



Sun Cluster 3.1 Software Installation Guide

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

Part No: 816-3388-10
May 2003, Revision A

Copyright 2003 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, JumpStart, Sun Fire, OpenBoot, SunPlex, and Solaris are trademarks, registered trademarks, or service marks 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(TM) is a registered trademark of Oracle Corporation. Netscape(TM) is a trademark or registered trademark of Netscape Communications Corporation in the United States and other countries. The Adobe PostScript(TM) 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.

Federal Acquisitions: Commercial Software—Government Users Subject to Standard License Terms and Conditions.

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 2003 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.

Des 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, JumpStart, Sun Fire, OpenBoot, SunPlex, et Solaris sont des marques de fabrique ou des marques déposées, ou marques de service, 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(TM) est une marque déposée registre de Oracle Corporation. Netscape(TM) est une marque de Netscape Communications Corporation aux Etats-Unis et dans d'autres pays. Le logo Adobe PostScript(TM) 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 REPOUDRE 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.



030519@5943



Contents

Preface	7
1 Planning the Sun Cluster Configuration	11
Where to Find Sun Cluster Installation Tasks	11
Planning the Solaris Operating Environment	13
Guidelines for Selecting Your Solaris Installation Method	13
Solaris Software Group Considerations	13
System Disk Partitions	14
Planning the Sun Cluster Environment	17
Licensing	17
Software Patches	17
IP Addresses	17
Sun Cluster Configurable Components	18
Planning the Global Devices and Cluster File Systems	23
Guidelines for Highly Available Global Devices and Cluster File Systems	23
Mount Information for Cluster File Systems	23
Planning Volume Management	24
Guidelines for Volume Manager Software	25
Guidelines for Solstice DiskSuite/Solaris Volume Manager Software	26
Guidelines for VERITAS Volume Manager Software	27
File-System Logging	28
Mirroring Guidelines	29
2 Installing and Configuring Sun Cluster Software	33
Installing the Software	34

▼ How to Prepare for Cluster Software Installation	35
▼ How to Install Cluster Control Panel Software on the Administrative Console	37
▼ How to Install Solaris Software	39
▼ How to Install Sun Cluster Software on the First Cluster Node (scinstall)	44
▼ How to Install Sun Cluster Software on Additional Cluster Nodes (scinstall)	53
Using SunPlex Manager to Install Sun Cluster Software	62
▼ How to Install Solaris and Sun Cluster Software (JumpStart)	74
▼ How to Configure the Name Service Switch	87
▼ How to Set Up the Root Environment	88
▼ How to Install Data Service Software Packages	89
▼ How to Perform Post-Installation Setup	90
▼ How to Uninstall Sun Cluster Software To Correct Installation Problems	92
Configuring the Cluster	93
▼ How to Add Cluster File Systems	94
▼ How to Configure Internet Protocol (IP) Network Multipathing Groups	98
▼ How to Change Private Hostnames	99
▼ How to Configure Network Time Protocol (NTP)	100
Installing the Sun Cluster Module for Sun Management Center	102
Installation Requirements for Sun Cluster Monitoring	103
▼ How to Install the Sun Cluster Module for Sun Management Center	103
▼ How to Start Sun Management Center	104
▼ How to Add a Cluster Node as a Sun Management Center Agent Host Object	105
▼ How to Load the Sun Cluster Module	106
3 Upgrading Sun Cluster Software	109
Upgrading From Sun Cluster 3.0 to Sun Cluster 3.1 Software	109
Overview of Upgrading From Sun Cluster 3.0 to Sun Cluster 3.1 Software	110
▼ How to Prepare the Cluster for Upgrade	111
▼ How to Upgrade the Solaris Operating Environment	113
▼ How to Upgrade Cluster Software	116
▼ How to Finish Upgrading Cluster Software	119
Sun Management Center Software Upgrade	122
▼ How to Upgrade Sun Management Center Software	122

A	Installing and Configuring Solstice DiskSuite/Solaris Volume Manager Software	125
	Installing and Configuring Solstice DiskSuite/Solaris Volume Manager Software	126
	Solstice DiskSuite/Solaris Volume Manager Configuration Example	127
	▼ How to Install Solstice DiskSuite Software	129
	▼ How to Set the Number of Metadevice or Volume Names and Disksets	130
	▼ How to Create State Database Replicas	132
	Mirroring the Root Disk	133
	▼ How to Mirror the Root (/) File System	133
	▼ How to Mirror the Global Namespace	137
	▼ How to Mirror File Systems Other than Root (/) That Cannot Be Unmounted	140
	▼ How to Mirror File Systems That Can Be Unmounted	144
	▼ How to Create a Diskset	147
	Adding Drives to a Diskset	149
	▼ How to Repartition Drives in a Diskset	151
	▼ How to Create an md.tab File	152
	▼ How to Activate Metadevices or Volumes	154
	Mediators Overview	155
	▼ How to Check the Status of Mediator Data	156
	▼ How to Fix Bad Mediator Data	157
B	Installing and Configuring VERITAS Volume Manager	159
	Installing and Configuring VxVM Software	159
	Setting Up a rootdg Disk Group Overview	160
	▼ How to Install VERITAS Volume Manager Software and Encapsulate the Root Disk	161
	▼ How to Mirror the Encapsulated Root Disk	165
	▼ How to Install VERITAS Volume Manager Software Only	167
	▼ How to Create a rootdg Disk Group on a Non-Root Disk	170
	▼ How to Create and Register a Disk Group	171
	▼ How to Assign a New Minor Number to a Disk Device Group	173
	▼ How to Verify the Disk Group Configuration	174
	▼ How to Unencapsulate the Root Disk	174
	Index	177

Preface

The *Sun Cluster 3.1 Software Installation Guide* contains guidelines for planning a Sun™ Cluster 3.1 configuration, and provides procedures for installing, upgrading, and configuring the Sun Cluster software.

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 environment and expertise with the volume manager software used with Sun Cluster software.

Using UNIX Commands

This document contains information on commands 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.

- AnswerBook2™ online documentation for the Solaris software environment
- Other software documentation that you received with your system
- Solaris operating environment man pages

Typographic Conventions

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

TABLE P-1 Typographic Conventions

Typeface or Symbol	Meaning	Example
AaBbCc123	The names of commands, files, and directories; on-screen 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 on-screen computer output	<code>machine_name%</code> su Password:
<i>AaBbCc123</i>	Command-line placeholder: replace with a real name or value	To delete a file, type rm <i>filename</i> .
<i>AaBbCc123</i>	Book titles, new words, or terms, or words to be emphasized.	Read Chapter 6 in <i>User's Guide</i> . These are called <i>class</i> options. You must be <i>root</i> to do this.

Shell Prompts in Command Examples

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

TABLE P-2 Shell Prompts

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

Related Documentation

Application	Title	Part Number
Concepts	<i>Sun Cluster 3.1 Concepts Guide</i>	816-3383
Hardware	<i>Sun Cluster 3.x Hardware Administration Manual</i>	817-0168
	Sun Cluster 3.x Hardware Administration Collection at http://docs.sun.com/db/coll/1024.1/	
Data Services	<i>Sun Cluster 3.1 Data Service Planning and Administration Guide</i>	817-1526
	Sun Cluster 3.1 Data Service Collection at http://docs.sun.com/db/coll/573.10/	
API Development	<i>Sun Cluster 3.1 Data Services Developer's Guide</i>	816-3385
Administration	<i>Sun Cluster 3.1 System Administration Guide</i>	816-3384
Error Messages	<i>Sun Cluster 3.1 Error Messages Guide</i>	816-3382
Man Pages	<i>Sun Cluster 3.1 Man Page Reference Manual</i>	816-5251
Release Notes	<i>Sun Cluster 3.1 Release Notes</i>	816-5317
	<i>Sun Cluster 3.1 Release Notes Supplement</i>	816-3381

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>.

Getting Help

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

- Your name and email address (if available)

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

Use the following commands to gather information 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
<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 11
- “Planning the Solaris Operating Environment” on page 13
- “Planning the Sun Cluster Environment” on page 17
- “Planning the Global Devices and Cluster File Systems” on page 23
- “Planning Volume Management” on page 24

Where to Find Sun Cluster Installation Tasks

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

TABLE 1-1 Location of Sun Cluster Software Installation Task Information

Task	For Instructions, Go To ...
Set up cluster hardware.	<i>Sun Cluster 3.x Hardware Administration Manual</i> Documentation shipped with your server and storage devices

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

Task	For Instructions, Go To ...
Plan cluster software installation.	This chapter "Sun Cluster Installation and Configuration Worksheets" in <i>Sun Cluster 3.1 Release Notes</i>
Install a new cluster or add nodes to an existing cluster: Install the Solaris operating environment, Cluster Control Panel (optional), SunPlex Manager (optional), cluster framework, and data service software packages.	"Installing the Software" on page 34
Install and configure Solstice DiskSuite/Solaris Volume Manager software.	"Installing and Configuring Solstice DiskSuite/Solaris Volume Manager Software" on page 126 Solstice DiskSuite/Solaris Volume Manager documentation
Install and configure VERITAS Volume Manager (VxVM) software.	"Installing and Configuring VxVM Software" on page 159 VxVM documentation
Configure cluster framework software and optionally install and configure Sun Management Center.	"Configuring the Cluster" on page 93
Plan, install, and configure resource groups and data services.	<i>Sun Cluster 3.1 Data Service Planning and Administration Guide</i> "Sun Cluster Installation and Configuration Worksheets" in <i>Sun Cluster 3.1 Data Service Release Notes</i>
Develop custom data services.	<i>Sun Cluster 3.1 Data Services Developer's Guide</i>
Upgrade from Sun Cluster 3.0 to Sun Cluster 3.1 software (Solaris operating environment, cluster framework, data services, and volume manager software).	"Upgrading From Sun Cluster 3.0 to Sun Cluster 3.1 Software" on page 109 "Installing and Configuring Solstice DiskSuite/Solaris Volume Manager Software" on page 126 or "Installing and Configuring VxVM Software" on page 159 Volume manager documentation

Planning the Solaris Operating Environment

This section provides guidelines for planning Solaris software installation in a cluster configuration. For more information about Solaris software, see the 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 operating environment and Sun Cluster software by using JumpStart. If you are installing several cluster nodes, consider a network installation.

See “How to Install Solaris and Sun Cluster Software (JumpStart)” on page 74 for details about the `scinstall` JumpStart installation method. See the Solaris installation documentation for details about standard Solaris installation methods.

Solaris Software Group Considerations

Sun Cluster 3.1 software requires at least the Solaris End User System Support 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 will install.

- See your server documentation for any Solaris software requirements. For example, Sun Enterprise 10000 servers require the Entire Distribution + OEM software group.
- If you install the Solaris 8 10/01 operating environment and intend to use SCI-PCI adapters or the Remote Shared Memory Application Programming Interface (RSMAPI), ensure that you install the RSMAPI software packages (`SUNWrsm`, `SUNWrsmx`, `SUNWrsmo`, and `SUNWrsmox`). The Solaris Developer System Support software group or higher includes these packages. If you install the End User System Support software group, use the `pkgadd(1M)` command to install these RSMAPI packages before you install Sun Cluster software. See the Solaris 8 10/01 section (3RSM) man pages for information about using the RSMAPI.
- You might need to install other Solaris software packages which are not part of the End User System Support software group, for example, the Apache HTTP server packages. Third-party software, such as ORACLE®, might also require additional Solaris packages. See your third-party documentation for any Solaris software

requirements.

System Disk Partitions

Add this information to “Local File Systems With Mirrored Root Worksheet” in *Sun Cluster 3.1 Release Notes* or “Local File Systems with Non-Mirrored Root Worksheet” in *Sun Cluster 3.1 Release Notes*.

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

- **swap** – The amount of swap space allocated for Solaris and Sun Cluster software combined must be no less than 750 Mbytes. For best results, add at least 512 Mbytes for Sun Cluster software to the amount required by the Solaris operating environment. In addition, allocate any additional amount required by applications that will run on the cluster node.
- **/globaldevices** – Create a 512-Mbyte file system that will be used by the `scinstall(1M)` utility for global devices.
- **Volume manager** – Create a 20-Mbyte partition for volume manager use on a slice at the end of the disk (slice 7). If your cluster uses VERITAS Volume Manager (VxVM) and you intend to encapsulate the root disk, you need 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 operating environment.

