“How to Provoke a Switchover” on page 248](#)

- “How to Update the DNS Entry” on page 248
 - “How to Configure the Application to Read and Write to the Secondary Volume” on page 249
-

Introduction to Data Replication

This section introduces disaster tolerance and describes the data replication methods used by Sun StorEdge Availability Suite software.

What Is Disaster Tolerance?

Disaster tolerance is the ability of a system to restore an application on an alternate cluster when the primary cluster fails. Disaster tolerance is based on *data replication* and *failover*.

Data replication is the copying of data from a primary cluster to a backup or secondary cluster. Through data replication, the secondary cluster has an up-to-date copy of the data on the primary cluster. The secondary cluster can be located far away from the primary cluster.

Failover is the automatic relocation of a resource group or device group from a primary cluster to a secondary cluster. If the primary cluster fails, the application and the data are immediately available on the secondary cluster.

Data Replication Methods Used by Sun StorEdge Availability Suite Software

This section describes the remote mirror replication method and the point-in-time snapshot method used by Sun StorEdge Availability Suite software. This software uses the `sndradm(1RPC)` and `iiadm(1II)` commands to replicate data. For more information about these commands, see the *Sun Cluster 3.0 and Sun StorEdge Software Integration Guide*.

Remote Mirror Replication

Remote mirror replication is illustrated in [Figure 6-1](#). Data from the master volume of the primary disk is replicated to the master volume of the secondary disk through a TCP/IP connection. A remote mirror bitmap tracks differences between the master volume on the primary disk and the master volume on the secondary disk.

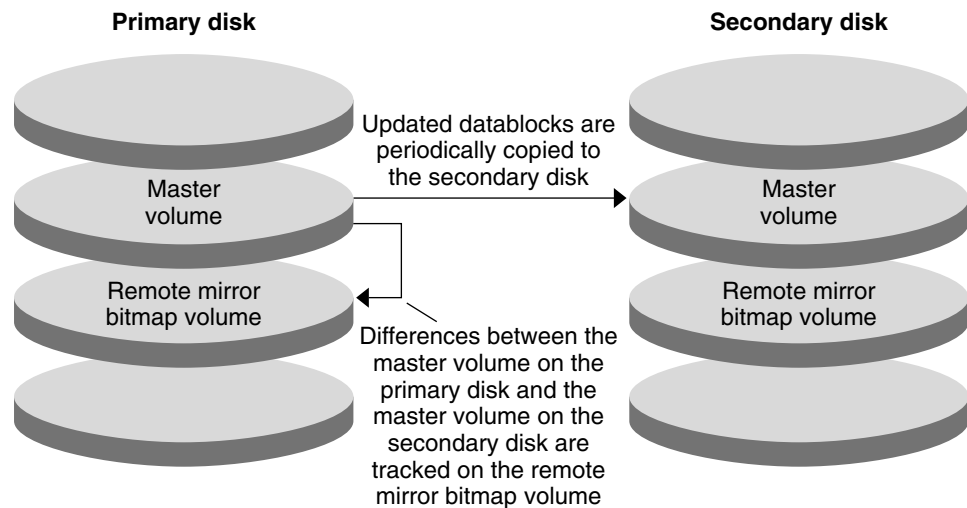


FIGURE 6-1 Remote Mirror Replication

Remote mirror replication can be performed synchronously in real time, or asynchronously. Each volume set in each cluster can be configured individually, for synchronous replication or asynchronous replication.

In synchronous data replication, a write operation is not confirmed as complete until the remote volume has been updated.

In asynchronous data replication, a write operation is confirmed as complete before the remote volume is updated. Asynchronous data replication provides greater flexibility over long distances and low bandwidth.

Point-in-Time Snapshot

Point-in-time snapshot is illustrated in [Figure 6-2](#). Data from the master volume of each disk is copied to the shadow volume on the same disk. The point-in-time bitmap tracks differences between the master volume and the shadow volume. When data is copied to the shadow volume, the point-in-time bitmap is reset.

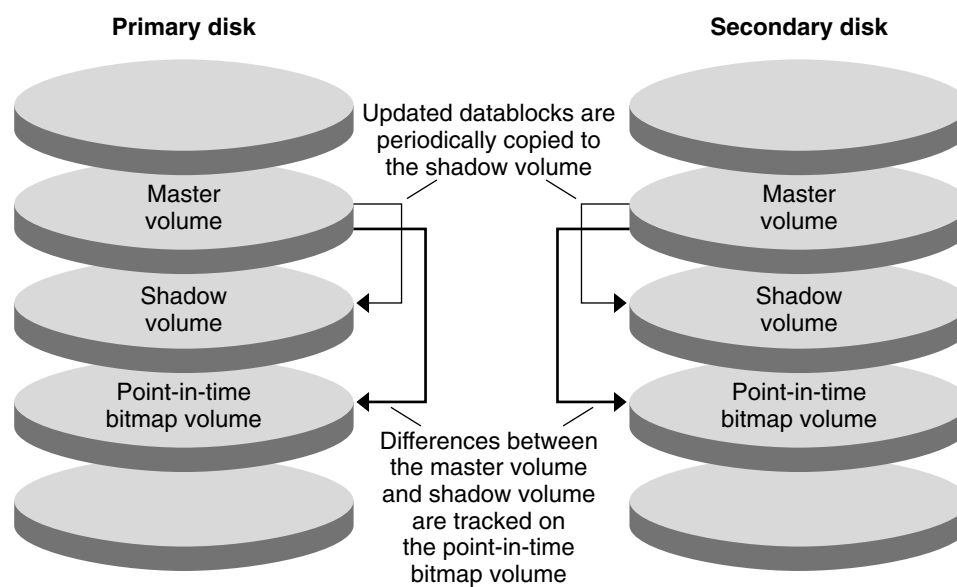


FIGURE 6-2 Point-in-Time Snapshot

Replication in the Example Configuration

The following figure illustrates how remote mirror replication and point-in-time snapshot are used in [“Example Configuration”](#) on page 230.




FIGURE 6-3 Replication in the Example Configuration

Guidelines for Configuring Data Replication

This section provides guidelines for configuring data replication between clusters. This section also contains tips for configuring replication resource groups and application resource groups. Use these guidelines when you are configuring data replication for your cluster.

This section discusses the following topics:

- [“Configuring Replication Resource Groups” on page 226](#)
- [“Configuring Application Resource Groups” on page 226](#)
- [“Configuring Resource Groups for a Failover Application” on page 227](#)
- [“Configuring Resource Groups for a Scalable Application” on page 228](#)
- [“Guidelines for Managing a Failover or Switchover” on page 229](#)

Configuring Replication Resource Groups

Replication resource groups colocate the device group under Sun StorEdge Availability Suite software control with the logical hostname resource. A replication resource group must have the following characteristics:

- Be a failover resource group

A failover resource can run on only one node at a time. When a failover occurs, failover resources take part in the failover.
- Have a logical hostname resource

The logical hostname must be hosted by the primary cluster. After a failover or switchover, the logical hostname must be hosted by the secondary cluster. The Domain Name System (DNS) is used to associate the logical hostname with a cluster.
- Have an HAStoragePlus resource

The HAStoragePlus resource enforces the switchover of the device group when the replication resource group is switched over or failed over. Sun Cluster software also enforces the switchover of the replication resource group when the device group is switched over. In this way, the replication resource group and the device group are always colocated, or mastered by the same node.

The following extension properties must be defined in the HAStoragePlus resource:

 - *GlobalDevicePaths*. This extension property defines the device group to which a volume belongs.
 - *AffinityOn property = True*. This extension property causes the device group to switch over or fail over when the replication resource group switches over or fails over. This feature is called an *affinity switchover*.

For more information about HAStoragePlus, see the `SUNW.HAStoragePlus(5)` man page.
- Be named after the device group with which it is colocated, followed by `-stor-rg`

For example, `devicegroup-stor-rg`.
- Be online on both the primary cluster and the secondary cluster

Configuring Application Resource Groups

To be highly available, an application must be managed as a resource in an application resource group. An application resource group can be configured for a failover application or a scalable application.

Application resources and application resource groups configured on the primary cluster must also be configured on the secondary cluster. Also, the data accessed by the application resource must be replicated to the secondary cluster.

This section provides guidelines for configuring the following application resource groups:

- “Configuring Resource Groups for a Failover Application” on page 227
- “Configuring Resource Groups for a Scalable Application” on page 228

Configuring Resource Groups for a Failover Application

In a failover application, an application runs on one node at a time. If that node fails, the application fails over to another node in the same cluster. A resource group for a failover application must have the following characteristics:

- Have an `HASStoragePlus` resource to enforce the switchover of the device group when the application resource group is switched over or failed over

The device group is colocated with the replication resource group and the application resource group. Therefore, the switchover of the application resource group enforces the switchover of the device group and replication resource group. The application resource group, the replication resource group, and the device group are mastered by the same node.

Note, however, that a switchover or failover of the device group or the replication resource group does not cause a switchover or failover of the application resource group.

- If the application data is globally mounted, the presence of an `HASStoragePlus` resource in the application resource group is not compulsory but is advised.
- If the application data is mounted locally, the presence of an `HASStoragePlus` resource in the application resource group is compulsory.

Without an `HASStoragePlus` resource, the switchover or failover of the application resource group would not trigger the switchover or failover of the replication resource group and device group. After a switchover or failover, the application resource group, replication resource group, and device group would not be mastered by the same node.

For more information about `HASStoragePlus`, see the `SUNW.HASStoragePlus(5)` man page.

- Must be online on the primary cluster and offline on the secondary cluster

The application resource group must be brought online on the secondary cluster when the secondary cluster takes over as the primary cluster.

The following figure illustrates the configuration of an application resource group and a replication resource group in a failover application.

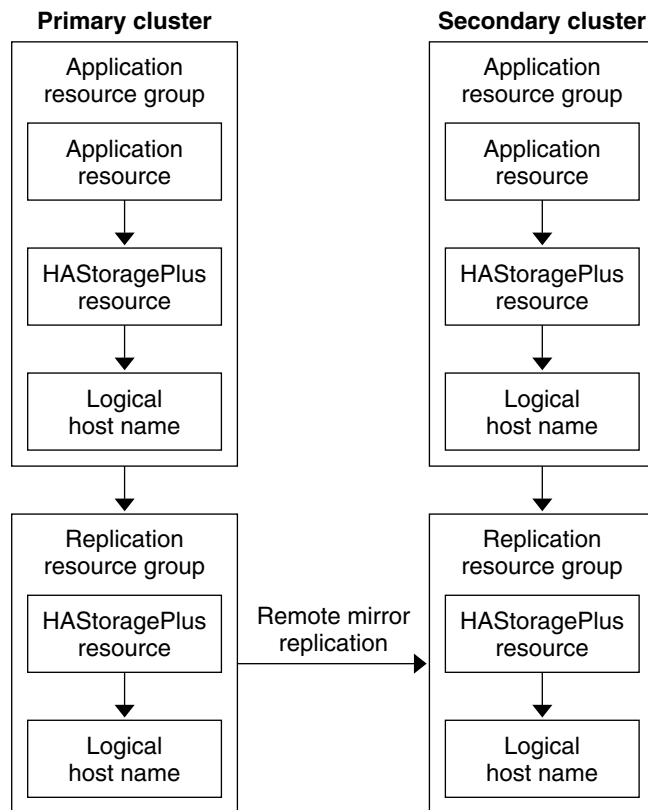


FIGURE 6-4 Configuration of Resource Groups in a Failover Application

Configuring Resource Groups for a Scalable Application

In a scalable application, an application runs on several nodes to create a single, logical service. If a node that is running a scalable application fails, failover does not occur. The application continues to run on the other nodes.