See the following guidelines for additional partition planning information.

Guidelines for the Root (/) File System

As with any other system running the Solaris operating environment, 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/Solaris Volume Manager software requires less than 5 Mbytes, and VxVM software requires less than 15 Mbytes. For best results, you need to configure ample additional space and inode capacity for the creation of both block special devices and character special devices used by either Solstice DiskSuite/Solaris Volume Manager or VxVM software, especially if a large number of shared disks are in the cluster. Therefore, add at least 100 Mbytes to the amount of space you would normally allocate for your root (/) file system.

- `/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/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/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. 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 swap Partition

The amount of swap space allocated for Solaris and Sun Cluster software combined must be no less than 750 Mbytes. For best results, add at least 512 Mbytes for Sun Cluster software to the amount required by the Solaris operating environment. In addition, allocate additional swap space for any third-party applications you install on the node that also have swap requirements. See your third-party application documentation for any swap requirements.

Guidelines for the `/globaldevices` File System

Sun Cluster software requires that you set aside a special file system on one of the local disks for use in managing global devices. This file system must be separate, as it will later be mounted as a cluster file system. Name this file system `/globaldevices`, which is the default name recognized by the `scinstall(1M)` command. The `scinstall` command later renames the file system `/global/.devices/node@nodeid`, where `nodeid` represents the number assigned to a node when it becomes a cluster member, and the original `/globaldevices` mount point is removed. The `/globaldevices` file system must have ample space and inode capacity for creating both block special devices and character special devices, especially if a large number of disks are in the cluster. A file system size of 512 Mbytes should be more than enough for most cluster configurations.

Volume Manager Requirements

If you use Solstice DiskSuite/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/Solaris Volume Manager software to function properly. See the Solstice DiskSuite/Solaris Volume Manager documentation for more information.

If you use VxVM and you intend to encapsulate the root disk, you need two unused slices available for use by VxVM, as well as some additional unassigned free space at either the beginning or the end of the disk. See the 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 will be installed with the Solaris operating environment End User System Support 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/Solaris Volume Manager software or VxVM. If you use Solstice DiskSuite/Solaris Volume Manager software, you use slice 7 for the state database replica. If you use VxVM, you later free slice 7 by assigning it a zero length. This layout provides the necessary two free slices, 4 and 7, and it 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 Solaris operating environment software, 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.
1	swap	1GB	512 Mbytes for Solaris operating environment software. 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 it as a cluster file system.
4	unused	-	Available as a free slice for encapsulating the root disk under VxVM.

TABLE 1-2 Example File System Allocation (Continued)

Slice	Contents	Allocation (in Mbytes)	Description
5	unused	-	-
6	unused	-	-
7	volume manager	20MB	Used by Solstice DiskSuite/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 for Sun Cluster software installation. For detailed information about Sun Cluster components, see the *Sun Cluster 3.1 Concepts Guide*.

Licensing

Ensure that you have any necessary license certificates available 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 3.1 Release Notes* or consult your Sun service provider. See "Patching Sun Cluster Software and Firmware" in *Sun Cluster 3.1 System Administration Guide* for general guidelines and procedures for applying patches.

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 to them. Add these IP addresses to any naming services used. Also add these IP addresses to the local `/etc/inet/hosts` file on each cluster node after you install Solaris software.

TABLE 1-3 Sun Cluster Components That Use IP Addresses

Component	IP Addresses Needed
Administrative console	1 per subnet
Network adapters	2 per adapter (1 primary IP address and 1 test IP address)
Cluster nodes	1 per node, per subnet
Domain console network interface (Sun Fire 15000)	1 per domain
Console-access device	1
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 used to communicate with the cluster nodes. A terminal concentrator can be 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. For more information about console access, see the *Sun Cluster 3.1 Concepts Guide*.

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. See the *Sun Cluster 3.1 Data Service Planning and Administration Guide* for information and “Sun Cluster Installation and Configuration Worksheets” in *Sun Cluster 3.1 Data Service Release Notes* for worksheets for planning resource groups. For more information about data services and resources, also see the *Sun Cluster 3.1 Concepts Guide*.

Sun Cluster Configurable Components

This section provides guidelines for the Sun Cluster components that you configure during installation.

Cluster Name

Add this planning information to the “Cluster and Node Names Worksheet” in *Sun Cluster 3.1 Release Notes*.

You specify a name for the cluster during Sun Cluster installation. The cluster name should be unique throughout the enterprise.

Node Names

Add this planning information to the “Cluster and Node Names Worksheet” in *Sun Cluster 3.1 Release Notes*. Information for most other worksheets is grouped by node name.

The node name is the name you assign to a machine when you install the Solaris operating environment. During Sun Cluster installation, you specify the names of all nodes that you are installing as a cluster.

Private Network

Add this planning information to the “Cluster and Node Names Worksheet” in *Sun Cluster 3.1 Release Notes*.

Sun Cluster software uses the private network for internal communication between nodes. Sun Cluster requires at least two connections to the cluster interconnect on the private network. You specify the private network address and netmask when you install 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 enterprise.

Note – After you have successfully installed the node as a cluster member, you cannot change the private network address and netmask.

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

- Use zeroes for the last two octets of the address.
- Follow the guidelines in RFC 1597 for network address assignments.

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 instructions on how to contact the InterNIC to obtain copies of RFCs.

If you specify a netmask other than the default, it must meet the following requirements:

- Minimally mask all bits given in the private network address
- Have no “holes”

Private Hostnames

Add this planning information to the “Cluster and Node Names Worksheet” in *Sun Cluster 3.1 Release Notes*.

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

Cluster Interconnect

Add this planning information to the “Cluster Interconnect Worksheet” in *Sun Cluster 3.1 Release Notes*.

The cluster interconnects provide the hardware pathways for private network communication between cluster nodes. Each interconnect consists of a cable connected between two transport adapters, a transport adapter and a transport junction, or two transport junctions. During Sun Cluster installation, you specify the following configuration 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. If you specify a transport junction, it will be easier to add another node to the cluster in the future.
- **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 automatically assigned during installation, or create other names.

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.

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 installation by using the `scsetup(1M)` utility.

For more information about the cluster interconnect, see the *Sun Cluster 3.1 Concepts Guide*.

Public Networks

Add this planning information to the “Public Networks Worksheet” in *Sun Cluster 3.1 Release Notes*.

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.
- The `local-mac-address?` variable must use the default value `true` for Ethernet adapters. Sun Cluster 3.1 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`.

See also “IP Network Multipathing Groups” on page 22 for guidelines on planning public network adapter backup groups. For more information about public network interfaces, see the *Sun Cluster 3.1 Concepts Guide*.

Disk Device Groups

Add this planning information to the “Disk Device Groups Worksheet” in *Sun Cluster 3.1 Release Notes*.

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 multiported disks and properly-configured volume manager devices as failover devices. Proper configuration of a volume manager device includes multiported disks and correct setup of the volume manager itself so that multiple nodes can host the exported device. You cannot configure tape

drives, CD-ROMs, or single-ported disks as failover devices.

- **Mirroring** – You must mirror the disks to protect the data from disk failure. See “Mirroring Guidelines” on page 29 for additional guidelines. See “Installing and Configuring Solstice DiskSuite/Solaris Volume Manager Software” on page 126 or “Installing and Configuring VxVM Software” on page 159 and your volume manager documentation for instructions on mirroring.

For more information about disk device groups, see the *Sun Cluster 3.1 Concepts Guide*.

IP Network Multipathing Groups

Add this planning information to the “Public Networks Worksheet” in *Sun Cluster 3.1 Release Notes*.

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. If a multipathing group is configured with two or more adapters and an 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.
- The `local-mac-address?` variable must have a value of `true` for Ethernet adapters. This is a change from the requirement for Sun Cluster 3.0 software.
- You must configure a test IP address for each multipathing group adapter.
- 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 they are not highly available.
- There are no requirements or restrictions for the name of a multipathing group.

For more information about IP Network Multipathing, see “Deploying Network Multipathing” in *IP Network Multipathing Administration Guide* or “Administering Network Multipathing (Task)” in *System Administration Guide: IP Services*.

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 assign quorum devices by using the `scsetup(1M)` utility.

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 so that the quorum devices have completely independent failure pathways.
- **Connection** – Do not connect a quorum device to more than two nodes.

For more information about quorum, see the *Sun Cluster 3.1 Concepts Guide*.

Planning the Global Devices and Cluster File Systems

This section provides guidelines for planning global devices and cluster file systems. For more information about global devices and cluster files systems, see the *Sun Cluster 3.1 Concepts Guide*.

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 global device and cluster file system layout.

- **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 they 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.

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 in the `/global` directory, unless prohibited by other software products. Using the `/global` directory enables you to easily distinguish cluster file systems, which are globally available, from local file systems.
- **VxFS mount requirement** – Globally mount and unmount a VxFS file system from the primary node (the node that masters the disk on which the VxFS file system resides) to ensure that the operation succeeds. A VxFS file system mount or unmount operation that is performed from a secondary node might fail.
- **Nesting mount points** – Normally, you should not nest the mount points for cluster file systems. For example, do not set up one file system mounted on `/global/a` and another file system mounted on `/global/a/b`. Ignoring this rule can cause availability and node boot order problems 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 (for example, different slices on the same disk).

Planning Volume Management

Add this planning information to the “Disk Device Groups Worksheet” in *Sun Cluster 3.1 Release Notes* and the “Volume Manager Configurations Worksheet” in *Sun Cluster 3.1 Release Notes*. For Solstice DiskSuite/Solaris Volume Manager, also add this planning information to the “Metadevices Worksheet (Solstice DiskSuite/Solaris Volume Manager)” in *Sun Cluster 3.1 Release Notes*.

This section provides guidelines for planning volume management of your cluster configuration.

Sun Cluster uses volume manager software to group disks into disk device groups that can then be administered as one unit. Sun Cluster supports Solstice DiskSuite/Solaris Volume Manager software and VERITAS Volume Manager (VxVM).

- If you use Solstice DiskSuite/Solaris Volume Manager software, you must install it on all nodes of the cluster, regardless of whether you use VxVM on some nodes to manage disks.
- If you use VxVM and enable the VxVM cluster feature, you must install and license VxVM on all nodes of the cluster.
- If you use VxVM and do *not* enable the VxVM cluster feature, you only need to install and license VxVM on those nodes which are attached to storage devices which VxVM will manage.
- If you install both Solstice DiskSuite/Solaris Volume Manager software and VxVM on a node, you must use Solstice DiskSuite/Solaris Volume Manager software to manage disks local to each node (such as the root disk) and you must use VxVM to manage all shared disks.

See your volume manager documentation and “Installing and Configuring Solstice DiskSuite/Solaris Volume Manager Software” on page 126 or “Installing and Configuring VxVM Software” on page 159 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 3.1 Concepts Guide*.

Guidelines for Volume Manager Software

Consider the following general guidelines when configuring your disks.

- **Mirrored multihost disks** – You must mirror all multihost disks across disk expansion units. See “Mirroring Multihost Disks” on page 29 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 disks.
- **Mirrored root** – Mirroring the root disk ensures high availability, but such mirroring is not required. See “Mirroring Guidelines” on page 29 for guidelines on deciding whether to mirror the root disk.
- **Unique naming** – On any cluster node, if a local Solstice DiskSuite metadvice or a local Solaris Volume Manager or VxVM volume is used as the device on which the `/global/.devices/node@nodeid` file system is mounted, the name of that metadvice or 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 3.1 Data Service Planning and Administration Guide* for information about node lists.
- **Multiported disks** – You must connect, or port, all disks used to construct a device group within the cluster to all of the nodes configured in the node list for that device group. Solstice DiskSuite/Solaris Volume Manager software is able to automatically check for this at the time that disks are added to a diskset. 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 they are not required.

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

Guidelines for Solstice DiskSuite/Solaris Volume Manager Software

Consider the following points when you plan Solstice DiskSuite/Solaris Volume Manager configurations.

- **Local metadvice or volume names** – The name of each local Solstice DiskSuite metadvice or Solaris Volume Manager volume must be unique throughout the cluster and cannot be the same as any device ID (DID) name.
- **Mediators** – Each diskset configured with exactly two disk strings and mastered by exactly two nodes must have Solstice DiskSuite/Solaris Volume Manager mediators configured for the diskset. A *disk string* consists of a disk enclosure, its physical disks, cables from the enclosure to the node(s), and the interface adapter cards. You must configure each diskset with exactly two nodes acting as mediator hosts. You must use the same two nodes for all disksets requiring mediators and those two nodes must master those disksets. Mediators cannot be configured for disksets 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 diskset are created in advance, at reconfiguration boot time, based on configuration parameters found in the `/kernel/drv/md.conf` file.

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 disksets that can be created for a system to meet the needs of the entire cluster. You must set the value of `md_nsets` to the expected number of disksets in the cluster, plus one to allow Solstice DiskSuite/Solaris Volume Manager software to manage the private disks on the local host (that is, those metadevices or volumes that are not in the local diskset).

The maximum number of disksets allowed per cluster is 32 (31 disksets for general use plus one for private disk management). The default value of `md_nsets` is 4.

- `nmd` – The `nmd` field defines the number of metadevices or volumes created for each diskset. You must set the value of `nmd` to the predicted highest value of metadvice or volume name used by any one of the disksets in the cluster. For example, if a cluster uses 10 metadevices or volumes in its first 15 disksets, but 1000 metadevices or volumes in the 16th diskset, you must set the value of `nmd` to at least 1000. Also, the value of `nmd` must be large enough to ensure that there are enough numbers for each DID name and for each local metadvice or volume name to be unique throughout the cluster.

The highest allowed value of a metadvice or volume name per diskset 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. Increasing the value of these fields after the cluster is in production is time consuming because it requires a reconfiguration reboot for each node. Raising 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 will use.



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

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.

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 (a feature introduced in VxVM version 3.2), 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. While failure to assign consistent names does not interfere with correct cluster behavior, it greatly complicates cluster administration and greatly increases the possibility of configuration errors, potentially leading to loss of data.
- **Root disk group** – You must create a default root disk group (`rootdg`) on each node. The `rootdg` disk group can be created on the following disks:
 - The root disk, which must be encapsulated
 - One or more local non-root disks, which can be encapsulated or initialized
 - A combination of root and local non-root disks

The `rootdg` disk group must be local to the node.

- **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 will 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** – Using Dirty Region Logging (DRL) is highly recommended but not required. Using DRL decreases volume recovery time after a node failure. Using DRL might decrease I/O throughput.

File-System Logging

Logging is required for cluster file systems. Sun Cluster software supports the following logging file systems:

- 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.
- VERITAS File System (VxFS) logging—See the `mount_vxfs` man page provided with VxFS software for more information.

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

TABLE 1-4 Supported File System Logging Matrix

Volume Manager	Supported File System Logging
Solstice DiskSuite/Solaris Volume Manager	Solaris UFS logging, Solstice DiskSuite trans-metadevice logging or Solaris Volume Manager transactional-volume logging, VxFS logging
VERITAS Volume Manager	Solaris UFS logging, VxFS logging

Consider the following points when you choose between Solaris UFS logging and trans-metadevice logging.

- Transactional volumes are scheduled to be removed from the Solaris operating environment in an upcoming Solaris release. Solaris UFS logging, available since the Solaris 8 release, 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 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, up to a maximum of 64 Mbytes.
- **Log metadvice** – 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 guidelines for planning the mirroring of your cluster configuration.

Mirroring Multihost Disks

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

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 disk expansion unit.
- **Disk space** – Mirroring doubles the amount of necessary disk space.
- **Three-way mirroring** – Solstice DiskSuite/Solaris Volume Manager software and VERITAS Volume Manager (VxVM) support three-way mirroring. However, Sun Cluster requires only two-way mirroring.
- **Number of metadvice or volumes** – Under Solstice DiskSuite/Solaris Volume Manager software, mirrors consist of other Solstice DiskSuite metadvice or Solaris Volume Manager volumes such as concatenations or stripes. Large configurations might contain a large number of metadvice or volumes.
- **Differing disk sizes** – If you mirror to a disk 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 the *Sun Cluster 3.1 Concepts Guide*.

Mirroring the Root Disk

Add this planning information to the “Local File Systems With Mirrored Root Worksheet” in *Sun Cluster 3.1 Release Notes*.

For maximum availability, you should 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 concerning 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/Solaris Volume Manager Software" on page 126 or "Installing and Configuring VxVM Software" on page 159 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 so that you can boot from the mirror if the primary boot disk fails.
- **Complexity** – Mirroring the root disk adds complexity to system administration and complicates booting in single-user mode.
- **Backups** – Regardless of whether or not 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 configured as a quorum device to mirror a root disk.
- **Quorum** – Under Solstice DiskSuite/Solaris Volume Manager software, in failure scenarios in which state database quorum is lost, you cannot reboot the system until maintenance is performed. See the Solstice DiskSuite/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. At a later point, the primary root disk might return to service (perhaps after a power cycle or transient I/O errors) and subsequent boots are performed by using the primary root disk specified in the OpenBoot™ PROM `boot-device` field. In this situation no manual repair task occurs, but the drive starts working well enough to boot. Note that a Solstice DiskSuite/Solaris Volume Manager 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 (causing a stale submirror). For example, changes to the `/etc/system` file would be lost. Some Solstice DiskSuite/Solaris Volume Manager 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 an underlying physical device, and the mirroring becomes active partway through the boot process (after the metadevices or volumes are loaded). Before this point, the system is 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.

The following procedures are in this chapter.

- “How to Prepare for Cluster Software Installation” on page 35
- “How to Install Cluster Control Panel Software on the Administrative Console” on page 37
- “How to Install Solaris Software” on page 39
- “How to Install Sun Cluster Software on the First Cluster Node (`scinstall`)” on page 44
- “How to Install Sun Cluster Software on Additional Cluster Nodes (`scinstall`)” on page 53
- “How to Install SunPlex Manager Software” on page 64
- “How to Add RBAC Authorization to an Existing User Account” on page 66
- “How to Create a New User Account” on page 67
- “How to Install Sun Cluster Software (SunPlex Manager)” on page 68
- “How to Install Solaris and Sun Cluster Software (JumpStart)” on page 74
- “How to Configure the Name Service Switch” on page 87
- “How to Set Up the Root Environment” on page 88
- “How to Install Data Service Software Packages” on page 89
- “How to Perform Post-Installation Setup” on page 90
- “How to Uninstall Sun Cluster Software To Correct Installation Problems” on page 92
- “How to Add Cluster File Systems” on page 94
- “How to Configure Internet Protocol (IP) Network Multipathing Groups” on page 98
- “How to Change Private Hostnames” on page 99
- “How to Configure Network Time Protocol (NTP)” on page 100
- “How to Install the Sun Cluster Module for Sun Management Center” on page 103
- “How to Start Sun Management Center” on page 104
- “How to Add a Cluster Node as a Sun Management Center Agent Host Object” on page 105
- “How to Load the Sun Cluster Module” on page 106

Installing the Software

The following table lists the tasks you perform to install the software.

TABLE 2-1 Task Map: Installing the Software

Task	For Instructions, Go To ...
Plan the layout of your cluster configuration and prepare to install software.	"How to Prepare for Cluster Software Installation" on page 35
(Optional) Install the Cluster Control Panel (CCP) software on the administrative console.	"How to Install Cluster Control Panel Software on the Administrative Console" on page 37
Install the Solaris operating environment and Sun Cluster software to establish new cluster nodes. Choose one of the following three methods:	
■ Method 1 – (New clusters or added nodes) Install Solaris software, then install the Sun Cluster software by using the <code>scinstall</code> utility.	"How to Install Solaris Software" on page 39 "How to Install Sun Cluster Software on the First Cluster Node (<code>scinstall</code>)" on page 44 "How to Install Sun Cluster Software on Additional Cluster Nodes (<code>scinstall</code>)" on page 53
■ Method 2 – (New clusters only) Install Solaris software, then install SunPlex™ Manager and use it to install the Sun Cluster software.	"How to Install Solaris Software" on page 39 "Using SunPlex Manager to Install Sun Cluster Software" on page 62
■ 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.	"How to Install Solaris and Sun Cluster Software (JumpStart)" on page 74
Configure the name service look-up order.	"How to Configure the Name Service Switch" on page 87
Set up directory paths.	"How to Set Up the Root Environment" on page 88
Install data service software packages.	"How to Install Data Service Software Packages" on page 89
Perform post-installation setup and assign quorum votes.	"How to Perform Post-Installation Setup" on page 90
Install and configure volume manager software:	

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

Task	For Instructions, Go To ...
■ Install and configure Solstice DiskSuite/Solaris Volume Manager software.	“Installing and Configuring Solstice DiskSuite/Solaris Volume Manager Software” on page 126 Solstice DiskSuite/Solaris Volume Manager documentation
■ Install and configure VERITAS Volume Manager software.	“Installing and Configuring VxVM Software” on page 159 VERITAS Volume Manager documentation
Configure the cluster.	“Configuring the Cluster” on page 93

▼ 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 will help you plan your cluster configuration and prepare your installation strategy.

- *Sun Cluster 3.1 Release Notes*—Restrictions, bug workarounds, and other late-breaking information.
- *Sun Cluster 3.1 Release Notes Supplement*—Post-release documentation about additional restrictions, bug workarounds, new features, and other late-breaking information. This document is regularly updated and is published online at the following Web site.
<http://docs.sun.com>
- *Sun Cluster 3.1 Concepts Guide*—Overview of the Sun Cluster product.
- *Sun Cluster 3.1 Software Installation Guide* (this manual)—Planning guidelines and procedures for installing and configuring Solaris, Sun Cluster, and volume manager software.
- *Sun Cluster 3.1 Data Service Planning and Administration Guide*—Planning guidelines and procedures for installing and configuring 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 product documentation you might need for reference during cluster installation.

- Solaris software
- Solstice DiskSuite/Solaris Volume Manager software
- VERITAS Volume Manager
- Sun Management Center
- Third-party applications such as ORACLE

3. Plan your cluster configuration.



Caution – Plan your cluster installation completely and 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 Parallel Fail Safe/Real Application Clusters Guard option of Oracle Parallel Server/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.

- Use the planning guidelines in Chapter 1 and in the *Sun Cluster 3.1 Data Service Planning and Administration Guide* to determine how you will install and configure your cluster.
- Fill out the cluster framework configuration worksheets in the *Sun Cluster 3.1 Release Notes* and the data services configuration worksheets in the *Sun Cluster 3.1 Data Service Release Notes*. 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.

5. Do you intend to use Cluster Control Panel software to connect from an administrative console to your cluster nodes?

- If yes, go to “How to Install Cluster Control Panel Software on the Administrative Console” on page 37.
- If no, go to one of the following procedures.
 - If you intend to install Sun Cluster software by using either SunPlex Manager (GUI-based method) or the `scinstall(1M)` utility (text-based method), go to “How to Install Solaris Software” on page 39 to first install Solaris software.
 - 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 74.

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

This procedure describes how to install the Cluster Control Panel (CCP) software on the 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, plus a common window that sends input to all nodes at one time.

You can use any desktop machine that runs the Solaris 8 or Solaris 9 operating environment as an administrative console. In addition, you can also use the administrative console as a Sun Management Center console and/or server, and as an AnswerBook server. See Sun Management Center documentation for information on how to install Sun Management Center software. See “Sun Cluster AnswerBooks Installation” in *Sun Cluster 3.1 Release Notes* for information on how to install an AnswerBook server.

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.

- 1. Ensure that the Solaris 8 or Solaris 9 operating environment and any Solaris patches are installed on the administrative console.**

All platforms require Solaris 8 with at least the End User System Support software group.