When a scalable application is managed as a resource in an application resource group, it is not necessary to colocate the application resource group with the device group. Therefore, it is not necessary to create an HAStoragePlus resource for the application resource group.

A resource group for a scalable application must have the following characteristics:

- Have a dependency on the shared address resource group
 - The shared address is used by the nodes that are running the scalable application, to distribute incoming data.
- Be online on the primary cluster and offline on the secondary cluster

The following figure illustrates the configuration of resource groups in a scalable application.

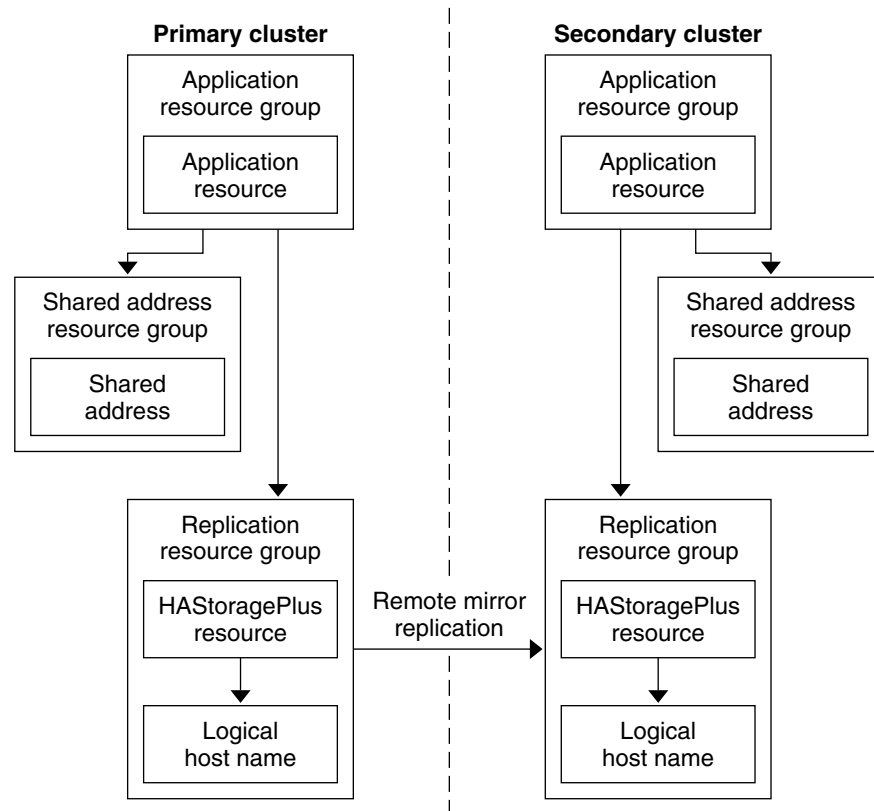


FIGURE 6-5 Configuration of Resource Groups in a Scalable Application

Guidelines for Managing a Failover or Switchover

If the primary cluster fails, the application must be switched over to the secondary cluster as soon as possible. To enable the secondary cluster to take over, the DNS must be updated. In addition, the secondary volume must be mounted on the mount point directory for the application file system.

The DNS associates a client with the logical hostname of an application. After a failover or switchover, the DNS mapping to the primary cluster must be removed, and a DNS mapping to the secondary cluster must be created. The following figure shows how the DNS maps a client to a cluster.




FIGURE 6-6 DNS Mapping of a Client to a Cluster

To update the DNS, use the `nsupdate` command. For information, see the `nsupdate(1M)` man page. For an example of how to cope with a failover or switchover, see [“Example of How to Cope With a Failover or Switchover”](#) on page 247.

After repair, the primary cluster can be brought back online. To switch back to the original primary cluster, perform the following steps:

1. Synchronize the primary cluster with the secondary cluster to ensure that the primary volume is up-to-date.
2. Update the DNS so that clients can access the application on the primary cluster.
3. Mount the primary volume on to the mount point directory for the application file system.

Example Configuration

This section provides a step-by-step example of how data replication was configured for an NFS application by using Sun StorEdge Availability Suite software.

Connecting and Installing the Clusters

[Figure 6-7](#) illustrates the cluster configuration used in the example configuration. The secondary cluster in the example configuration contains one node, but other cluster configurations can be used.




FIGURE 6-7 Example Cluster Configuration

Table 6-1 summarizes the hardware and software required by the example configuration. The Solaris OS, Sun Cluster software, and volume manager software must be installed on the cluster nodes *before* you install Sun StorEdge Availability Suite software and patches.

TABLE 6-1 Required Hardware and Software

Hardware or Software	Requirement
Node hardware	Sun StorEdge Availability Suite software is supported on all servers using the Solaris OS. For information about which hardware to use, see the <i>Sun Cluster 3.x Hardware Administration Manual</i> .
Disk space	Approximately 11 Mbytes.

TABLE 6-1 Required Hardware and Software (Continued)

Hardware or Software	Requirement
Solaris OS	Solaris 8 or Solaris 9 releases that are supported by Sun Cluster software. All nodes must use the same version of the Solaris OS. For information about installation, see “Installing the Software” on page 46.
Sun Cluster software	Sun Cluster 3.1 9/04 software. For information about installation, see Chapter 2 and “How to Install Sun Cluster Software on a Single-Node Cluster” on page 87.
Volume manager software	Solstice DiskSuite or Solaris Volume Manager software or VERITAS Volume Manager (VxVM) software. All nodes must use the same version of volume manager software. Information about installation is in “Installing and Configuring Solstice DiskSuite or Solaris Volume Manager Software” on page 129 and “SPARC: Installing and Configuring VxVM Software” on page 163.
Sun StorEdge Availability Suite software	For information about how to install the software, see the <i>Sun StorEdge Availability Suite 3.1 Point-in-Time Copy Software Installation Guide</i> and the <i>Sun StorEdge Availability Suite 3.1 Remote Mirror Software Installation Guide</i> .
Sun StorEdge Availability Suite software patches	For information about the latest patches, see http://sunsolve.sun.com .

Example of How to Configure Device Groups and Resource Groups

This chapter describes how disk device groups and resource groups were configured for an NFS application. The following table lists the names of the groups and resources that were created for the example configuration.

TABLE 6-2 Summary of the Groups and Resources in the Example Configuration

Group or Resource	Name	Description
Disk device group	<i>devicegroup</i>	The disk device group.

TABLE 6-2 Summary of the Groups and Resources in the Example Configuration
(Continued)

Group or Resource	Name	Description
Replication resource group and resources	<i>devicegroup-stor-rg</i>	The replication resource group.
	<i>lhost-reprg-prim,</i> <i>lhost-reprg-sec</i>	The logical hostnames for the replication resource group on the primary cluster and the secondary cluster.
	<i>devicegroup-stor</i>	The HAStoragePlus resource for the replication resource group.
Application resource group and resources	<i>nfs-rg</i>	The application resource group.
	<i>lhost-nfsrg-prim,</i> <i>lhost-nfsrg-sec</i>	The logical hostnames for the application resource group on the primary cluster and the secondary cluster.
	<i>nfs-dg-rs</i>	The HAStoragePlus resource for the application.
	<i>nfs-rs</i>	The NFS resource.

With the exception of *devicegroup-stor-rg*, the names of the groups and resources are example names that can be changed as required. The replication resource group must have a name with the format *devicegroup-stor-rg*.

Configuring Disk Device Groups

This section describes how to configure a disk device group on the primary cluster and the secondary cluster. This example configuration uses VxVM software. For information about Solstice DiskSuite or Solaris Volume Manager software, see [Chapter 3](#).

The following figure illustrates the volumes that were created in the disk device group.




FIGURE 6-8 Volumes for the Disk Device Group

Note – The volumes defined in this section must not include disk label private areas, for example, cylinder 0. The VxVM software manages this constraint automatically.

▼ *How to Configure a Disk Device Group on the Primary Cluster*

1. **Create a disk group that contains four volumes, volume 1 through volume 4.**
For information about configuring a disk group by using the VxVM software, see [Chapter 4](#).
2. **Access `nodeA` as superuser.**
`nodeA` is the first node of the primary cluster. For a reminder of which node is `nodeA`, see [Figure 6-7](#).
3. **Configure the disk group to create a disk device group.**

```
nodeA# /usr/cluster/bin/scconf -a -D type=vxvm,name=devicegroup \
,nodeList=nodeA:nodeB
```

The disk device group is called *devicegroup*.

4. Start the disk device group.

```
nodeA# /usr/cluster/bin/scswitch -z -D devicegroup -h nodeA
```

5. Synchronize the disk device group with the Sun Cluster software.

```
nodeA# /usr/cluster/bin/scconf -c -D name=devicegroup, sync
```

6. Create the file system for the disk device group.

```
nodeA# /usr/sbin/newfs /dev/vx/rdisk/devicegroup/vol101 < /dev/null
nodeA# /usr/sbin/newfs /dev/vx/rdisk/devicegroup/vol102 < /dev/null
nodeA# /usr/sbin/newfs /dev/vx/rdisk/devicegroup/vol103 < /dev/null
nodeA# /usr/sbin/newfs /dev/vx/rdisk/devicegroup/vol104 < /dev/null
```

7. Enable remote access between the nodes in the primary cluster and secondary cluster by adding the following entities to the `.rhosts` file on `nodeA` and `nodeB`.

```
nodeC +
+ root
```

▼ *How to Configure a Disk Device Group on the Secondary Cluster*

● Follow the procedure in “How to Configure a Disk Device Group on the Primary Cluster” on page 234, with these exceptions:

- Replace `nodeA` with `nodeC`.
- Do not use `nodeB`.
- In Step 3, include `nodeC` only in the node list. For example:

```
nodeC# /usr/cluster/bin/scconf -a -D type=vxvm,name=devicegroup \
, nodelist=nodeC
```

- In Step 7, add the following entities to the `.rhosts` file on `nodeC` only:

```
nodeA +
nodeB +
+ root
```

Configuring File Systems for the NFS Application

This section describes how the file systems were configured for the NFS application.

▼ *How to Configure the File System on the Primary Cluster for the NFS Application*

1. On `nodeA` and `nodeB`, create a mount point directory for the NFS file system.

For example:

```
nodeA# mkdir /global/mountpoint
```

2. **On nodeA and nodeB, configure the master volume to be mounted automatically on the mount point.**

Add or replace the following text to the `/etc/vfstab` file on nodeA and nodeB. The text must be on a single line.

```
/dev/vx/dsk/devicegroup/vol01 /dev/vx/rdisk/devicegroup/vol01 \  
/global/mountpoint ufs 3 no global,logging
```

For a reminder of the volumes names and volume numbers used in the disk device group, see [Figure 6–8](#).