- 2. If you install from the CD-ROM, insert the Sun Cluster 3.1 CD-ROM into the CD-ROM drive of the administrative console.**

If the Volume Management daemon `vold(1M)` is running and configured to manage CD-ROM devices, it automatically mounts the CD-ROM on the `/cdrom/suncluster_3_1` directory.

- 3. Change to the `/cdrom/suncluster_3_1/SunCluster_3.1/Sol_ver/Packages` directory, where `ver` is 8 (for Solaris 8) or 9 (for Solaris 9).**

```
# cd /cdrom/suncluster_3_1/SunCluster_3.1/Sol_ver/Packages
```

- 4. Install the `SUNWcccon` package.**

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

- 5. (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.

6. If you installed from a CD-ROM, eject the CD-ROM.

7. Create an `/etc/clusters` file.

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.

8. 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 (such as a terminal concentrator (TC), System Service Processor (SSP), or Sun Fire system controller), and the port number.

- 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.

```
# 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

9. For convenience, add the `/opt/SUNWcluster/bin` directory to the `PATH` and the `/opt/SUNWcluster/man` directory to the `MANPATH` on the administrative console.

If you installed the `SUNWscman` package, also add the `/usr/cluster/man` directory to the `MANPATH`.

10. Start the CCP utility.

```
# /opt/SUNWcluster/bin/ccp clustername
```

See the procedure “How to Remotely Log In to Sun Cluster” in “Beginning to Administer the Cluster” in *Sun Cluster 3.1 System Administration Guide* for information about how to use the CCP utility. Also see the `ccp(1M)` man page.

11. Install the Solaris operating environment.

- To install Solaris software, go to “How to Install Solaris Software” on page 39.
- To install Solaris and Sun Cluster software by using the `scinstall` JumpStart option, go to “How to Install Solaris and Sun Cluster Software (JumpStart)” on page 74.

▼ How to Install Solaris Software

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

Note – If your nodes are already installed with the Solaris operating environment, you must still reinstall the Solaris software as described in this procedure to ensure successful installation of Sun Cluster software.

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

See the *Sun Cluster 3.x Hardware Administration Manual* 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 35 for requirements and guidelines.

3. Have available your completed “Local File Systems With Mirrored Root Worksheet” in *Sun Cluster 3.1 Release Notes* or “Local File Systems with Non-Mirrored Root Worksheet” in *Sun Cluster 3.1 Release Notes*.

4. Are you using a naming service?

- If no, go to Step 5. You will set up local hostname information in Step 16.
- If yes, add address-to-name mappings for all public hostnames and logical addresses to any naming services (such as NIS or DNS) used by clients for access to cluster services. See “IP Addresses” on page 17 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) is installed and configured on your administrative console, you can use the `cconsole(1M)` utility to display the individual console screens. CCP also opens a master window from which you can send your input to all individual console windows at the same time.

If you do not use CCP, connect to the consoles of each node individually.

Tip – To save time, you can install the Solaris operating environment on each node at the same time.

6. Do the nodes in the cluster use Ethernet adapters?

- If no, go to Step 7.
- If yes, ensure that the `local-mac-address?` variable is correctly set to `true` for Ethernet adapters.

Sun Cluster 3.1 software does not support the `local-mac-address?` variable set to `false` for Ethernet adapters. This is a change from the requirement for Sun Cluster 3.0 software.

a. Display the value of the `local-mac-address?` variable.

- If the node is preinstalled with Solaris software, as superuser run the following command.

```
# /usr/sbin/eeprom local-mac-address?
```

- If the node is not yet installed with Solaris software, run the following command from the `ok` prompt.

```
ok printenv local-mac-address?
```

b. Does the command return `local-mac-address?=true` on each node?

- If yes, the variable settings are correct. Go to Step 7.
- If no, change the variable setting on any node that is not set to `true`.
 - If the node is preinstalled with Solaris software, as superuser run the following command.

```
# /usr/sbin/eeprom local-mac-address?=true
```

- If the node is not yet installed with Solaris software, run the following command from the `ok` prompt.

```
ok setenv local-mac-address? true
```

c. Repeat Step a to verify any changes you made in Step b.

The new setting becomes effective at the next system reboot.

7. Install the Solaris operating environment as instructed in the Solaris installation documentation.

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

You can use any method normally used to install the Solaris operating environment to install the software on new nodes to be installed into a clustered environment. These methods include the Solaris interactive installation program, Solaris JumpStart, and Solaris Web Start.

During Solaris software installation, do the following.

a. Install at least the End User System Support software group.

- If you intend to use the Remote Shared Memory Application Programming Interface (RSMAPI) or use SCI-PCI adapters for the interconnect transport, the required RSMAPI software packages (SUNW`rsm`, SUNW`rsmx`, SUNW`rsmo`, and SUNW`rsmox`) are included with the higher-level software groups. If you install the End User System Support software group, you must install the SUNW`rsm*` packages manually from the Solaris CD-ROM at Step 12.
- If you intend to use SunPlex Manager, the required Apache software packages (SUNW`apchr` and SUNW`apchu`) are included with the higher-level software groups. If you install the End User System Support software group, you must install the SUNW`apch*` packages manually from the Solaris CD-ROM at Step 13.

See “Solaris Software Group Considerations” on page 13 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-devices subsystem. If you intend to use SunPlex Manager 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 used by `scinstall`.

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

- If you intend to use SunPlex Manager to install Solstice DiskSuite software (Solaris 8), configure Solaris Volume Manager software (Solaris 9), or install Sun Cluster HA for NFS or Sun Cluster HA for Apache in addition to installing Sun Cluster software, create a 20-Mbyte file system on slice 7 with a mount-point name of `/sds`.

Otherwise, create any file system partitions needed to support your volume manager software as described in “System Disk Partitions” on page 14.

c. Choose auto reboot.

Note – The Solaris installation tool installs Solaris software and reboots the node before it displays the next prompts.

d. For ease of administration, set the same root password on each node.

e. Answer no when asked whether to enable automatic power-saving shutdown.

You must disable automatic shutdown in Sun Cluster configurations. See the `pmconfig(1M)` and `power.conf(4)` man pages for more information.

Note – The Solaris interface groups feature is disabled by default during Solaris software installation. Interface groups are not supported in a Sun Cluster configuration and should not be enabled. See the `ifconfig(1M)` man page for more information about Solaris interface groups.

8. Are you installing a new node to an existing cluster?

- If yes, go to Step 9.
- If no, skip to Step 12.

9. Have you added the new node to the cluster’s authorized-node list?

- If yes, go to Step 10.
- If no, run `scsetup(1M)` from another, active cluster node to add the new node’s name to the list of authorized cluster nodes. See “Task Map: Adding a Cluster Node to an Existing Cluster” in *Sun Cluster 3.1 System Administration Guide* for procedures.

10. Create a mount point on the new node for each cluster file system in the cluster.

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

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

b. 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.

11. Is VERITAS Volume Manager (VxVM) installed on any nodes that are already in the cluster?

- If yes, ensure that the same `vxio` number is used on the VxVM-installed nodes and 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 by changing the `/etc/name_to_major` entry to use a different number.

- If no, go to Step 12.

12. Do you intend to use the Remote Shared Memory Application Programming Interface (RSMAPI) or use SCI-PCI adapters for the interconnect transport?

- If yes **and** you installed the End User System Support software group, install the `SUNWrsm*` packages from the Solaris CD-ROM.

```
# pkgadd -d . SUNWrsm SUNWrsmx SUNWrsmo SUNWrsmox
```

- If no, or if you installed a higher-level software group, go to Step 13.

13. Do you intend to use SunPlex Manager?

- If yes **and** you installed the End User System Support software group, install the `SUNWapch*` packages from the Solaris CD-ROM.

```
# pkgadd -d . SUNWapchr SUNWapchu
```

- If no, or if you installed a higher-level software group, go to Step 14.

Apache software packages must already be installed before SunPlex Manager is installed.

14. Install any Solaris software patches.

See “Patches and Required Firmware Levels” in *Sun Cluster 3.1 Release Notes* for the location of patches and installation instructions. If necessary, view the `/etc/release` file to see the exact version of Solaris software that is installed on a node.

15. Install any hardware-related patches and download any needed firmware 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.

16. 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.

17. Do you intend to use dynamic reconfiguration on Sun Enterprise 10000 servers?

- If yes, on each node add the following entry to the `/etc/system` file.

```
set kernel_cage_enable=1
```

This entry becomes effective after the next system reboot. See the *Sun Cluster 3.1 System Administration Guide* for procedures to perform dynamic reconfiguration tasks in a Sun Cluster configuration. See your server documentation for more information about dynamic reconfiguration.

- If no, go to Step 18.

18. Install Sun Cluster software on your cluster nodes.

- To use SunPlex Manager, go to “Using SunPlex Manager to Install Sun Cluster Software” on page 62.
- To use `scinstall`, go to “How to Install Sun Cluster Software on the First Cluster Node (`scinstall`)” on page 44.

▼ How to Install Sun Cluster Software on the First Cluster Node (`scinstall`)

After you install the Solaris operating environment, perform this task on one node of the cluster to install Sun Cluster software and establish the cluster.

Note – If you used the `scinstall(1M)` custom JumpStart or SunPlex Manager installation method, the Sun Cluster software is already installed. Go to “How to Configure the Name Service Switch” on page 87.

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

You must install Solaris software as described in “How to Install Solaris Software” on page 39. 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 you intend to install on the cluster.

2. Do you intend to use SunPlex Manager?

- If yes, ensure that the Apache software packages are installed on the node. If you installed the Solaris End User System Support software group, install the `SUNwapch*` packages from the Solaris CD-ROM.

```
# pkgadd -d . SUNWapchr SUNWapchu
```

The Apache software packages are automatically installed if you installed a higher-level Solaris software group.

- If no, go to Step 3.

3. Have available the following completed configuration planning worksheets:

- “Cluster and Node Names Worksheet” in *Sun Cluster 3.1 Release Notes*
- “Cluster Interconnect Worksheet” in *Sun Cluster 3.1 Release Notes*

See “Planning the Sun Cluster Environment” on page 17 for planning guidelines.

4. Become superuser on the cluster node to install.

5. If you install from the CD-ROM, insert the Sun Cluster 3.1 CD-ROM into the CD-ROM drive of the node to install and configure.

If the Volume Management daemon `vold(1M)` is running and configured to manage CD-ROM devices, it automatically mounts the CD-ROM on the `/cdrom/suncluster_3_1` directory.

6. Do you intend to use the Remote Shared Memory Application Programming Interface (RSMAPI) or use SCI-PCI adapters for the interconnect transport?

- If no, go to Step 7.
- If yes, install additional packages from the `Packages` directory of the Sun Cluster 3.1 CD-ROM. These packages are not automatically installed by `scinstall(1M)`.

```
# cd /cdrom/suncluster_3_1/SunCluster_3.1/Sol_ ver/Packages
# pkgadd -d . packages
```

Note – In the CD-ROM path, replace *ver* with 8 (for Solaris 8) or 9 (for Solaris 9)

The following table lists the Sun Cluster 3.1 packages required to support the RSMAPI or SCI-PCI adapters and the order in which you must install them.

TABLE 2-2 Sun Cluster 3.1 Packages to Support the RSMAPI and SCI-PCI Adapters

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

7. Change to the `/cdrom/suncluster_3_1/SunCluster_3.1/Sol_ ver/Tools` directory, where *ver* is 8 (for Solaris 8) or 9 (for Solaris 9).

The following example uses the path to the Solaris 8 version of Sun Cluster software.

```
# cd /cdrom/suncluster_3_1/SunCluster_3.1/Sol_8/Tools
```

8. Start the `scinstall(1M)` utility.

```
# ./scinstall
```

Follow the prompts to install Sun Cluster software, using the information from your configuration planning worksheets.

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

- Interactive `scinstall` enables you to type ahead. Therefore, do not press Return 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. If you press Control-D to abort the session after Sun Cluster software is installed, `scinstall` asks you whether you want it to de-install those packages.
- Your session answers are stored as defaults for the next time you run this menu option. Default answers display between brackets ([]) at the end of the prompt.

Tip – Until the node is successfully booted in cluster mode, you can rerun `scinstall` and change the configuration information as needed. However, if bad configuration data for the node was pushed over to the established portion of the cluster, you might first need to remove the bad information. To do this, log in to one of the active cluster nodes, then use the `scconf(1M)` command to remove the bad adapter, junction, or cable information.

9. From the Main Menu, type **1** (Establish a new cluster).

After you read the information in the Establishing a New Cluster screen, type **yes** to continue to the next screen.

```
*** Main Menu ***

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

* 1) Establish a new cluster using this machine as the first node
* 2) Add this machine as a node in an established cluster
  3) Configure a cluster to be JumpStarted from this install server
  4) Add support for new data services to this cluster node
  5) Print release information for this cluster node

* ?) Help with menu options
* q) Quit

Option: 1

*** Establishing a New Cluster ***
...
```

Do you want to continue (yes/no) [yes]? **yes**

10. When prompted whether to continue to install Sun Cluster software packages, type yes.

>>> Software Package Installation <<<

Installation of the Sun Cluster framework software packages will take a few minutes to complete.

Is it okay to continue (yes/no) [yes]? **yes**

** Installing SunCluster 3.0 **
SUNWscr.....done
...Hit ENTER to continue:

After all packages are installed, press Return to continue to the next screen.

11. Specify the cluster name.

>>> Cluster Name <<<

...

What is the name of the cluster you want to establish? *clustername*

12. Run the preinstallation check.

>>> Check <<<

This step runs `sccheck(1M)` to verify that certain basic hardware and software pre-configuration requirements have been met. If `sccheck(1M)` detects potential problems with configuring this machine as a cluster node, a list of warnings is printed.

Hit ENTER to continue:

If the `sccheck(1M)` command detects any problems, it displays information about the problems found and prompts you for your next action.

The `sccheck` utility has detected the following potential problems:

...

Hit ENTER to continue:

...

What would you like to do?

- 1) Proceed with the installation anyway
- 2) Retry `sccheck`
- q) Return to the main menu

Option:

If you see this `sccheck` error message, type one of the following responses.

- If you know that you can safely ignore the `sccheck` messages, for example from notification in the *Sun Cluster 3.1 Release Notes* that a recent qualification will trigger a harmless error message, type **1** to continue installation.
- If you can correct the problem without quitting the `scinstall` utility, correct the problem and type **2** to run `sccheck` again.

- Type **q** to stop installation. Correct the problem detected by `sccheck`, then return to Step 8 and restart `scinstall`.

13. Specify the names of the *other* nodes that will become part of this cluster.

```
>>> Cluster Nodes <<<
...
Node name: node2
Node name (Ctrl-D to finish): <Control-D>

This is the complete list of nodes:
...
Is it correct (yes/no) [yes]?
```

14. Specify whether to use data encryption standard (DES) authentication.

By default, Sun Cluster software permits a node to connect to the cluster only if the node is physically connected to the private interconnect and if the node name was specified in Step 13. However, the node actually communicates with the sponsoring node over the public network, since the private interconnect is not yet fully configured. DES authentication provides an additional level of security at installation time by enabling the sponsoring node to more reliably authenticate nodes that attempt to contact it to update the cluster configuration.

If you choose to use DES authentication for additional security, you must configure all necessary encryption keys before any node can join the cluster. See the `keyserv(1M)` and `publickey(4)` man pages for details.

```
>>> Authenticating Requests to Add Nodes <<<
...
Do you need to use DES authentication (yes/no) [no]?
```

15. Specify the private network address and netmask.

```
>>> Network Address for the Cluster Transport <<<
...
Is it okay to accept the default network address (yes/no) [yes]?
Is it okay to accept the default netmask (yes/no) [yes]?
```

Note – You cannot change the private network address after the cluster is successfully formed.

16. Specify whether the cluster uses transport junctions.

- If this is a two-node cluster, specify whether you intend to use transport junctions.

```
>>> Point-to-Point Cables <<<
...
Does this two-node cluster use transport junctions (yes/no) [yes]?
```

Tip – You can specify that the cluster uses transport junctions, regardless of whether the nodes are directly connected to each other. If you specify that the cluster uses transport junctions, you can more easily add new nodes to the cluster in the future.

- If this cluster has three or more nodes, you must use transport junctions. Press Return to continue to the next screen.

```
>>> Point-to-Point Cables <<<
```

```
...
```

```
Since this is not a two-node cluster, you will be asked to configure  
two transport junctions.
```

```
Hit ENTER to continue:
```

17. Does this cluster use transport junctions?

- If yes, specify names for the transport junctions. You can use the default names `switchN` or create your own names.

```
>>> Cluster Transport Junctions <<<
```

```
...
```

```
What is the name of the first junction in the cluster [switch1]?
```

```
What is the name of the second junction in the cluster [switch2]?
```

- If no, skip to Step 18.

18. Specify the first cluster interconnect transport adapter.

```
>>> Cluster Transport Adapters and Cables <<<
```

```
...
```

```
Select the first cluster transport adapter to use:
```

```
1) adapter
```

```
2) adapter
```

```
...
```

```
N) Other
```

```
Option:
```

The `scinstall` utility lists all Ethernet adapters found through autodiscovery. To configure adapters that are not listed, such as SCI-PCI adapters, type the number for Other, then specify the adapter information requested in the subsequent menus.

19. If your cluster uses transport junctions, specify the name of the first transport junction and its port.

Otherwise, skip to Step 20.

```
Name of the junction to which "adapter" is connected [switch1]?
```

```
Use the default port name for the "adapter" connection (yes/no) [yes]?
```

```
Hit ENTER to continue:
```

Note – If your configuration uses SCI adapters, do not accept the default when you are prompted for the adapter connection (the port name). Instead, provide the port name (0, 1, 2, or 3) found on the Dolphin switch itself, to which the node is *physically* cabled. The following example shows the prompts and responses for declining the default port name and specifying the Dolphin switch port name 0.

```
Use the default port name for the "adapter" connection
(yes/no) [yes]? no
What is the name of the port you want to use? 0
```

20. Choose the second cluster interconnect transport adapter.

```
Select the second cluster transport adapter to use:
  1) adapter
  2) adapter
...
  N) Other
```

Option:

You configure two adapters by using the `scinstall` command. You can configure additional adapters after Sun Cluster software is installed by using the `scsetup` utility.

21. If your cluster uses transport junctions, specify the name of the second transport junction and its port.

Otherwise, skip to Step 22.

```
Name of the junction to which "adapter" is connected [switch2]?
Use the default port name for the "adapter" connection (yes/no) [yes]?
```

Hit ENTER to continue:

Note – If your configuration uses SCI adapters, do not accept the default when you are prompted for the adapter port name. Instead, provide the port name (0, 1, 2, or 3) found on the Dolphin switch itself, to which the node is *physically* cabled. The following example shows the prompts and responses for declining the default port name and specifying the Dolphin switch port name 0.

```
Use the default port name for the "adapter" connection
(yes/no) [yes]? no
What is the name of the port you want to use? 0
```

22. Specify the global devices file system name.

```
>>> Global Devices File System <<<
...
The default is to use /globaldevices.
```

Is it okay to use this default (yes/no) [yes]?

23. Do you have any Sun Cluster software patches to install?

```
>>> Automatic Reboot <<<
...
Do you want scinstall to reboot for you (yes/no) [yes]?
```

- If yes, type **no** in the Automatic Reboot screen to decline automatic reboot.
- If no, type **yes** to accept automatic reboot.

24. Accept or decline the generated `scinstall` command.

The `scinstall` command generated from your input is displayed for confirmation.

```
>>> Confirmation <<<

Your responses indicate the following options to scinstall:

scinstall -ik \
...
Are these the options you want to use (yes/no) [yes]?
Do you want to continue with the install (yes/no) [yes]?
```

- If you accept the command and continue the installation, `scinstall` processing continues. “Example—Installing Sun Cluster Software on the First Node” on page 52 shows an example of the output you might see during `scinstall` processing.

Sun Cluster installation output is logged in the `/var/cluster/logs/install/scinstall.log.pid` file, where *pid* is the process ID number of the `scinstall` instance.

Note – Unless you have installed your own `/etc/inet/ntp.conf` file, the `scinstall` command installs a default `ntp.conf` file for you. Because the default file is shipped with references to the maximum possible number of nodes, 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 100 for information on how to suppress these messages under otherwise normal cluster conditions.

- If you decline the command, `scinstall` asks if you want to de-install the Sun Cluster software.

```
Do you want to de-install the Sun Cluster software (yes/no) [no]?
```

After `scinstall` returns you to the Main Menu, you can rerun menu option 1 and provide different answers. Your previous session answers display as the defaults.

25. Install any Sun Cluster software patches.

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

26. Reboot the node to establish the cluster.

If you rebooted the node after you installed patches in Step 25, you do not need to reboot the node a second time.

The first node reboot after Sun Cluster software installation forms the cluster and establishes this node as the *first-installed* node of the cluster.

27. Install additional nodes in the cluster.

Go to "How to Install Sun Cluster Software on Additional Cluster Nodes (scinstall)" on page 53.

Example—Installing Sun Cluster Software on the First Node

The following example shows the `scinstall` command executed and progress messages displayed as `scinstall` installation tasks are completed on the node `phys-schost-1`, which is the first node to be installed in the cluster. The cluster will have four nodes and uses Ethernet adapters and transport junctions.

```
>>> Confirmation <<<

Your responses indicate the following options to scinstall:

scinstall -ik \
  -C sc-cluster \
  -F \
  -T node=phys-schost-1,node=phys-schost-2,node=phys-schost-3,node=phys
-schost-4,authtype=sys \
  -A trtype=dlpi,name=hme1 -A trtype=dlpi,name=hme3 \
  -B type=switch,name=switch1 -B type=switch,name=switch2 \
  -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

Initializing cluster name to "sc-cluster" ... done
Initializing authentication options ... done
Initializing configuration for adapter "hme1" ... done
Initializing configuration for adapter "hme3" ... done
Initializing configuration for junction "switch1" ... done
Initializing configuration for junction "switch2" ... done
Initializing configuration for cable ... done
Initializing configuration for cable ... 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 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.061401232831
Power management is incompatible with the HA goals of the cluster.
Please do not attempt to re-configure power management.

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.10559

Rebooting ...
```

▼ How to Install Sun Cluster Software on Additional Cluster Nodes (scinstall)

After you install Sun Cluster software on the first-installed node and establish the cluster, perform this task on each additional node to install in the cluster. You can also use this procedure to add new nodes to an existing cluster.

Note – If you used the `scinstall(1M)` custom JumpStart or SunPlex Manager installation method, the Sun Cluster software is already installed. Go to “How to Configure the Name Service Switch” on page 87.

1. Ensure that the first-installed node is successfully installed with Sun Cluster software and that the cluster is established.

For installation procedures to establish a cluster, see “How to Install Sun Cluster Software on the First Cluster Node (scinstall)” on page 44.

2. If you are adding a new node to an existing, fully installed cluster, ensure that you have performed the following tasks.

a. Prepare the cluster to accept a new node.

Follow instructions in the procedure “How to Add a Cluster Node to the Authorized Node List” in “Adding and Removing a Cluster Node” in *Sun Cluster 3.1 System Administration Guide*.

b. Install Solaris software on the new node.

You must install Solaris software as described in “How to Install Solaris Software” on page 39. 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 you intend to install on the cluster.

3. Do you intend to use SunPlex Manager?

- If yes, ensure that the Apache software packages are installed on the node. If you installed the Solaris End User System Support software group, install the SUNWapch* packages from the Solaris CD-ROM.