3. **On nodeA, create a volume for the file system information that is used by Sun StorEdge Availability Suite software.**

```
nodeA# /usr/sbin/vxassist -g devicegroup make vol05 120m disk1
```

Volume 5 contains the file system information that is used by Sun StorEdge Availability Suite software.

4. **On nodeA, resynchronize the device group with the Sun Cluster software.**

```
nodeA# /usr/cluster/bin/scconf -c -D name=devicegroup, sync
```

5. **On nodeA, create the file system for volume 5.**

```
nodeA# /usr/sbin/newfs /dev/vx/rdisk/devicegroup/vol05
```

6. **On nodeA and nodeB, create a mount point for volume 5.**

For example:

```
nodeA# mkdir /global/etc
```

7. **On nodeA and nodeB, configure volume 5 to be mounted automatically on the mount point.**

Add or replace the following text to the `/etc/vfstab` file on nodeA and nodeB. The text must be on a single line.

```
/dev/vx/dsk/devicegroup/vol05 /dev/vx/rdisk/devicegroup/vol05 \  
/global/etc ufs 3 yes global,logging
```

8. **Mount volume 5 on nodeA.**

```
nodeA# mount /global/etc
```

9. **Make volume 5 accessible to remote systems.**

- a. **Create a directory called `/global/etc/SUNW.nfs` on nodeA.**

```
nodeA# mkdir -p /global/etc/SUNW.nfs
```

- b. **Create the file `/global/etc/SUNW.nfs/dfstab.nfs-rs` on nodeA.**

```
nodeA# touch /global/etc/SUNW.nfs/dfstab.nfs-rs
```

- c. **Add the following line to the `/global/etc/SUNW.nfs/dfstab.nfs-rs` file on nodeA:**

```
share -F nfs -o rw -d "HA NFS" /global/mountpoint
```

▼ *How to Configure the File System on the Secondary Cluster for the NFS Application*

- Repeat the procedure in “[How to Configure the File System on the Primary Cluster for the NFS Application](#)” on page 235, with these exceptions:
 - Replace `nodeA` with `nodeC`.
 - Do not use `nodeB`.

Creating Replication Resource Groups

This section describes how a replication resource group was created on the primary cluster and on the secondary cluster.

▼ *How to Create a Replication Resource Group on the Primary Cluster*

1. Access `nodeA` as superuser.
2. Register `SUNW.HAStoragePlus` as a resource type.

```
nodeA# /usr/cluster/bin/scrgadm -a -t SUNW.HAStoragePlus
```

3. Create a replication resource group for the disk device group.

```
nodeA# /usr/cluster/bin/scrgadm -a -g devicegroup-stor-rg -h nodeA,nodeB  
devicegroup
```

The name of the disk device group.

```
devicegroup-stor-rg
```

The name of the replication resource group.

```
-h nodeA, nodeB
```

Specifies the cluster nodes that can master the replication resource group.

4. Add a `SUNW.HAStoragePlus` resource to the replication resource group.

```
nodeA# /usr/cluster/bin/scrgadm -a -j devicegroup-stor \  
-g devicegroup-stor-rg -t SUNW.HAStoragePlus \  
-x GlobalDevicePaths=devicegroup \  
-x AffinityOn=True
```

```
devicegroup-stor
```

The `HAStoragePlus` resource for the replication resource group.

```
-x GlobalDevicePaths=
```

Specifies the extension property that Sun StorEdge Availability Suite software relies on.

-x AffinityOn=True

Specifies that the SUNW.HAStoragePlus resource must perform an affinity switchover for the global devices and cluster file systems defined by -x GlobalDevicePaths=. Therefore, when the replication resource group fails over or is switched over, the associated device group is switched over.

For more information about these extension properties, see the SUNW.HAStoragePlus(5) man page.

5. Add a logical hostname resource to the replication resource group.

```
nodeA# /usr/cluster/bin/scrgadm -a -L \  
-j lhost-reprg-prim -g devicegroup-stor-rg -l lhost-reprg-prim
```

Where *lhost-reprg-prim* is the logical hostname for the replication resource group on the primary cluster.

6. Enable the resources, manage the resource group, and bring the resource group online.

```
nodeA# /usr/cluster/bin/scswitch -Z -g devicegroup-stor-rg  
nodeA# /usr/cluster/bin/scswitch -z -g devicegroup-stor-rg -h nodeA
```

7. Verify that the resource group is online.

```
nodeA# /usr/cluster/bin/scstat -g
```

Examine the resource group state field to confirm that the replication resource group is online for nodeA and nodeB.

▼ *How to Create a Replication Resource Group on the Secondary Cluster*

● Repeat the procedure in “How to Create a Replication Resource Group on the Primary Cluster” on page 237, with these exceptions:

- Replace nodeA with nodeC.
- Do not use nodeB.
- Replace references to *lhost-reprg-prim* with *lhost-reprg-sec*.

Creating Application Resource Groups

This section describes how application resource groups were created for an NFS application. The procedures in this section are specific to the application. The procedures cannot be used for another type of application.

▼ *How to Create an Application Resource Group on the Primary Cluster*

1. Access nodeA as superuser.
2. Register SUNW.nfs as a resource type.

```
nodeA# scrgadm -a -t SUNW.nfs
```

3. If `SUNW.HAStoragePlus` has not been registered as a resource type, register it.

```
nodeA# scrgadm -a -t SUNW.HAStoragePlus
```

4. Create an application resource group for the *devicegroup*.

```
nodeA# scrgadm -a -g nfs-rg \  
-y Pathprefix=/global/etc \  
-y Auto_start_on_new_cluster=False \  
-y RG_dependencies=devicegroup-stor-rg
```

nfs-rg

Is the name of the application resource group.

Pathprefix=/global/etc

Specifies a directory into which the resources in the group can write administrative files.

Auto_start_on_new_cluster=False

Specifies that the application resource group is not started automatically.

RG_dependencies=devicegroup-stor-rg

Specifies the resource groups that the application resource group depends on. In this example, the application resource group depends on the replication resource group.

If the application resource group is switched over to a new primary node, the replication resource group is automatically switched over. However, if the replication resource group is switched over to a new primary node, the application resource group must be manually switched over.

5. Add a `SUNW.HAStoragePlus` resource to the application resource group.

```
nodeA# scrgadm -a -j nfs-dg-rs -g nfs-rg \  
-t SUNW.HAStoragePlus \  
-x FileSystemMountPoints=/global/mountpoint \  
-x AffinityOn=True
```

nfs-dg-rs

Is the name of the HAStoragePlus resource for the NFS application.

-x FileSystemMountPoints=/global/

Specifies that the mount point for the file system is global.

-t SUNW.HAStoragePlus

Specifies that the resource is of the type `SUNW.HAStoragePlus`.

-x AffinityOn=True

Specifies that the application resource must perform an affinity switchover for the global devices and cluster file systems defined by `-x GlobalDevicePaths=`. Therefore, when the application resource group fails over or is switched over, the associated device group is switched over.

For more information about these extension properties, see the `SUNW.HAStoragePlus(5)` man page.

6. Add a logical hostname resource to the application resource group.

```
nodeA# /usr/cluster/bin/scrgadm -a -L -j lhost-nfsrg-prim -g nfs-rg \  
-l lhost-nfsrg-prim
```

Where *lhost-nfsrg-prim* is the logical hostname of the application resource group on the primary cluster.

7. Enable the resources, manage the application resource group, and bring the application resource group online.

a. Bring the HAStoragePlus resource for the NFS application online.

```
nodeA# /usr/cluster/bin/scrgadm -a -g nfs-rg \  
-j nfs-rs -t SUNW.nfs -y Resource_dependencies=nfs-dg-rs
```

b. Bring the application resource group online on nodeA .

```
nodeA# /usr/cluster/bin/scswitch -Z -g nfs-rg  
nodeA# /usr/cluster/bin/scswitch -z -g nfs-rg -h nodeA
```

8. Verify that the application resource group is online.

```
nodeA# /usr/cluster/bin/scstat -g
```

Examine the resource group state field to determine whether the application resource group is online for nodeA and nodeB.

▼ *How to Create an Application Resource Group on the Secondary Cluster*

1. Create the application group resource as described in Step 1 through Step 6 of “How to Create an Application Resource Group on the Primary Cluster” on page 238, with the following exceptions:

- Replace nodeA with nodeC.
- Ignore references to nodeB.
- Replace references to *lhost-nfsrg-prim* with *lhost-nfsrg-sec*.

2. Ensure that the application resource group does not come online on nodeC.

```
nodeC# /usr/cluster/bin/scswitch -n -j nfs-rs  
nodeC# /usr/cluster/bin/scswitch -n -j nfs-dg-rs  
nodeC# /usr/cluster/bin/scswitch -n -j lhost-nfsrg-sec  
nodeC# /usr/cluster/bin/scswitch -z -g nfs-rg -h ""
```

The resource group remains offline after a reboot, because *Auto_start_on_new_cluster=False*.

3. If the global volume is mounted on the primary cluster, unmount the global volume from the secondary cluster.

```
nodeC# umount /global/mountpoint
```

If the volume is mounted on a secondary cluster, the synchronization fails.

Example of How to Enable Data Replication

This section describes how data replication was enabled for the example configuration. This section uses the Sun StorEdge Availability Suite software commands `sndradm` and `iiadm`. For more information about these commands, see the *Sun Cluster 3.0 and Sun StorEdge Software Integration Guide*.

▼ How to Enable Replication on the Primary Cluster

1. Access `nodeA` as superuser.

2. Flush all transactions.

```
nodeA# /usr/sbin/lockfs -a -f
```

3. Confirm that the logical hostnames `lhost-reprg-prim` and `lhost-reprg-sec` are online.

```
nodeA# /usr/cluster/bin/scstat -g
```

Examine the state field of the resource group.

4. Enable remote mirror replication from the primary cluster to the secondary cluster.

This step enables replication from the master volume of the primary cluster to the master volume of the secondary cluster. In addition, this step enables replication to the remote mirror bitmap on volume 4.

- If the primary cluster and secondary cluster are unsynchronized, run this command:

```
nodeA# /usr/opt/SUNWesm/sbin/sndradm -n -e lhost-reprg-prim \  
/dev/vx/rdisk/devicegroup/vol101 \  
/dev/vx/rdisk/devicegroup/vol104 lhost-reprg-sec \  
/dev/vx/rdisk/devicegroup/vol101 \  
/dev/vx/rdisk/devicegroup/vol104 ip sync
```

- If the primary cluster and secondary cluster are synchronized, run this command:

```
nodeA# /usr/opt/SUNWesm/sbin/sndradm -n -E lhost-reprg-prim \  
/dev/vx/rdisk/devicegroup/vol101 \  
/dev/vx/rdisk/devicegroup/vol104 lhost-reprg-sec \  
/dev/vx/rdisk/devicegroup/vol101 \  
/dev/vx/rdisk/devicegroup/vol104 ip sync
```

5. Enable autosynchronization.

```
nodeA# /usr/opt/SUNWesm/sbin/sndradm -n -a on lhost-reprg-prim \  
/dev/vx/rdisk/devicegroup/vol101 \  
/dev/vx/rdisk/devicegroup/vol104 lhost-reprg-sec \  
/dev/vx/rdisk/devicegroup/vol101 \  
/dev/vx/rdisk/devicegroup/vol104 ip sync
```

This step enables autosynchronization. When the active state of autosynchronization is set to on, the volume sets are resynchronized if the system

reboots or a failure occurs.