```
# pkgadd -d . SUNWapchr SUNWapchu
```

The Apache software packages are automatically installed if you installed a higher-level Solaris software group.

- If no, go to Step 4.

4. Have available the following completed configuration planning worksheets:

- “Cluster and Node Names Worksheet” in *Sun Cluster 3.1 Release Notes*
- “Cluster Interconnect Worksheet” in *Sun Cluster 3.1 Release Notes*

See “Planning the Sun Cluster Environment” on page 17 for planning guidelines.

5. Become superuser on the cluster node to install.

6. If you install from the CD-ROM, insert the Sun Cluster 3.1 CD-ROM into the CD-ROM drive of the node to install and configure.

If the Volume Management daemon `vold(1M)` is running and configured to manage CD-ROM devices, it automatically mounts the CD-ROM on the `/cdrom/suncluster_3_1` directory.

7. Do you intend to use the Remote Shared Memory Application Programming Interface (RSMAPI) or use SCI-PCI adapters for the interconnect transport?

- If no, go to Step 8.
- If yes, install additional packages from the `Packages` directory of the Sun Cluster 3.1 CD-ROM. These packages are not automatically installed by `scinstall(1M)`.

```
# cd /cdrom/suncluster_3_1/SunCluster_3.1/Sol_ver/Packages
# pkgadd -d . packages
```

Note – In the CD-ROM path, replace *ver* with 8 (for Solaris 8) or 9 (for Solaris 9)

The following table lists the Sun Cluster 3.1 packages required to support the RSMAPI or SCI-PCI adapters and the order in which you must install them.

TABLE 2-3 Sun Cluster 3.1 Packages to Support the RSMAPI and SCI-PCI Adapters

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

8. Change to the `/cdrom/suncluster_3_1/SunCluster_3.1/Sol_ver/Tools` directory, where *ver* is 8 (for Solaris 8) or 9 (for Solaris 9).

The following example uses the path to the Solaris 8 version of Sun Cluster software.

```
# cd /cdrom/suncluster_3_1/SunCluster_3.1/Sol_8/Tools
```

9. Start the `scinstall` utility.

```
# ./scinstall
```

You can start this step while software is still being installed on the first-installed node. If necessary, the second node waits for the first node to complete installation.



Caution – *Do not* reboot or shut down the first-installed node while any other nodes are being installed, even if you use another node as the sponsoring node. Until quorum votes are assigned to the cluster nodes and cluster install mode is disabled, the first-installed node, which established the cluster, is the only node that has a quorum vote. If the cluster is still in install mode, you will cause a system panic because of lost quorum if you reboot or shut down the first-installed node.

10. From the Main Menu, type 2 (Add this machine as a node).

After you read the information in the Establishing a New Cluster screen, type **yes** to continue to the next screen.

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

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

- * 1) Establish a new cluster using this machine as the first node
- * 2) Add this machine as a node in an established cluster
- 3) Configure a cluster to be JumpStarted from this install server

```
4) Add support for new data services to this cluster node
5) Print release information for this cluster node
```

```
* ?) Help with menu options
* q) Quit
```

```
Option: 2
```

```
*** Adding a Node to an Established Cluster ***
```

```
...
```

```
Do you want to continue (yes/no) [yes]? yes
```

11. When prompted whether to continue to install Sun Cluster software packages, type *yes*.

```
>>> Software Installation <<<
```

```
Installation of the Sun Cluster framework software packages will only
take a few minutes to complete.
```

```
Is it okay to continue (yes/no) [yes]? yes
```

```
** Installing SunCluster 3.0 **
```

```
SUNWscr.....done
```

```
...Hit ENTER to continue:
```

After all packages are installed, press Return to continue to the next screen.

12. Specify the name of any existing cluster node, referred to as the *sponsoring node*.

```
>>> Sponsoring Node <<<
```

```
...
```

```
What is the name of the sponsoring node? node1
```

13. Specify the cluster name.

```
>>> Cluster Name <<<
```

```
...
```

```
What is the name of the cluster you want to join? clustername
```

14. Run the preinstallation check.

```
>>> Check <<<
```

```
This step runs sccheck(1M) to verify that certain basic hardware and
software pre-configuration requirements have been met. If sccheck(1M)
detects potential problems with configuring this machine as a cluster
node, a list of warnings is printed.
```

```
Hit ENTER to continue:
```

If the `sccheck` command detects any problems, it displays information about the problems found and prompts you for your next action.

```
The sccheck utility has detected the following potential problems:
```

```
...
```

```
Hit ENTER to continue:
```

```
...
```


What would you like to do?

- 1) Proceed with the installation anyway
- 2) Retry sccheck
- q) Return to the main menu

Option:

If you see this `sccheck` error message, type one of the following responses.

- If you know that you can safely ignore the `sccheck` messages, for example from notification in the *Sun Cluster 3.1 Release Notes* that a recent qualification will trigger a harmless error message, type **1** to continue installation.
- If you can correct the problem without quitting the `scinstall` utility, correct the problem and type **2** to run `sccheck` again.
- Type **q** to stop installation. Correct the problem detected by `sccheck`, then return to Step 9 and restart `scinstall`.

15. Specify whether to use autodiscovery to configure the cluster transport.

If your configuration does not use Ethernet adapters, answer **no** and go to Step 17.

```
>>> Autodiscovery of Cluster Transport <<<
```

```
If you are using ethernet adapters as your cluster transport
adapters, autodiscovery is the best method for configuring the
cluster transport.
```

```
Do you want to use autodiscovery (yes/no) [yes]?
```

```
...
```

```
The following connections were discovered:
```

```
node1:adapter switch node2:adapter
node1:adapter switch node2:adapter
```

```
Is it okay to add these connections to the configuration (yes/no) [yes]?
```

16. Did you use autodiscovery in Step 15?

- If yes, skip to Step 24.
- If no, go to Step 17.

17. Specify whether this is a two-node cluster.

```
>>> Point-to-Point Cables <<<
```

```
...
```

```
Is this a two-node cluster (yes/no) [yes]?
```

18. Did you specify that this is a two-node cluster?

- If yes, specify whether to use transport junctions.

```
Does this two-node cluster use transport junctions (yes/no) [yes]?
```

- If no, press Return to continue. You must use transport junctions if a cluster contains three or more nodes.

Since this is not a two-node cluster, you will be asked to configure two transport junctions.

Hit ENTER to continue:

19. Did you specify that the cluster will use transport junctions?

- If yes, specify the transport junctions.

```
>>> Cluster Transport Junctions <<<
...
What is the name of the first junction in the cluster [switch1]?
What is the name of the second junction in the cluster [switch2]?
```

- If no, skip to Step 20.

20. Specify the first cluster interconnect transport adapter.

Type **help** to list all transport adapters available to the node.

```
>>> Cluster Transport Adapters and Cables <<<
...
What is the name of the first cluster transport adapter (help)? adapter
```

21. Specify what the first transport adapter connects to.

- If the transport adapter uses a transport junction, specify the name of the junction and its port.

```
Name of the junction to which "adapter" is connected [switch1]?
...
Use the default port name for the "adapter" connection (yes/no) [yes]?
```

- If the transport adapter does not use a transport junction, specify the name of the other transport adapter it connects to.

```
Name of adapter on "node1" to which "adapter" is connected? adapter
```

22. Specify the second cluster interconnect transport adapter.

Type **help** to list all transport adapters available to the node.

```
What is the name of the second cluster transport adapter
(help)? adapter
```

23. Specify what the second transport adapter connects to.

- If the transport adapter uses a transport junction, specify the name of the junction and its port.

```
Name of the junction to which "adapter" is connected [switch2]?
Use the default port name for the "adapter" connection (yes/no) [yes]?
```

Hit ENTER to continue:

- If the transport adapter does not use a transport junction, specify the name of the other transport adapter it connects to.

```
Name of adapter on "node1" to which "adapter" is connected? adapter
```

24. Specify the global devices file system name.

```
>>> Global Devices File System <<<
...
The default is to use /globaldevices.

Is it okay to use this default (yes/no) [yes]?
```

25. Do you have any Sun Cluster software patches to install?

```
>>> Automatic Reboot <<<
...
Do you want scinstall to reboot for you (yes/no) [yes]?
```

- If yes, type **no** in the Automatic Reboot screen to decline automatic reboot.
- If no, type **yes** to accept automatic reboot.

26. Accept or decline the generated `scinstall` command.

The `scinstall` command generated from your input is displayed for confirmation.

```
>>> Confirmation <<<

Your responses indicate the following options to scinstall:

scinstall -i \
...
Are these the options you want to use (yes/no) [yes]?
Do you want to continue with the install (yes/no) [yes]?
```

- If you accept the command and continue the installation, `scinstall` processing continues. “Example—Installing Sun Cluster Software on the First Node” on page 52 shows an example of the output you might see during `scinstall` processing. If the sponsoring node is not yet established in the cluster, `scinstall` waits for the sponsoring node to become available.

Sun Cluster installation output is logged in the `/var/cluster/logs/install/scinstall.log.pid` file, where *pid* is the process ID number of the `scinstall` instance.

Note – Unless you have installed your own `/etc/inet/ntp.conf` file, the `scinstall` command installs a default `ntp.conf` file for you. Because the default file is shipped with references to eight nodes, 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 100 for information on how to suppress these messages under otherwise normal cluster conditions.

- If you decline the command, `scinstall` asks if you want to de-install the Sun Cluster software.

Do you want to de-install the Sun Cluster software (yes/no) [no]?
After `scinstall` returns you to the Main Menu, you can rerun menu option 2 and provide different answers. Your previous session answers display as the defaults.

27. Install any Sun Cluster software patches.

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

28. Reboot the node to establish the cluster.

If you rebooted the node after you installed patches in Step 27, you do not need to reboot the node a second time.



Caution – *Do not* reboot or shut down the first-installed node while any other nodes are being installed, even if you use another node in the cluster as the sponsoring node. Until quorum votes are assigned to the cluster nodes and cluster install mode is disabled, the first-installed node, which established the cluster, is the only node that has a quorum vote. If the cluster is still in install mode, you will cause a system panic because of lost quorum if you reboot or shut down the first-installed node. Cluster nodes remain in install mode until the first time you run the `scsetup(1M)` command, during the procedure “How to Perform Post-Installation Setup” on page 90.

29. Repeat this procedure on each additional node until all nodes are fully configured.

You do not need to wait for the second node to complete installation and reboot into the node before you begin installation on additional nodes.

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

Go to “How to Configure the Name Service Switch” on page 87.

Example—Installing Sun Cluster Software on the Second Node

The following example shows the `scinstall` command executed and progress messages displayed as `scinstall` installation tasks are completed on the node `phys-schost-2`, which is the second node to be installed in the cluster.

```
>>> 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 \  
-B type=switch,name=switch1 -B type=switch,name=switch2 \  
-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-2" 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-2" ... done (id=2)

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 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 ...

Using SunPlex Manager to Install Sun Cluster Software

Note – To add a new node to an existing cluster, do not use SunPlex Manager. Instead, go to “How to Install Sun Cluster Software on the First Cluster Node (`scinstall`)” on page 44.

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

- (Solaris 8 only) Solstice DiskSuite software—Also configures up to three metaset and associated metadevices, and creates and mounts cluster file systems for each.
- Sun Cluster HA for NFS data service
- Sun Cluster HA for Apache scalable data service

For Solaris 9, Solaris Volume Manager software is already installed as part of Solaris software installation. You can use SunPlex Manager to configure up to three metasets and associated Solstice DiskSuite metadevices or Solaris Volume Manager volumes, and to create and mount cluster file systems for each.

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

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

Software Package	Installation Requirements
Solstice DiskSuite/Solaris Volume Manager	A partition that uses <code>/sds</code> as the mount-point name. The partition must be at least 20 Mbytes.
Sun Cluster HA for NFS data service	<ul style="list-style-type: none">■ At least two shared disks of the same size which are connected to the same set of nodes.■ Solstice DiskSuite software installed, or Solaris Volume Manager software configured, by SunPlex Manager.■ 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 and is on the same subnet as the base hostnames of the cluster nodes.

TABLE 2-4 Requirements to Use SunPlex Manager to Install Software (Continued)

Software Package	Installation Requirements
Sun Cluster HA for Apache scalable data service	<ul style="list-style-type: none"> ■ At least two shared disks of the same size which are connected to the same set of nodes. ■ Solstice DiskSuite software installed, or Solaris Volume Manager software configured, by SunPlex Manager. ■ 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 and is on the same subnet as the base hostnames of the cluster nodes.

The following table lists each metaset name and cluster file system mount point created by SunPlex Manager, depending on the number of shared disks connected to the node. For example, if a node has four shared disks connected to it, SunPlex Manager creates the `mirror-1` and `mirror-2` metasets, but 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 Manager

Shared Disks	Metaset Name	Cluster File System Mount Point	Purpose
First pair of shared disks	<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 of shared disks	<code>mirror-2</code>	<code>/global/mirror-2</code>	unused
Third pair of shared disks	<code>mirror-3</code>	<code>/global/mirror-3</code>	unused

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

SunPlex Manager 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 Manager server. The following characters are accepted by SunPlex Manager.

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

This filter will potentially cause problems in the following two areas:

- **Password entry for Sun Open Net Environment (Sun ONE) services** – If the password contains unusual characters, these characters will be stripped out resulting in two problems. Either the resulting password will be under 8 characters and will fail, or the application will be configured with a different password than the user expects.
- **Localization** – Alternative character sets (for example: accented characters or Asian characters) will not work for input.

▼ How to Install SunPlex Manager Software

The SunPlex Manager graphical user interface (GUI) provides an easy way to install and administer Sun Cluster software. Follow this procedure to install SunPlex Manager software on your cluster.

Note – If you intend to install Sun Cluster software by using another method, you do not need to perform this procedure. The `scinstall` command installs SunPlex Manager for you as part of the installation process.

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 39. 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 you intend to install on the cluster.

2. Become superuser on a cluster node.

3. Are Apache software packages installed on the node?

- If yes, go to Step 4.

The Apache software packages are included in the Solaris Entire Distribution software group and all higher-level software groups. If you installed a lower-level software group, use the `pkginfo(1)` command to determine whether the software packages in Step c are already installed.

- If no, install Apache software packages.

a. If you install from the CD-ROM, insert the Solaris 8 Software 2 of 2 CD-ROM into the CD-ROM drive of the node.

If the Volume Management daemon `vold(1M)` is running and configured to manage CD-ROM devices, it automatically mounts the CD-ROM.

b. 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 following order.

```
# pkgadd -d . SUNWapchr SUNWapchu SUNWapchd
```

d. Eject the Solaris CD-ROM.

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.

4. Install the SunPlex Manager software packages.

a. If you install from the CD-ROM, insert the Sun Cluster 3.1 CD-ROM into the CD-ROM drive of the node.

If the Volume Management daemon `vold(1M)` is running and configured to manage CD-ROM devices, it automatically mounts the CD-ROM on the /cdrom/suncluster_3_1 directory.

b. Change to the

/cdrom/suncluster_3_1/SunCluster_3.1/Sol_*ver*/Packages directory, where *ver* is 8 (for Solaris 8) or 9 (for Solaris 9) .

The following example uses the path to the Solaris 8 version of Sun Cluster software.

```
# cd /cdrom/suncluster_3_1/SunCluster_3.1/Sol_8/Packages
```

c. Install the SunPlex Manager software packages and answer yes for all prompts.

```
# pkgadd -d . SUNWscva SUNWscvr SUNWscvw
```

d. Eject the Sun Cluster CD-ROM.

5. Repeat Step 2 through Step 4 on each node of the cluster.

6. Is the root password the same on every node of the cluster?

- If yes, go to Step 7.
- If no, set the root password to the same value on each node of the cluster. If necessary, also use `chkey(1)` to update the RPC key pair.

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

The root password must be the same on all nodes in the cluster to use the root password to access SunPlex Manager.

7. **Do you intend to set up additional user accounts to access SunPlex Manager?**

- If yes, go to Step 8.
- If no, go to “How to Install Sun Cluster Software (SunPlex Manager)” on page 68 to install Sun Cluster software on your cluster nodes.

Users who do not use the `root` system account nor have a user account set up on a particular node cannot access the cluster through SunPlex Manager from that node. Also, users cannot manage that node through another cluster node to which the users do have access.

8. **Determine how to set up user accounts to access SunPlex Manager.**

In addition to `root`-user access, users can log in to SunPlex Manager with a user account that has role-based access control (RBAC). Go to one of the procedures listed in the following table to set up user accounts.

TABLE 2-6 Methods to Set Up SunPlex Manager User Accounts

Method	Go to This Procedure
Add RBAC authorization to an existing user account.	“How to Add RBAC Authorization to an Existing User Account” on page 66
Create a new user account that has RBAC authorization.	“How to Create a New User Account” on page 67

Note – If you assign RBAC authorization to a non-`root` user account, that user account can perform administrative actions usually performed only by `root`.

See “Role-Based Access Control” in *System Administration Guide, Volume 2* or “Using the Solaris Management Tools With RBAC (Task Map)” in *System Administration Guide: Basic Administration* for more information.

▼ How to Add RBAC Authorization to an Existing User Account

Add RBAC authorization to an existing user account. This enables the user to log in to SunPlex Manager by using the user’s regular system password and have access to full SunPlex Manager functionality.

Note – If you assign RBAC authorization to a non-`root` user account, that user account can perform a set of administrative actions usually performed only by `root`.

1. **Become superuser on a node of the cluster.**

2. Add the following entry to the `/etc/user_attr` file.

```
# vi /etc/user_attr
username:::type=normal;auths=solaris.cluster.admin
```

3. Repeat on each remaining node of the cluster.

4. Use SunPlex Manager to install Sun Cluster software.

Go to “How to Install Sun Cluster Software (SunPlex Manager)” on page 68.

▼ How to Create a New User Account

Create a new user account on all nodes the cluster.

Note – If you assign RBAC authorization to a non-root user account, that user account can perform a set of administrative actions usually performed only by root.

1. Become superuser on a node of the cluster.

2. Create the new user account.

```
# useradd -d dir -A solaris.cluster.admin login
```

`-d dir` Specifies the home directory of the new user

`-A solaris.cluster.admin` Assigns `solaris.cluster.admin` authorization to the new user account

`login` Name of the new user account

Note – The user name must be unique and must not already exist either on the local machine or in the network name service.

See the `useradd(1M)` man page for more information about creating user accounts.

3. Set the password.

```
# passwd login
```

4. Repeat on each remaining node of the cluster.

Ensure that the password for the user account is the same on all nodes of the cluster.

5. Use SunPlex Manager to install Sun Cluster software.

Go to “How to Install Sun Cluster Software (SunPlex Manager)” on page 68.

▼ How to Install Sun Cluster Software (SunPlex Manager)

Note – To add a new node to an existing cluster, do not use SunPlex Manager. Instead, go to “How to Install Sun Cluster Software on Additional Cluster Nodes (scinstall)” on page 53.

Perform this procedure to use SunPlex Manager to install 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 disksets (Solaris 9). If you use SunPlex Manager to install Solstice DiskSuite software or to configure Solaris Volume Manager disksets, you can also install the Sun Cluster HA for NFS data service or the scalable Sun Cluster HA for Apache data service or both.

The installation process might take from 30 minutes to two or more hours, depending on the number of cluster nodes, choice of data services, and number of disks in your cluster configuration.

- 1. 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 64.
- 2. Ensure that the cluster configuration meets the requirements to use SunPlex Manager to install software.**
See “Using SunPlex Manager to Install Sun Cluster Software” on page 62 for installation requirements and restrictions.
- 3. Do you intend to install Sun Cluster HA for NFS or Sun Cluster HA for Apache?**
 - If no, go to Step 4.
 - If yes, ensure that the cluster configuration meets all applicable requirements. See “Using SunPlex Manager to Install Sun Cluster Software” on page 62.
- 4. Prepare file system paths to a CD-ROM image of each software product you intend to install.**
 - a. Provide each CD-ROM image in a location that is available to each node.**
The CD-ROM images must be 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 exported to the network from machines outside the cluster.
 - Exported file systems on machines outside the cluster.
 - CD-ROM images copied to local file systems on each node of the cluster. The local file system must use the same name on each node.

b. Record the path to each CD-ROM image.

You will provide this information to SunPlex Manager in Step 17.

5. Do you intend to use the Remote Shared Memory Application Programming Interface (RSMAPI) or use SCI-PCI adapters for the interconnect transport?

- If no, go to Step 6.
- If yes, install additional packages from the `Packages` directory of the Sun Cluster 3.1 CD-ROM. These packages are not automatically installed by `scinstall(1M)`.

```
# cd /cdrom/suncluster_3_1/SunCluster_3.1/Sol_ver/Packages
# pkgadd -d . packages
```

Note – In the CD-ROM path, replace *ver* with 8 (for Solaris 8) or 9 (for Solaris 9)

The following table lists the Sun Cluster 3.1 packages required to support the RSMAPI or SCI-PCI adapters and the order in which you must install them.

TABLE 2-7 Sun Cluster 3.1 Packages to Support the RSMAPI and SCI-PCI Adapters

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

6. Are there any patches required to support Sun Cluster or Solstice DiskSuite software?

- If yes, go to Step 7.
- If no, go to Step 9.

7. Do you intend to use SunPlex Manager to install patches?

- If yes, go to Step 8.
- If no, manually install all patches required to support Sun Cluster or Solstice DiskSuite software *before* you use SunPlex Manager, then go to Step 9.

8. Copy patches required for Sun Cluster or Solstice DiskSuite software into a single directory on a file system that is available to each node.

- a. Ensure that only one version of each patch is present in this patch directory.**
If the patch directory contains multiple versions of the same patch, SunPlex Manager cannot determine the correct patch dependency order.
- b. Ensure that the patches are uncompressed.**

c. Record the path to the patch directory.

You will provide this information to SunPlex Manager in Step 17.

9. Have available the following completed configuration planning worksheets:

- “Cluster and Node Names Worksheet” in *Sun Cluster 3.1 Release Notes*
- “Cluster Interconnect Worksheet” in *Sun Cluster 3.1 Release Notes*
- “Network Resources Worksheet” in *Sun Cluster 3.1 Release Notes*

See Chapter 1 and the *Sun Cluster 3.1 Data Service Planning and Administration Guide* for planning guidelines.

10. Start SunPlex Manager.

a. From the administrative console or any other machine outside the cluster, launch a browser.

b. Disable the browser’s Web proxy.

SunPlex Manager 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 displays in the browser window.

Note – If SunPlex Manager displays the administration interface instead of the Sun Cluster Installation screen, Sun Cluster software is already installed 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.

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

- The Solaris End User Software Group or higher is installed.
- Root disk partitions include a 512-MByte slice with the mount point `/globaldevices`.
- Root disk partitions include a 20-MByte slice with the mount point `/sds`.
- File system paths to all needed CD-ROM images and patches are set up, as described in Step 4 through Step 8.

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

12. Type a name for the cluster and select the number of nodes in your cluster.

The default number of nodes displayed might be higher than the number of nodes you intend to install in your cluster. If this happens, select the correct number of nodes you will install. This might occur if other nodes that are ready to be installed by SunPlex Manager use the same public network as the nodes you intend to install.

Click Next to continue.

Tip – You can use the Back button to return to a previous screen and change your information. However, SunPlex Manager does not save the information you supplied in the later screens. When you click Next, you must again type or select your configuration information in those screens.

13. Type the name of each cluster node.

SunPlex Manager supplies as defaults the names of nodes it finds on the public network which are ready to be installed by SunPlex Manager. If you specify a larger number of nodes to install than exist on the network, SunPlex Manager supplies additional default names. These additional names follow the naming convention *phys-clustername-N*.

Note – SunPlex Manager might list nodes other than the ones you intend to install in your cluster. This would happen if those other nodes use the same public network as the nodes you are installing, and if those nodes are installed with SunPlex Manager software but not yet installed with Sun Cluster software. If SunPlex Manager supplies the name of a node that you do not want to install in your cluster, type over it with the correct node name.

Click Next to continue.

14. From the pull-down lists for each node, select the names of the two adapters used for the private interconnects.

Refer to your completed “Cluster Interconnect Worksheet” for the appropriate adapter names for each node.

Click Next to continue.

15. Choose whether to install Solstice DiskSuite software (Solaris 8) or to configure Solaris Volume Manager mirrored disksets (Solaris 9).