6. Verify that the cluster is in logging mode.

```
nodeA# /usr/opt/SUNWesm/sbin/sndradm -P
```

The output should look like this:

```
/dev/vx/rdisk/devicegroup/vol01 ->  
lhost-reprg-sec:/dev/vx/rdisk/devicegroup/vol01  
autosync: off, max q writes:4194304, max q fbas:16384, mode:sync,ctag:  
devicegroup, state: logging
```

In logging mode, the state is logging, and the active state of autosynchronization is off. When the data volume on the disk is written to, the bitmap file on the same disk is updated.

7. Enable point-in-time snapshot.

```
nodeA# /usr/opt/SUNWesm/sbin/iiadm -e ind \  
/dev/vx/rdisk/devicegroup/vol01 \  
/dev/vx/rdisk/devicegroup/vol02 \  
/dev/vx/rdisk/devicegroup/vol03  
nodeA# /usr/opt/SUNWesm/sbin/iiadm -w \  
/dev/vx/rdisk/devicegroup/vol02
```

This step enables the master volume of the primary disk to be copied to the shadow volume on the same disk. In this example, the master volume is volume 1, the shadow volume is volume 2, and the point-in-time bitmap volume is volume 3.

8. Attach the point-in-time snapshot to the remote mirror set.

```
nodeA# /usr/opt/SUNWesm/sbin/sndradm -I a \  
/dev/vx/rdisk/devicegroup/vol01 \  
/dev/vx/rdisk/devicegroup/vol02 \  
/dev/vx/rdisk/devicegroup/vol03
```

This step associates the point-in-time snapshot with the remote mirror volume set. Sun StorEdge Availability Suite software ensures that a point-in-time snapshot is taken before remote mirror replication can occur.

▼ How to Enable Replication on the Secondary Cluster

1. Access nodeC as superuser.

2. Flush all transactions.

```
nodeC# /usr/sbin/lockfs -a -f
```

3. Enable remote mirror replication from the primary cluster to the secondary cluster.

```
nodeC# /usr/opt/SUNWesm/sbin/sndradm -n -e lhost-reprg-prim \  
/dev/vx/rdisk/devicegroup/vol01 \  
/dev/vx/rdisk/devicegroup/vol04 lhost-reprg-sec \  
/dev/vx/rdisk/devicegroup/vol04
```

```
/dev/vx/rdisk/devicegroup/vol01 \  
/dev/vx/rdisk/devicegroup/vol04 ip sync
```

The primary cluster detects the presence of the secondary cluster and starts synchronization. Refer to the system log file `/var/opt/SUNWesm/ds.log` for information about the status of the clusters.

4. Enable independent point-in-time snapshot.

```
nodeC# /usr/opt/SUNWesm/sbin/iiadm -e ind \  
/dev/vx/rdisk/devicegroup/vol01 \  
/dev/vx/rdisk/devicegroup/vol02 \  
/dev/vx/rdisk/devicegroup/vol03  
nodeC# /usr/opt/SUNWesm/sbin/iiadm -w \  
/dev/vx/rdisk/devicegroup/vol02
```

5. Attach the point-in-time snapshot to the remote mirror set.

```
nodeC# /usr/opt/SUNWesm/sbin/sndradm -I a \  
/dev/vx/rdisk/devicegroup/vol01 \  
/dev/vx/rdisk/devicegroup/vol02 \  
/dev/vx/rdisk/devicegroup/vol03
```

Example of How to Perform Data Replication

This section describes how data replication was performed for the example configuration. This section uses the Sun StorEdge Availability Suite software commands `sndradm` and `iiadm`. For more information about these commands, see the *Sun Cluster 3.0 and Sun StorEdge Software Integration Guide*.

▼ How to Perform a Remote Mirror Replication

In this procedure, the master volume of the primary disk is replicated to the master volume on the secondary disk. The master volume is volume 1 and the remote mirror bitmap volume is volume 4.

1. Access `nodeA` as superuser.
2. Verify that the cluster is in logging mode.

```
nodeA# /usr/opt/SUNWesm/sbin/sndradm -P
```

The output should look like this:

```
/dev/vx/rdisk/devicegroup/vol01 ->  
lhost-reprg-sec:/dev/vx/rdisk/devicegroup/vol01  
autosync: off, max q writes:4194304, max q fbas:16384, mode:sync,ctag:  
devicegroup, state: logging
```

In logging mode, the state is `logging`, and the active state of autosynchronization is `off`. When the data volume on the disk is written to, the bitmap file on the same disk is updated.

3. Flush all transactions.

```
nodeA# /usr/sbin/lockfs -a -f
```

4. Repeat Step 1 through Step 3 on nodeC.

5. Copy the master volume of nodeA to the master volume of nodeC.

```
nodeA# /usr/opt/SUNWesm/sbin/sndradm -n -m lhost-reprg-prim \  
/dev/vx/rdisk/devicegroup/vol101 \  
/dev/vx/rdisk/devicegroup/vol104 lhost-reprg-sec \  
/dev/vx/rdisk/devicegroup/vol101 \  
/dev/vx/rdisk/devicegroup/vol104 ip sync
```

6. Wait until the replication is complete and the volumes are synchronized.

```
nodeA# /usr/opt/SUNWesm/sbin/sndradm -n -w lhost-reprg-prim \  
/dev/vx/rdisk/devicegroup/vol101 \  
/dev/vx/rdisk/devicegroup/vol104 lhost-reprg-sec \  
/dev/vx/rdisk/devicegroup/vol101 \  
/dev/vx/rdisk/devicegroup/vol104 ip sync
```

7. Confirm that the cluster is in replicating mode.

```
nodeA# /usr/opt/SUNWesm/sbin/sndradm -P
```

The output should look like this:

```
/dev/vx/rdisk/devicegroup/vol101 ->  
lhost-reprg-sec:/dev/vx/rdisk/devicegroup/vol101  
autosync: on, max q writes:4194304, max q fbas:16384, mode:sync,ctag:  
devicegroup, state: replicating
```

In replicating mode, the state is `replicating`, and the active state of autosynchronization is `on`. When the primary volume is written to, the secondary volume is updated by Sun StorEdge Availability Suite software.

▼ How to Perform a Point-in-Time Snapshot

In this procedure, point-in-time snapshot was used to synchronize the shadow volume of the primary cluster to the master volume of the primary cluster. The master volume is volume 1 and the shadow volume is volume 2.

1. Access nodeA as superuser.

2. Quiesce the application that is running on nodeA.

```
nodeA# /usr/cluster/bin/scswitch -n -j nfs-rs
```

3. Put the primary cluster in to logging mode.

```
nodeA# /usr/opt/SUNWesm/sbin/sndradm -n -l lhost-reprg-prim \  
/dev/vx/rdisk/devicegroup/vol101 \  
/dev/vx/rdisk/devicegroup/vol104 lhost-reprg-sec \  
/dev/vx/rdisk/devicegroup/vol101 \  
/dev/vx/rdisk/devicegroup/vol101 \  
/dev/vx/rdisk/devicegroup/vol101
```

```
/dev/vx/rdisk/devicegroup/vol04 ip sync
```

When the data volume on the disk is written to, the bitmap file on the same disk is updated. No replication occurs.

4. Synchronize the shadow volume of the primary cluster to the master volume of the primary cluster.

```
nodeA# /usr/opt/SUNWesm/sbin/iiadm -u s /dev/vx/rdisk/devicegroup/vol02
nodeA# /usr/opt/SUNWesm/sbin/iiadm -w /dev/vx/rdisk/devicegroup/vol02
```

5. Synchronize the shadow volume of the secondary cluster to the master volume of the secondary cluster.

```
nodeC# /usr/opt/SUNWesm/sbin/iiadm -u s /dev/vx/rdisk/devicegroup/vol02
nodeC# /usr/opt/SUNWesm/sbin/iiadm -w /dev/vx/rdisk/devicegroup/vol02
```

6. Restart the application on nodeA.

```
nodeA# /usr/cluster/bin/scswitch -e -j nfs-rs
```

7. Resynchronize the secondary volume with the primary volume.

```
nodeA# /usr/opt/SUNWesm/sbin/sndradm -n -u lhost-reprg-prim \
/dev/vx/rdisk/devicegroup/vol01 \
/dev/vx/rdisk/devicegroup/vol04 lhost-reprg-sec \
/dev/vx/rdisk/devicegroup/vol01 \
/dev/vx/rdisk/devicegroup/vol04 ip sync
```

Example of How to Verify That Replication Is Configured Correctly

This section describes how the replication configuration was confirmed in the example configuration.

▼ How to Verify That Replication Is Configured Correctly

1. Verify that the primary cluster is in replicating mode, with autosynchronization on.

```
nodeA# /usr/opt/SUNWesm/sbin/sndradm -P
```

The output should look like this:

```
/dev/vx/rdisk/devicegroup/vol01 ->
lhost-reprg-sec:/dev/vx/rdisk/devicegroup/vol01
autosync: on, max q writes:4194304, max q fbas:16384, mode:sync,ctag:
devicegroup, state: replicating
```

In replicating mode, the state is `replicating`, and the active state of autosynchronization is on. When the primary volume is written to, the secondary volume is updated by Sun StorEdge Availability Suite software.

If the primary cluster is not in replicating mode, put it in to replicating mode, as follows:

```
nodeA# /usr/opt/SUNWesm/sbin/sndradm -n -u lhost-reprg-prim \  
/dev/vx/rdisk/devicegroup/vol101 \  
/dev/vx/rdisk/devicegroup/vol104 lhost-reprg-sec \  
/dev/vx/rdisk/devicegroup/vol101 \  
/dev/vx/rdisk/devicegroup/vol104 ip sync
```

2. Make a directory on a client machine.

a. Log in to a client machine as superuser.

You see a prompt like this:

```
client-machine#
```

b. Make a directory on the client machine.

```
client-machine# mkdir /dir
```

3. Mount the directory to the application on the primary cluster, and display the mounted directory.

a. Mount the directory to the application on the primary cluster.

```
client-machine# mount -o rw lhost-nfsrg-prim:/global/mountpoint /dir
```

b. Display the mounted directory.

```
client-machine# ls /dir
```

4. Mount the directory to the application on the secondary cluster, and display the mounted directory.

a. Unmount the directory to the application on the primary cluster.

```
client-machine# umount /dir
```

b. Take the application resource group offline on the primary cluster.

```
nodeA# /usr/cluster/bin/scswitch -n -j nfs-rs  
nodeA# /usr/cluster/bin/scswitch -n -j nfs-dg-rs  
nodeA# /usr/cluster/bin/scswitch -n -j lhost-nfsrg-prim  
nodeA# /usr/cluster/bin/scswitch -z -g nfs-rg -h ""
```

c. Put the primary cluster in to logging mode.

```
nodeA# /usr/opt/SUNWesm/sbin/sndradm -n -l lhost-reprg-prim \  
/dev/vx/rdisk/devicegroup/vol101 \  
/dev/vx/rdisk/devicegroup/vol104 lhost-reprg-sec \  
/dev/vx/rdisk/devicegroup/vol101 \  
/dev/vx/rdisk/devicegroup/vol104 ip sync
```

When the data volume on the disk is written to, the bitmap file on the same disk is updated. No replication occurs.

d. Bring the application resource group online on the secondary cluster.

```
nodeC# /usr/cluster/bin/scswitch -Z -g nfs-rg
```

e. Access the client machine as superuser.

You see a prompt like this:

```
client-machine#
```

f. Mount the directory that was created in Step 2 to the application on the secondary cluster.

```
client-machine# mount -o rw lhost-nfsrg-sec:/global/mountpoint /dir
```

g. Display the mounted directory.

```
client-machine# ls /dir
```

5. Ensure that the directory displayed in Step 3 is the same as that displayed in Step 4.

6. Return the application on the primary cluster to the mounted directory.

a. Take the application resource group offline on the secondary cluster.

```
nodeC# /usr/cluster/bin/scswitch -n -j nfs-rs
nodeC# /usr/cluster/bin/scswitch -n -j nfs-dg-rs
nodeC# /usr/cluster/bin/scswitch -n -j lhost-nfsrg-sec
nodeC# /usr/cluster/bin/scswitch -z -g nfs-rg -h ""
```

b. Ensure that the global volume is unmounted from the secondary cluster.

```
nodeC# umount /global/mountpoint
```

c. Bring the application resource group online on the primary cluster.

```
nodeA# /usr/cluster/bin/scswitch -Z -g nfs-rg
```

d. Put the primary cluster in to replicating mode.

```
nodeA# /usr/opt/SUNWesm/sbin/sndradm -n -u lhost-reprg-prim \  
/dev/vx/rdisk/devicegroup/vol101 \  
/dev/vx/rdisk/devicegroup/vol104 lhost-reprg-sec \  
/dev/vx/rdisk/devicegroup/vol101 \  
/dev/vx/rdisk/devicegroup/vol104 ip sync
```

When the primary volume is written to, the secondary volume is updated by Sun StorEdge Availability Suite software.

Example of How to Cope With a Failover or Switchover

This section describes how a switchover was provoked and how the application was transferred to the secondary cluster. After a switchover or failover, you must update the DNS entry and configure the application to read and write to the secondary volume.

▼ How to Provoke a Switchover

1. Put the primary cluster into logging mode.

```
nodeA# /usr/opt/SUNWesm/sbin/sndradm -n -l lhost-reprg-prim \  
/dev/vx/rdisk/devicegroup/vol101 \  
/dev/vx/rdisk/devicegroup/vol104 lhost-reprg-sec \  
/dev/vx/rdisk/devicegroup/vol101 \  
/dev/vx/rdisk/devicegroup/vol104 ip sync
```

When the data volume on the disk is written to, the bitmap file on the same disk is updated. No replication occurs.

2. Confirm that the primary cluster and the secondary cluster are in logging mode, with autosynchronization off.

a. On nodeA, run this command:

```
nodeA# /usr/opt/SUNWesm/sbin/sndradm -P
```

The output should look like this:

```
/dev/vx/rdisk/devicegroup/vol101 ->  
lhost-reprg-sec: /dev/vx/rdisk/devicegroup/vol101  
autosync:off, max q writes:4194304,max q fbas:16384,mode:sync,ctag:  
devicegroup, state: logging
```

b. On nodeC, run this command:

```
nodeC# /usr/opt/SUNWesm/sbin/sndradm -P
```

The output should look like this:

```
/dev/vx/rdisk/devicegroup/vol101 <-  
lhost-reprg-prim: /dev/vx/rdisk/devicegroup/vol101  
autosync:off, max q writes:4194304,max q fbas:16384,mode:sync,ctag:  
devicegroup, state: logging
```

For nodeA and nodeC, the state should be logging, and the active state of autosynchronization should be off.

3. Confirm that the secondary cluster is ready to take over from the primary cluster.

```
nodeC# /usr/sbin/fsck -y /dev/vx/rdisk/devicegroup/vol101
```

4. Switch over to the secondary cluster.

```
nodeC# scswitch -Z -g nfs-rg  
nodeC# scswitch -Z -g nfs-rg -h nodeC
```

▼ How to Update the DNS Entry

For an illustration of how the DNS maps a client to a cluster, see [Figure 6-6](#).

1. Start the nsupdate command.

For information, see the nsupdate(1M) man page.

2. **Remove the current DNS mapping between the client machine and the logical hostname of the application resource group on the primary cluster.**

```
> update delete client-machine A
> update delete IPaddress1.in-addr.arpa TTL PTR client-machine
```

client-machine

Is the full name of the client. For example, `mymachine.mycompany.com`.

IPaddress1

Is the IP address is of the logical hostname `lhost-nfsrg-prim`, in reverse order.

TTL

Is the time to live, in seconds. A typical value is 3600.

3. **Create the new DNS mapping between the client machine and the logical hostname of the application resource group on the secondary cluster.**

```
> update add client-machine TTL A IPaddress2
> update add IPaddress3.in-addr.arpa TTL PTR client-machine
```

IPaddress2

Is the IP address is of the logical hostname `lhost-nfsrg-sec`, in forward order.

IPaddress3

Is the IP address is of the logical hostname `lhost-nfsrg-sec`, in reverse order.

▼ How to Configure the Application to Read and Write to the Secondary Volume

1. **Configure the secondary volume to be mounted to the mount point directory for the NFS file system.**

```
client-machine# mount -o rw lhost-nfsrg-sec:/global/mountpoint /xxx
```

The mount point was created in [Step 1](#) of “[How to Configure the File System on the Primary Cluster for the NFS Application](#)” on page 235.

2. **Confirm that the secondary cluster has write access to the mount point.**

```
client-machine# touch /xxx/data.1
client-machine# umount /xxx
```


Sun Cluster Installation and Configuration Worksheets

This appendix provides worksheets to plan various components of your cluster configuration and examples of completed worksheets for your reference. See “Installation and Configuration Worksheets” in *Sun Cluster Data Services Planning and Administration Guide for Solaris OS* for configuration worksheets for resources, resource types, and resource groups.

Installation and Configuration Worksheets

If necessary, make additional copies of a worksheet to accommodate all the components in your cluster configuration. Follow planning guidelines in [Chapter 1](#) to complete these worksheets. Then refer to your completed worksheets during cluster installation and configuration.

Note – The data used in the worksheet examples is intended as a guide only. The examples do not represent a complete configuration of a functional cluster.

The following table lists the planning worksheets and examples provided in this appendix, as well as the titles of sections in [Chapter 1](#) that contain related planning guidelines.

TABLE A-1 Cluster Installation Worksheets and Related Planning Guidelines

Worksheet	Example	Section Titles of Related Planning Guidelines
"Local File System Layout Worksheet" on page 254	"Example: Local File System Layout Worksheets, With and Without Mirrored Root" on page 255	"System Disk Partitions" on page 18 "Guidelines for Mirroring the Root Disk" on page 39
"Public Networks Worksheet" on page 256	"Example: Public Networks Worksheet" on page 257	"Public Networks" on page 23 "IP Network Multipathing Groups" on page 29
"Local Devices Worksheets" on page 258	"Example: Local Devices Worksheets" on page 259	---
"Disk Device Group Configurations Worksheet" on page 260	"Example: Disk Device Group Configurations Worksheet" on page 261	"Disk Device Groups" on page 31 "Planning Volume Management" on page 33
"Volume-Manager Configurations Worksheet" on page 262	"Example: Volume-Manager Configurations Worksheet" on page 263	"Planning Volume Management" on page 33 Your volume manager documentation

TABLE A-1 Cluster Installation Worksheets and Related Planning Guidelines (Continued)

Worksheet	Example	Section Titles of Related Planning Guidelines
"Metadevices Worksheet (Solstice DiskSuite or Solaris Volume Manager)" on page 264	"Example: Metadevices Worksheet (Solstice DiskSuite or Solaris Volume Manager)" on page 265	"Planning Volume Management" on page 33 <i>Solstice DiskSuite 4.2.1 Installation and Product Notes</i> or <i>Solaris Volume Manager Administration Guide</i>

Local File System Layout Worksheet

Node name: _____

TABLE A-2 Local File Systems With Mirrored Root Worksheet

Volume Name	Component	Component	File System	Size
			/	
			swap	
			/globaldevices	

TABLE A-3 Local File Systems with Nonmirrored Root Worksheet

Device Name	File System	Size
	/	
	swap	
	/globaldevices	

Example: Local File System Layout Worksheets, With and Without Mirrored Root

Node name: **phys-schost-1**

TABLE A-4 Example: Local File Systems With Mirrored Root Worksheet

Volume Name	Component	Component	File System	Size
d1	c0t0d0s0	c1t0d0s0	/	6.75 GB
d2	c0t0d0s1	c1t0d0s1	swap	750 MB
d3	c0t0d0s3	c1t0d0s3	/globaldevices	512 MB
d7	c0t0d0s7	c1t0d0s7	SDS replica	20 MB

TABLE A-5 Example: Local File Systems With Nonmirrored Root Worksheet

Device Name	File System	Size
c0t0d0s0	/	6.75 GB
c0t0d0s1	swap	750 MB
c0t0d0s3	/globaldevices	512 MB
c0t0d0s7	SDS replica	20 MB

Public Networks Worksheet

TABLE A-6 Public Networks Worksheet

Component	Name
Node name	
Primary hostname	
IP Network Multipathing group	
Adapter name	
Backup adapter(s) (optional)	
Network name	
Secondary hostname	
IP Network Multipathing group	
Adapter name	
Backup adapter(s) (optional)	
Network name	
Secondary hostname	
IP Network Multipathing group	
Adapter name	
Backup adapter(s) (optional)	
Network name	
Secondary hostname	
IP Network Multipathing group	
Adapter name	
Backup adapter(s) (optional)	
Network name	

Example: Public Networks Worksheet

TABLE A-7 Example: Public Networks Worksheet

Component	Name
Node name	phys-schost-1
Primary hostname	phys-schost-1
IP Network Multipathing group	ipmp0
Adapter name	qfe0
Backup adapter(s) (optional)	qfe4
Network name	net-85
Secondary hostname	phys-schost-1-86
IP Network Multipathing group	ipmp1
Adapter name	qfe1
Backup adapter(s) (optional)	qfe5
Network name	net-86
Secondary hostname	
IP Network Multipathing group	
Adapter name	
Backup adapter(s) (optional)	
Network name	
Secondary hostname	
IP Network Multipathing group	
Adapter name	
Backup adapter(s) (optional)	
Network name	

Local Devices Worksheets

Node name: _____

TABLE A-8 Local Disks Worksheet

Local Disk Name	Size

TABLE A-9 Other Local Devices Worksheet

Device Type	Name

Example: Local Devices Worksheets

Node name: `phys-schost-1`

TABLE A-10 Example: Local Disks Worksheet

Local Disk Name	Size
<code>c0t0d0</code>	2G
<code>c0t1d0</code>	2G
<code>c1t0d0</code>	2G
<code>c1t1d0</code>	2G

TABLE A-11 Example: Other Local Devices Worksheet

Device Type	Name
tape	<code>/dev/rmt/0</code>

Disk Device Group Configurations Worksheet

Volume manager (circle one):

Solstice DiskSuite | Solaris Volume Manager | VxVM

TABLE A-12 Disk Device Groups Worksheet

Disk Group/ Disk Set Name	Node Names (Indicate priority if ordered list)	Ordered priority? (circle one)	Failback? circle one)
		Yes No	Yes No
		Yes No	Yes No
		Yes No	Yes No
		Yes No	Yes No
		Yes No	Yes No
		Yes No	Yes No
		Yes No	Yes No
		Yes No	Yes No
		Yes No	Yes No
		Yes No	Yes No
		Yes No	Yes No
		Yes No	Yes No
		Yes No	Yes No
		Yes No	Yes No
		Yes No	Yes No

Example: Disk Device Group Configurations Worksheet

Volume manager (circle one):

Solstice DiskSuite

TABLE A-13 Example: Disk Device Groups Configurations Worksheet

Disk Group/ Disk Set Name	Node Names (indicate priority if ordered list)	Ordered priority? (circle one)	Failback? circle one)
dg-schost-1	1) phys-schost-1, 2) phys-schost-2	Yes	Yes
		Yes No	Yes No
		Yes No	Yes No
		Yes No	Yes No
		Yes No	Yes No
		Yes No	Yes No
		Yes No	Yes No
		Yes No	Yes No
		Yes No	Yes No
		Yes No	Yes No
		Yes No	Yes No
		Yes No	Yes No

Example: Metadevices Worksheet (Solstice DiskSuite or Solaris Volume Manager)

TABLE A-17 Example: Metadevices Worksheet (Solstice DiskSuite or Solaris Volume Manager)

File System	Metatrans	Metamirrors		Submirrors		Hot Spare Pool	Physical Device	
		(Data)	(Log)	(Data)	(Log)		(Data)	(Log)
/A	d10	d11		d12, d13		hsp000	c1t0d0s0, c2t0d1s0	
			d14		d15	hsp006		c1t0d1s6, c2t1d1s6

Index

A

- adapters
 - IP Network Multipathing groups
 - requirements, 29
 - test IP addresses, 22
 - local-mac-address? variable
 - changes during upgrade, 194
 - requirements, 24
 - SBus SCI restriction, 28
 - SCI-PCI
 - installing Solaris packages, 55
 - installing Sun Cluster packages, 60
 - package requirements, 17
- adding
 - See also* installing
 - drives to a disk set, 153-155
 - mediator hosts, 160
 - nodes to the Sun Cluster module to Sun Management Center, 121
- administrative console
 - installing CCP software, 48-52
 - IP addresses, 22
 - MANPATH, 51
 - PATH, 51
- affinity switchover
 - configuring, 237
 - extension property, 226
- alternate boot path, displaying, 138
- Apache
 - installing packages
 - by using pkgadd, 67
 - modifying links during upgrade, 189

- application resource groups
 - creating, 238
 - guidelines for configuring, 226
- asynchronous data replication, 223
- authentication, *See* authorized-node list
- authorized-node list
 - adding nodes, 123
 - removing nodes, 95
- autoscinstall.class file, 82
- Availability Suite, preparing for cluster upgrade, 187

B

- backup cluster, role in data replication, 222
- bitmap
 - point-in-time snapshot, 223
 - remote mirror, 222
- boot devices, alternate boot path, 138

C

- C locale, 101
- cconsole command, 51
 - installing the software, 48-52
 - using, 53, 83
- ccp command, 51
- Cluster Control Panel (CCP) software
 - installing, 48-52
 - starting, 51

- cluster file systems
 - See also* shared file systems
 - caution notice, 108
 - communication end-points restriction, 31
 - configuring, 108-113
 - `fattach` command restriction, 31
 - `forcedirectio` restriction, 33
 - LOFS restriction, 31
 - planning, 30-33
 - quotas restriction, 31
 - required mount options, 110
 - verifying the configuration, 112
 - VxFS restrictions, 32
- cluster interconnects
 - configuring on a single-node cluster, 92
 - planning, 28-29
- cluster mode, verifying, 196
- cluster name, 26
- cluster nodes
 - adding to the Sun Cluster module to Sun Management Center, 121
 - determining the node-ID number, 177
 - installing and configuring
 - additional nodes, 90-97
 - by using `JumpStart`, 76-87
 - by using `scinstall`, 57-63
 - by using SunPlex Installer, 70
 - single-node clusters, 87-90
 - planning, 26
 - upgrading
 - nonrolling, 184
 - rolling, 201-214
 - verifying
 - cluster mode, 196
 - installation mode, 89
- `clusters` file, administrative console, 50
- common agent container
 - package installation, 68
 - security files, 68
- communication end-points, restriction on cluster file systems, 31
- configuring
 - cluster file systems, 108-113
 - cluster interconnects on a single-node cluster, 92
 - data replication, 221-249
 - disk sets, 151-153
 - IP Network Multipathing groups, 113-114, 114

- configuring (Continued)
 - `md.tab` file, 156-157
 - name-service switch, 98-99
 - Network Time Protocol (NTP), 116-117
 - quorum devices, 104
 - Solaris Volume Manager, 129-150
 - Solstice DiskSuite, 129-150
 - state database replicas, 135-136
 - user work environment, 99-100
 - VERITAS Volume Manager (VxVM), 163-179
- console-access devices
 - IP addresses, 22
 - planning, 23
 - serial-port numbers, 50
- creating
 - See* adding
 - See* configuring
 - See* installing

D

- data replication
 - configuring disk device groups, 233
 - configuring file systems for an NFS application, 235
 - coping with failover, 247
 - creating resource groups
 - application, 238
 - replication, 237
 - definition, 222
 - enabling, 241
 - example configuration, 230
 - guidelines for configuring, 225
 - in the example configuration, 224
 - introduction to, 222
 - mounting the secondary volume, 249
 - performing, 243
 - point-in-time snapshot, 223, 244
 - remote mirror, 222, 243
 - required hardware and software, 231
 - synchronous and asynchronous, 223
 - updating a DNS entry, 248
 - verifying the configuration, 245
- data services
 - installing
 - by using `scinstall`, 102-104
 - by using SunPlex Installer, 70

- data services (Continued)
 - upgrading
 - nonrolling, 195
 - rolling, 209
 - Sun Cluster HA for SAP liveCache
 - configuration change, 199
 - deporting disk device groups, 173
 - device groups
 - See also* disk device groups
 - See also* raw-disk device groups
 - moving, 173, 204
 - device-ID names
 - determining, 105
 - displaying, 144
 - migrating after upgrade, 215
 - DID driver, updating, 216
 - Dirty Region Logging (DRL), planning, 37
 - disabling
 - installation mode, 106
 - resources, 186
 - disaster tolerance, definition, 222
 - disk device groups
 - See also* device groups
 - See also* disk groups
 - See also* raw-disk device groups
 - configuring, 233
 - importing and deporting, 173
 - planning, 31-32
 - registering changes to, 175
 - registering disk groups as, 174
 - reminoring, 175-176
 - status, 176
 - verifying
 - evacuation, 204
 - registration, 174
 - disk drives, *See* drives
 - disk groups
 - See also* disk device groups
 - configuring, 173-175
 - registering as disk device groups, 174
 - verifying the configuration, 176
 - disk sets
 - adding drives, 153-155
 - configuring, 151-153
 - planning the maximum number, 36
 - repartitioning drives, 155
 - setting the maximum number, 133-134
 - disk strings, dual-string mediator
 - requirements, 159
 - disks, *See* drives
 - disksets, *See* disk sets
 - domain console network interfaces, IP
 - addresses, 22
 - Domain Name System (DNS), 229
 - guidelines for updating, 229
 - updating, 248
 - drives
 - adding to disk sets, 153-155
 - mirroring differing device sizes, 39
 - repartitioning, 155
 - DRL, planning, 37
 - dual-string mediators
 - adding hosts, 160
 - overview, 159-161
 - planning, 35
 - repairing data, 161
 - restoring during upgrade
 - nonrolling, 199
 - rolling, 213
 - status, 160-161
 - unconfiguring during upgrade
 - nonrolling, 186
 - rolling, 205
 - Dynamic Multipathing (DMP), 37
- ## E
- enabling the kernel cage, 55
 - encapsulated root disks
 - configuring, 165-168
 - mirroring, 168-170
 - planning, 37
 - Enclosure-Based Naming, planning, 37
 - error messages
 - cluster, 13
 - metainit command, 136
 - scconf command, 174
 - scdidadm command, 197
 - scgdevs command, 132
 - SunPlex Installer, 75
 - /etc/clusters file, 50
 - /etc/inet/hosts file
 - configuring, 55
 - configuring for JumpStart, 83

- /etc/inet/hosts file (Continued)
 - planning, 22
- /etc/inet/ntp.conf.cluster file
 - configuring, 116-117
 - starting, 117
- /etc/inet/ntp.conf file
 - changes during upgrade, 194
 - configuring, 116-117
 - starting, 117
- /etc/init.d/xntpd.cluster start command, 117
- /etc/init.d/xntpd start command, 117
- /etc/lvm/md.tab file, 156-157
- /etc/name_to_major file
 - non-VxVM nodes, 55, 167, 171
 - VxVM-installed nodes, 167, 171
- /etc/nsswitch.conf file, 98-99
- /etc/release file, 48
- /etc/serialports file, 50
- /etc/system file
 - kernel_cage_enable variable, 55
 - secure ports for NFS, 24
 - stack-size setting, 97
 - thread stack-size setting, 175
- /etc/vfstab file
 - adding mount points, 110
 - modifying during upgrade
 - nonrolling, 190
 - rolling, 206
 - verifying the configuration, 112
- Ethernet adapters
 - local-mac-address? variable
 - changes during upgrade, 194
 - requirements, 24
- evacuating, *See* moving
- example configuration
 - cluster configuration used in, 231
 - data replication, 224, 230
 - groups and resources used in, 232
- extension properties
 - application resource, 239
 - replication resource, 237

F

- failover
 - affinity switchover, 226

- failover (Continued)
 - coping with, 247
 - definition, 222
 - guidelines for managing, 229
 - resource groups, 226
- failover applications, guidelines for configuring
 - resource groups, 227-228
- failover resource groups, role in failover, 226
- fattach command, restriction on cluster file systems, 31
- file systems for NFS application, configuring for
 - data replication, 235
- file-system logging, planning, 38-39
- forcedirectio command, restriction, 33

G

- global devices
 - caution notice, 178
 - /global/.devices/ directory
 - mirroring, 140-143
 - /global/.devices directory
 - node@nodeid file system, 34
 - /globaldevices partition
 - creating, 53
 - planning, 18
 - planning, 30-33
 - updating the namespace, 152
 - /global directory, 32
- global file systems, *See* cluster file systems

H

- help, 12
- high-priority processes, restriction for cluster nodes, 25
- hosts file
 - configuring, 55
 - configuring for JumpStart, 83
 - planning, 22
- hot spare disks, planning, 35

I

- importing disk device groups, 173

- initialization files, 99-100
- installation mode, 106
- installing
 - See also* adding
 - Apache packages
 - by using `pkgadd`, 67
 - Cluster Control Panel (CCP), 48-52
 - data services
 - by using `installer`, 100-102
 - by using `scinstall`, 102-104
 - by using SunPlex Installer, 70
 - RSMAPI
 - Solaris packages, 55
 - Sun Cluster package, 60
 - Solaris
 - alone, 52-55
 - with Sun Cluster, 76-87
 - Solstice DiskSuite, 129-150
 - by using SunPlex Installer, 70
 - from the Solaris CD-ROM, 131-132
 - Sun Cluster
 - all nodes, 57-63
 - by using JumpStart, 76-87
 - by using SunPlex Installer, 70
 - module to Sun Management Center, 119-120
 - single-node clusters, 87-90
 - status, 75
 - verifying, 107
 - Sun Management Center
 - requirements, 118-119
 - Sun Cluster module, 119-120
 - Sun StorEdge QFS, 62, 75, 87
 - Sun Web Console, 56
 - SunPlex Installer, 65
 - VERITAS File System (VxFS), 97
 - VERITAS Volume Manager (VxVM), 163-179
 - and encapsulating the root disk, 165-168
 - without encapsulating the root disk, 170-172
- IP addresses, planning, 22-23
- IP Network Multipathing groups
 - configuring, 113-114, 114
 - planning, 29-30
 - test-IP-address requirements
 - planning, 29
 - upgrade, 185
 - upgrading from NAFO groups, 182

- IP Network Multipathing groups (Continued)
 - upgrading NAFO groups, 194
- IPMP, *See* IP Network Multipathing groups
- IPv6 addresses, use on the public network, 23

J

- Java Dynamic Management Kit (JDMK),
 - package installation, 67
- JumpStart
 - `/etc/inet/hosts` file, 83
 - installing Solaris and Sun Cluster, 76-87
- junctions, *See* transport junctions

K

- Kerberos with NFS, restriction, 24
- `kernel_cage_enable` variable, 55
- `/kernel/drv/md.conf` file, 35
 - caution notice, 36, 134
 - configuring, 133-134

L

- licenses, planning, 21-22
- loading the Sun Cluster module to Sun Management Center, 121-123
- `local-mac-address?` variable
 - changes during upgrade, 194
 - requirements, 24
- locales, 100, 101
- `localonly` property, enabling, 169
- LOFS, restriction on cluster file systems, 31
- log files
 - package installation, 101
 - Sun Cluster installation, 62
 - SunPlex Installer installation, 75
- logging for cluster file systems, planning, 38-39
- logical addresses, planning, 23
- logical hostname resource, role in failover, 226
- logical network interfaces, restriction, 28
- loopback file system (LOFS), restriction on cluster file systems, 31

M

MANPATH

- administrative console, 51
- cluster nodes, 99

md.conf file

- caution notice, 134
- configuring, 133-134
- planning, 35

md_nsets field

- configuring, 133-134
- planning, 36

md.tab file, configuring, 156-157

mediators, *See* dual-string mediators

messages files

- See also* error messages
- cluster, 13
- SunPlex Installer, 75

metadevices

- activating, 157-159
- planning the maximum number, 36
- setting the maximum number, 133-134

minor-number conflicts, repairing, 175-176

mirroring

- differing device sizes, 39
- global namespace, 140-143
- multihost disks, 39
- planning, 39-41
- root disks, 136

- caution notice, 168

- planning, 39-41

mount options for cluster file systems

- requirements, 110
- VxFS, 32

mount points

- cluster file systems, 32-33
- modifying the `/etc/vfstab` file, 110
- nested, 33

moving, resource groups and device groups, 204

multihost disks

- mirroring, 39
- planning, 35

multiported disks, *See* multihost disks

N

NAFO groups

- See also* IP Network Multipathing groups
- upgrading to IP Network Multipathing groups, 194

name-service switch, configuring, 98-99

name_to_major file

- non-VxVM nodes, 55, 167, 171
- VxVM-installed nodes, 171

naming convention, replication resource groups, 226

Network File System (NFS)

- configuring application file systems for data replication, 235
- guidelines for cluster nodes, 24

Network Time Protocol (NTP),

- configuring, 116-117

NFS, *See* Network File System (NFS)

NIS servers, restriction for cluster nodes, 24

nmd field

- configuring, 133-134
- planning, 36

node lists

- disk device groups, 35
- raw-disk device groups
- removing nodes from, 169
- viewing, 169

nodes, *See* cluster nodes

noncluster mode

- rebooting into, 124
- rebooting into single-user, 190, 210

nonrolling upgrade, requirements, 183

nsswitch.conf file, configuring, 98-99

NTP, configuring, 116-117

ntp.conf.cluster file

- configuring, 116-117
- starting, 117

ntp.conf file

- changes during upgrade, 194
- configuring, 116-117
- starting, 117

O

online help, Sun Cluster module to Sun Management Center, 122

`/opt/SUNWcluster/bin/` directory, 51

- /opt/SUNWcluster/bin/cconsole
 - command, 51
 - installing the software, 48-52
 - using, 53, 83
- /opt/SUNWcluster/bin/ccp command, 51
- /opt/SUNWcluster/man/ directory, 51
- Oracle Parallel Server, *See* Oracle Real Application Clusters

P

- package installation
 - Apache, 67
 - Cluster Control Panel (CCP) software, 48-52
 - common agent container, 68
 - data services
 - by using installer, 100-102
 - by using scinstall, 102-104
 - Java Dynamic Management Kit (JDMK), 67
 - RSMRDT driver, 60
 - SCI-PCI, 60
 - Sun Cluster software, 56-57
 - Sun Web Console, 56
 - SunPlex Manager, 69
- partitions
 - /globaldevices, 18, 53
 - repartitioning drives, 155
 - root (/), 19
 - /sds, 53
 - swap, 18
 - volume manager, 18
- patches
 - default installation directory, 57
 - patch-list file, 58
 - planning, 22
- PATH
 - administrative console, 51
 - cluster nodes, 99
- PCI adapters, *See* SCI-PCI adapters
- point-in-time snapshot
 - definition, 223
 - performing, 244
- ports, *See* serial ports
- primary cluster, role in data replication, 222
- private hostnames
 - changing, 115
 - planning, 27

- private hostnames (Continued)
 - verifying, 115
- private network, planning, 26-27
- public network
 - IPv6 support, 23
 - planning, 23-24

Q

- QFS
 - See* Sun StorEdge QFS
- quorum devices
 - caution notice, 168
 - initial configuration, 104
 - and mirroring, 40
 - planning, 30
 - verifying, 107
- quotas, restriction on cluster file systems, 31

R

- RAID, restriction, 34
- rarpd service, restriction for cluster nodes, 25
- raw-disk device group node lists
 - removing nodes, 169
 - viewing, 169
- raw-disk device groups, *See* disk device groups
- rebooting
 - into noncluster mode, 124
 - into single-user noncluster mode, 190, 210
- registering, VxVM disk device groups, 174
- release file, 48
- remote mirror replication
 - definition, 222
 - performing, 243
- Remote Shared Memory Application Programming Interface (RSMAPI)
 - package requirements, 17
 - Solaris packages installation, 55
 - Sun Cluster packages
 - installing by using pkgadd, 60
- removing Sun Cluster software, 123-124
- repairing
 - mediator data, 161
 - minor-number conflicts, 175-176

- repairing (Continued)
 - storage reconfiguration during upgrade, 214-216
- replication, *See* data replication
- replication resource groups
 - creating, 237
 - guidelines for configuring, 226
 - naming convention, 226
- resource groups
 - application, 226
 - failover, 226
 - failover applications, 227-228
 - guidelines for configuring, 225
 - moving, 173, 204
 - replication, 226
 - scalable applications, 228-229
 - shared address, 228
 - taking offline, 186
 - verifying, 204
- resource types
 - reregistering after upgrade, 198, 212
- resources, disabling, 186
- rolling upgrade, requirements, 183
- root (/) file systems, mirroring, 136-140
- root disk groups
 - configuring
 - on encapsulated root disks, 165-168
 - on nonroot disks, 172-173
 - planning, 37
 - simple, 37
- root disks
 - encapsulating, 165-168
 - mirroring, 136
 - caution notice, 168
 - planning, 39-41
 - unencapsulating, 177-179
- root environment, configuring, 99-100
- rootdg, *See* root disk groups
- routers, restriction for cluster nodes, 24
- RPC service, restricted program numbers, 25
- rpcmod settings, 97
- RSMAPI, *See* Remote Shared Memory Application Programming Interface (RSMAPI)
- RSMRDT driver
 - package installation, 60
 - requirements, 71

S

- SBus SCI adapters, restriction, 28
- scalable applications, guidelines for configuring
 - resource groups, 228-229
- sccheck command, vfstab file check, 112
- scconf command
 - adding nodes to the authorized-node list, 123
 - enabling the localonly property, 137
 - error messages, 174
 - removing nodes from a node list
 - authorized-node list, 95
 - raw-disk device groups, 138, 169
 - verifying installation mode, 107
 - viewing private hostnames, 115
- scdidadm command
 - determining device-ID names, 105
 - displaying device-ID names, 144
 - error messages, 197
 - migrating device IDs after upgrade, 197, 215
 - verifying device-ID migration, 196
- scgdevs command
 - error messages, 132
 - updating the global-devices namespace, 152
 - verifying command processing, 152
- SCI-PCI adapters
 - package requirements, 17
 - Solaris packages installation, 55
 - Sun Cluster packages installation, 60
- scinstall command
 - installing and configuring Sun Cluster
 - additional nodes, 90-97
 - all nodes, 57-63
 - by using JumpStart, 76-87
 - single-node clusters, 87-90
 - installing Sun Cluster
 - data services, 102-104
 - uninstalling Sun Cluster, 123-124
 - upgrading data services, 195
 - upgrading Sun Cluster
 - nonrolling, 193
 - rolling, 208
 - verifying Sun Cluster software, 196
- scsetup command
 - adding cluster interconnects, 92
 - changing private hostnames, 115
 - postinstallation setup, 106
 - registering disk device groups, 174

- scshutdown command, 188
- scstat command
 - verifying cluster mode, 196
 - verifying disk-group configurations, 176
- scswitch command
 - moving resource groups and device groups, 173, 204
 - taking resource groups offline, 186
- scversions command, rolling upgrade, 213
- scvxinstall command
 - installing VxVM and encapsulating the root disk, 165-168
 - installing VxVM only, 170-172
- /sds partition, 53
- secondary cluster, role in data replication, 222
- secondary root disks, 40
- Secure NFS, restriction, 24
- secure ports for NFS, 24
- security files, 68
- serial ports
 - configuring on the administrative console, 50
 - Simple Network Management Protocol (SNMP), 119
- serialports file, 50
- shared address resource groups, guidelines for configuring, 228
- shared file systems
 - See also* cluster file systems
 - required mount parameters for QFS, 111
- shutting down the cluster, 188
- Simple Network Management Protocol (SNMP), port for Sun Management Center, 119
- single-node clusters, 87-90
- single-user noncluster mode
 - rebooting into, 190, 210
- snapshot, point-in-time, 223
- SNMP, port for Sun Management Center, 119
- software RAID, restriction, 34
- Solaris
 - installing
 - alone, 52-55
 - with Sun Cluster, 76-87
 - planning, 16-21
 - /globaldevices file system, 19-20
 - partitions, 18-21
 - root (/) file system, 19
 - software groups, 17-18
- Solaris, planning (Continued)
 - volume managers, 20
 - restrictions
 - automatic power-saving shutdown, 17
 - interface groups, 17
 - upgrading
 - nonrolling, 189-192
 - rolling, 206-207
 - verifying device-ID migration, 196
 - version, 48
- Solaris Volume Manager
 - coexistence with VxVM, 167, 171
 - configuring, 129-150
 - disk sets
 - adding drives, 153-155
 - configuring, 151-153
 - repartitioning drives, 155
 - setting the maximum number, 133-134
 - dual-string mediators
 - adding hosts, 160
 - overview, 159-161
 - repairing bad data, 161
 - status, 160-161
 - error messages, 136
 - md.tab file, 156-157
 - mediators
 - See* dual-string mediators
 - mirroring
 - global namespace, 140-143
 - root (/) file system, 136-140
 - root disks, 136
 - planning, 35-36
 - state database replicas, 135-136
 - transactional-volume logging
 - planning, 38
 - volumes
 - activating, 157-159
 - planning the maximum number, 36
 - setting the maximum number, 133-134
- Solstice DiskSuite
 - coexistence with VxVM, 167, 171
 - configuring, 129-150
 - disk sets
 - adding drives, 153-155
 - configuring, 151-153
 - repartitioning drives, 155
 - setting the maximum number, 133-134

- Solstice DiskSuite (Continued)
 - dual-string mediators
 - adding hosts, 160
 - overview, 159-161
 - repairing bad data, 161
 - status, 160-161
 - error messages, 136
 - installing, 129-150
 - from the Solaris CD-ROM, 131-132
 - installing and configuring
 - by using SunPlex Installer, 70
 - md.tab file, 156-157
 - mediators
 - See* dual-string mediators
 - metadevices
 - activating, 157-159
 - planning the maximum number, 36
 - setting the maximum number, 133-134
 - mirroring
 - root (/) file system, 136-140
 - root disks, 136
 - planning, 35-36
 - state database replicas, 135-136
 - trans-metadevice logging
 - planning, 38
- SSP, *See* console-access devices
- stack-size setting, 97, 175
- starting
 - Cluster Control Panel (CCP), 51
 - Sun Management Center, 120
 - SunPlex Installer, 74
- state database replicas, configuring, 135-136
- status
 - disk device groups, 176
 - dual-string mediators, 160-161
 - Sun Cluster
 - installation logs, 75
 - verifying, 107
- Sun Cluster HA for SAP liveCache, upgrade
 - configuration change, 199
- Sun Cluster module to Sun Management Center, 118-123
 - adding nodes, 121
 - installing, 119-120
 - loading, 121-123
 - online help, 122
 - requirements, 118-119
- Sun Cluster module to Sun Management Center (Continued)
 - upgrade
 - nonrolling, 217-218
- Sun Enterprise 10000 servers
 - dynamic reconfiguration support, 55
 - kernel_cage_enable variable, 55
 - serialports file, 51
- Sun Explorer
 - upgrading, 194, 208
- Sun Fire 15000 servers
 - IP addresses, 22
 - serial-port numbers, 51
- Sun Management Center
 - starting, 120
 - Sun Cluster module, 118-123
 - adding nodes, 121
 - installing, 119-120
 - loading, 121-123
 - online help, 122
 - upgrading, 217-218
 - upgrading, 218-220
- Sun StorEdge Availability Suite, preparing for
 - cluster upgrade, 187
- Sun StorEdge QFS
 - installing, 62, 75, 87
 - mounting shared file systems, 111
- Sun Web Console, installing, 56
- SunPlex Installer
 - guidelines, 63-65
 - installing, 65
 - starting, 74
 - using to install software, 70
- SunPlex Manager, package installation, 69
- swap, planning, 18
- switchback, guidelines for performing, 230
- switchover
 - affinity switchover, 226
 - guidelines for managing, 229
 - performing, 247
- SyMON, *See* Sun Management Center
- synchronous data replication, 223
- system controllers (SC), *See* console-access devices
- system file
 - kernel_cage_enable variable, 55
 - secure ports for NFS, 24
 - stack-size setting, 97

- system file (Continued)
 - thread stack-size setting, 175
- System Service Processor (SSP), *See*
 - console-access devices

T

- technical support, 12
- telnet command, serial-port numbers, 51
- terminal concentrators (TC), *See* console-access devices
- test-IP-address requirements
 - new installations, 29
 - upgrades, 182, 185
- thread stack-size setting, 175
- three-way mirroring, 39
- transport adapters, *See* adapters
- transport junctions, planning, 28

U

- UFS logging, planning, 38
- unencapsulating the root disk, 177-179
- uninstalling Sun Cluster software, 123-124
- upgrading, 181-220
 - choosing an upgrade method, 182-183
 - guidelines for, 182
 - nonrolling, 184
 - data services, 195
 - preparing the cluster, 184-189
 - requirements, 183
 - resource types, 198
 - restoring mediators, 199
 - Solaris, 189-192
 - unconfiguring mediators, 186
 - recovering from storage changes, 214-216
 - rolling, 201-214
 - data services, 209
 - preparing the cluster, 202-206
 - requirements, 183
 - resource types, 212
 - restoring mediators, 213
 - Solaris, 206-207
 - unconfiguring mediators, 205
- Sun Cluster HA for SAP liveCache configuration, 199

upgrading (Continued)

- Sun Cluster module to Sun Management Center, 217-218
- Sun Explorer, 194, 208
- Sun Management Center, 218-220
- Sun StorEdge Availability Suite configuration device, 187
- verifying
 - cluster status, 211
 - device-ID conversion, 196
 - successful upgrade, 211
 - version, 196
- user-initialization files, modifying, 99-100
 - /usr/cluster/bin/ directory, 99
 - /usr/cluster/bin/sccheck command, vfstab file check, 112
 - /usr/cluster/bin/scconf command
 - adding nodes to the authorized-node list, 123
 - enabling the localonly property, 137
 - error messages, 174
 - removing nodes from a node list
 - authorized-node list, 95
 - raw-disk device groups, 138, 169
 - verifying installation mode, 107
 - viewing private hostnames, 115
 - /usr/cluster/bin/scdidadm command
 - determining device-ID names, 105
 - displaying device-ID names, 144
 - error messages, 197
 - migrating device IDs after upgrade, 197, 215
 - verifying device-ID migration, 196
 - /usr/cluster/bin/scgdevs command
 - error messages, 132
 - updating the global-devices namespace, 152
 - verifying command processing, 152
 - /usr/cluster/bin/scinstall command
 - installing and configuring Sun Cluster
 - additional nodes, 90-97
 - all nodes, 57-63
 - by using JumpStart, 76-87
 - single-node clusters, 87-90
 - installing Sun Cluster
 - data services, 102-104
 - uninstalling Sun Cluster, 123-124
 - verifying Sun Cluster software, 196
 - /usr/cluster/bin/scsetup command
 - adding cluster interconnects, 92

- /usr/cluster/bin/scsetup command
- (Continued)
 - changing private hostnames, 115
 - postinstallation setup, 106
 - registering disk device groups, 174
- /usr/cluster/bin/scshutdown command, 188
- /usr/cluster/bin/scstat command
 - verifying cluster mode, 196
 - verifying disk-group configurations, 176
- /usr/cluster/bin/scswitch command
 - moving resource groups and device groups, 173, 204
 - taking resource groups offline, 186
- /usr/cluster/bin/scversions command,
 - rolling upgrade, 213
- /usr/cluster/bin/scvxinstall command
 - installing VxVM and encapsulating the root disk, 165-168
 - installing VxVM only, 170-172
- /usr/cluster/man/ directory, 99

V

- /var/sadm/install/logs/ directory, 101
- /var/adm/messages file, 13
- /var/cluster/spm/messages file, 75
- verifying, 115
 - cluster status, 211
 - data replication configuration, 245
 - device group configurations, 204
 - device-ID migration, 196
 - installation mode, 107
 - quorum configurations, 107
 - resource group configurations, 204
 - scgdevs command processing, 152
 - upgrade, 196, 211
 - vfstab configuration, 112
 - VxVM disk-group configurations, 176
- VERITAS File System (VxFS)
 - administering, 113
 - installing, 97
 - mounting cluster file systems, 32, 112
 - planning, 32, 38
- VERITAS Volume Manager (VxVM)
 - cluster feature, 174

- VERITAS Volume Manager (VxVM)
 - (Continued)
 - configuring, 163-179
 - disk groups, 173-175
 - non-VxVM nodes, 167, 171
 - volumes, 173-175
 - disk device groups
 - importing and deporting, 173
 - reminoring, 175-176
 - disk-group registration, 174
 - encapsulation, 37
 - Enclosure-Based Naming, 37
 - installing, 163-179
 - and encapsulating the root disk, 165-168
 - VxVM only, 170-172
 - mirroring the encapsulated root disk, 168-170
 - planning, 20, 37-38
 - removing man pages, 167, 171
 - root disk groups
 - configuring on nonroot disks, 172-173
 - planning, 37, 164-165
 - simple, 37
 - root disks
 - caution when unencapsulating, 178
 - unencapsulating, 177-179
 - verifying disk-group configurations, 176
- vfstab file
 - adding mount points, 110
 - modifying during upgrade
 - nonrolling, 190
 - rolling, 206
 - verifying the configuration, 112
- vold daemon, 67
- volume managers
 - See also* VERITAS Volume Manager (VxVM)
 - partitions for, 18
 - planning
 - general, 33-41
 - Solaris Volume Manager, 35-36
 - Solstice DiskSuite, 35-36
 - VERITAS Volume Manager, 37-38
- volumes
 - Solaris Volume Manager
 - activating, 157-159
 - planning the maximum number, 36
 - setting the maximum number, 133-134

volumes (Continued)

VxVM

 configuring, 173-175

 verifying, 176

VxFS, *See* VERITAS File System (VxFS)

vxio driver major number

 non-VxVM nodes, 167, 171

 VxVM-installed nodes, 167, 171

VxVM, *See* VERITAS Volume Manager (VxVM)

X

xntpd.cluster start command, 117

xntpd start command, 117