You must install Solstice DiskSuite software (Solaris 8) or configure Solaris Volume Manager mirrored disksets (Solaris 9) if you intend to install the Sun Cluster HA for NFS or Sun Cluster HA for Apache data service.



Caution – When SunPlex Manager installs Solstice DiskSuite software or configures Solaris Volume Manager software, any data on all shared disks will be lost.

Click Next to continue.

16. Choose whether to install Sun Cluster HA for NFS, Sun Cluster HA for Apache, or both.

Refer to your completed “Network Resources” worksheet for the appropriate logical hostname or shared address.

- For Sun Cluster HA for NFS, also specify the logical hostname the data service will use.
- For Sun Cluster HA for Apache, also specify the shared address the data service will use.

Click Next to continue.

17. Type the path for each CD-ROM image needed to install the packages you specified, and optionally the path for the patch directory.

- Type each path in the appropriate path field for each software package, as shown in Table 2-8.
- Each specified path for a CD-ROM image must be the directory which contains the `.cdtoc` file for the CD-ROM.
- For any software package you do not install, leave the relevant path field blank.
- If you have already installed the required patches, leave the Patch Directory Path field blank.

TABLE 2-8 CD-ROM Image Path Fields for Software Packages

Software Package to Install	Name of CD-ROM Image Path Field
Solstice DiskSuite	Solaris CD-ROM Path
Sun Cluster	Sun Cluster 3.1 CD-ROM Path
Sun Cluster HA for NFS, Sun Cluster HA for Apache	Sun Cluster 3.1 Agents CD-ROM Path
Sun Cluster patches, Solstice DiskSuite patches	Patch Directory Path

Click Next to continue.

18. Is the information you supplied correct as displayed in the Confirm Information screen?

- If yes, go to Step 19.

- If no, perform the following steps to correct the configuration information.
- a. **Click Back until you return to the screen with the information to change.**

Note – When you click Back to back up to a previous screen, any information you typed in the subsequent screens is lost.

- b. **Type the correct information and click Next.**
 - c. **Retype or reselect the information in each screen until you return to the Confirm Information screen.**
 - d. **Verify that the information in the Confirm Information screen is now correct.**
19. **Click Begin Installation to start the installation process.**

Note – Do *not* close the browser window or change the URL during the installation process.

- a. **If the browser displays a New Site Certification window, follow the onscreen instructions to accept the certificate.**
- b. **If the browser prompts for login information, type the appropriate user ID and password for the node you connect to.**

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

SunPlex Manager installation output is logged in the `/var/cluster/spm` directory. Sun Cluster installation output is logged in the `/var/cluster/logs/install/scinstall.log.pid` file, where *pid* is the process ID number of the `scinstall` instance.

If during installation the `sccheck` utility detects any problems, it displays information about the problems found and prompts you for your next action.

- Click Cancel to stop installation, correct the problem detected by `sccheck`, then return to Step d and restart SunPlex Manager.
- Alternately, if you can correct the problem without quitting SunPlex Manager, correct the problem and click Retry Validation to run `sccheck` again.
- If you know that you can safely ignore the `sccheck` messages, for example, from notification in the *Sun Cluster 3.1 Release Notes* that a recent qualification will trigger a harmless error message, click Continue Installation.

20. **Use SunPlex Manager to verify quorum assignments and modify them, if necessary.**

For clusters with three or more nodes, using shared quorum devices is optional. SunPlex Manager 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 reassign quorum votes in the cluster.

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

Go to “How to Configure the Name Service Switch” on page 87.

▼ 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 and Sun Cluster software on all cluster nodes in a single operation and establish 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 connections are verified before you install Solaris software.

See the *Sun Cluster 3.x Hardware Administration Manual* 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 35 for requirements and guidelines.

3. Have available the following information.

- The Ethernet address of each cluster node
- The following completed configuration planning worksheets:
 - “Local File Systems With Mirrored Root” in *Sun Cluster 3.1 Release Notes* or “Local File Systems with Non-Mirrored Root” in *Sun Cluster 3.1 Release Notes*
 - “Cluster and Node Names Worksheet” in *Sun Cluster 3.1 Release Notes*
 - “Cluster Interconnect Worksheet” in *Sun Cluster 3.1 Release Notes*

See “Planning the Solaris Operating Environment” on page 13 and “Planning the Sun Cluster Environment” on page 17 for planning guidelines.

4. Are you using a naming service?

- If no, go to Step 5. You will set up the necessary hostname information in Step 31.
- If yes, add address-to-name mappings for all public hostnames and logical addresses, as well as the IP address and hostname of the JumpStart server, to any naming services (such as NIS or DNS) used by clients for access to cluster services. See “IP Addresses” on page 17 for planning guidelines. See your

Solaris system administrator documentation for information about using Solaris naming services.

5. Are you installing a new node to an existing cluster?

- If yes, run `scsetup(1M)` from another, active cluster node to add the new node's name to the list of authorized cluster nodes. See "How to Add a Cluster Node to the Authorized Node List" in "Adding and Removing a Cluster Node" in *Sun Cluster 3.1 System Administration Guide*.
- If no, go to Step 6.

6. As superuser, set up the JumpStart install server for Solaris operating environment installation.

See the `setup_install_server(1M)` and `add_install_client(1M)` man pages and "Preparing Custom JumpStart Installations" in *Solaris 8 Advanced Installation Guide* or "Preparing Custom JumpStart Installations (Tasks)" in *Solaris 9 Installation Guide* for instructions on how to set up a JumpStart install server.

When you set up the install server, ensure that the following requirements are met.

- The install server is on the same subnet as the cluster nodes, but is not itself a cluster node.
 - The install server installs the release of the Solaris operating environment required by the Sun Cluster software.
 - A custom JumpStart directory exists for JumpStart installation of Sun Cluster. This `jumpstart-dir` directory must contain a copy of the `check(1M)` utility and be NFS exported for reading by the JumpStart install server.
 - Each new cluster node is configured as a custom JumpStart install client that uses the custom JumpStart directory set up for Sun Cluster installation.
- 7. Create a directory on the JumpStart install server to hold your copy of the Sun Cluster 3.1 CD-ROM, if one does not already exist.**

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

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

8. Copy the Sun Cluster CD-ROM to the JumpStart install server.

a. Insert the Sun Cluster 3.1 CD-ROM into the CD-ROM drive on the JumpStart install server.

If the Volume Management daemon `vold(1M)` is running and configured to manage CD-ROM devices, it automatically mounts the CD-ROM on the `/cdrom/suncluster_3_1` directory.

b. Change to the

`/cdrom/suncluster_3_1/SunCluster_3.1/Sol_ver/Tools` directory, where `ver` is 8 (for Solaris 8) or 9 (for Solaris 9).

The following example uses the path to the Solaris 8 version of Sun Cluster software.

```
# cd /cdrom/suncluster_3_1/SunCluster_3.1/Sol_8/Tools
```

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

The `scinstall` command creates the new installation directory as it copies the CD-ROM files. The installation directory name `/export/suncluster/sc31` is used here as an example.

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

d. Eject the CD-ROM.

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

e. Ensure that the Sun Cluster 3.1 CD-ROM image on the JumpStart install server is NFS exported for reading by the JumpStart install 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* and the `share(1M)` and `dfstab(4)` man pages for more information about automatic file sharing.

9. Are you installing a new node to an existing cluster?

- If yes, go to Step 10.
- If no, skip to Step 11.

10. Have you added the node to the cluster’s authorized-node list?

- If yes, go to Step 11.
- If no, run `scsetup(1M)` from any existing cluster node to add the new node’s name to the list of authorized cluster nodes. See “How to Add a Cluster Node to the Authorized Node List” in “Adding and Removing a Cluster Node” in *Sun Cluster 3.1 System Administration Guide*.

11. From the JumpStart install server, start the `scinstall(1M)` utility.

The path `/export/suncluster/sc31` is used here as an example of the installation directory you created.

```
# cd /export/suncluster/sc30/SunCluster_3.1/Sol_ver/Tools  
# ./scinstall
```

Note – In the CD-ROM path, replace `ver` with 8 (for Solaris 8) or 9 (for Solaris 9)

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

- Interactive `scinstall` enables you to type ahead. Therefore, do not press Return 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. If you press Control-D to abort the session after Sun Cluster software is installed, `scinstall` asks you whether you want it to de-install those packages.
- Your session answers are stored as defaults for the next time you run this menu option. Default answers display between brackets ([]) at the end of the prompt.

12. From the Main Menu, type 3 (Configure a cluster to be JumpStarted from this install server).

This option is used to configure customer 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) Establish a new cluster using this machine as the first node
2) Add this machine as a node in an established cluster
* 3) Configure a cluster to be JumpStarted from this install server
4) Add support for new data services to this cluster node
5) Print release information for this cluster node

* ?) Help with menu options
* q) Quit

```

```
Option: 3
```

```
*** Custom JumpStart ***
```

```
...
```

```
Do you want to continue (yes/no) [yes]?
```

---

**Note** – If option 3 does not have an asterisk in front, the option is disabled because JumpStart setup is not complete or has an error. Exit the `scinstall` utility, repeat Step 6 through Step 8 to correct JumpStart setup, then restart the `scinstall` utility.

---

## 13. Specify the JumpStart directory name.

```
>>> Custom JumpStart Directory <<<
....
What is your JumpStart directory name? jumpstart-dir
```

## 14. Specify the name of the cluster.

```
>>> Cluster Name <<<
...
What is the name of the cluster you want to establish? clustername
```

## 15. Specify the names of all cluster nodes.

```

>>> Cluster Nodes <<<
...
Please list the names of all cluster nodes planned for the initial
cluster configuration. You must enter at least two nodes. List one
node name per line. When finished, type Control-D:

Node name: node1
Node name: node2
Node name (Ctrl-D to finish): <Control-D>

This is the complete list of nodes:
...
Is it correct (yes/no) [yes]?

```

#### 16. Specify whether to use data encryption standard (DES) authentication.

By default, Sun Cluster software permits a node to connect to the cluster only if the node is physically connected to the private interconnect and if the node name was specified in Step 15. However, the node actually communicates with the sponsoring node over the public network, since the private interconnect is not yet fully configured. DES authentication provides an additional level of security at installation time by enabling the sponsoring node to more reliably authenticate nodes that attempt to contact it to update the cluster configuration.

If you choose to use DES authentication for additional security, you must configure all necessary encryption keys before any node can join the cluster. See the `keyserv(1M)` and `publickey(4)` man pages for details.

```

>>> Authenticating Requests to Add Nodes <<<
...
Do you need to use DES authentication (yes/no) [no]?

```

#### 17. Specify the private network address and netmask.

```

>>> Network Address for the Cluster Transport <<<
...
Is it okay to accept the default network address (yes/no) [yes]?
Is it okay to accept the default netmask (yes/no) [yes]?

```

---

**Note** – You cannot change the private network address after the cluster is successfully formed.

---

#### 18. Specify whether the cluster uses transport junctions.

- If this is a two-node cluster, specify whether you intend to use transport junctions.

```

>>> Point-to-Point Cables <<<
...
Does this two-node cluster use transport junctions (yes/no) [yes]?

```

---

**Tip** – You can specify that the cluster uses transport junctions, regardless of whether the nodes are directly connected to each other. If you specify that the cluster uses transport junctions, you can more easily add new nodes to the cluster in the future.

---

- If this cluster has three or more nodes, you must use transport junctions. Press Return to continue to the next screen.

```
>>> Point-to-Point Cables <<<
```

```
...
```

```
Since this is not a two-node cluster, you will be asked to configure
two transport junctions.
```

```
Hit ENTER to continue:
```

#### 19. Does this cluster use transport junctions?

- If yes, specify names for the transport junctions. You can use the default names `switchN` or create your own names.

```
>>> Cluster Transport Junctions <<<
```

```
...
```

```
What is the name of the first junction in the cluster [switch1]?
```

```
What is the name of the second junction in the cluster [switch2]?
```

- If no, go to Step 20.

#### 20. Specify the first cluster interconnect transport adapter of the first node.

```
>>> Cluster Transport Adapters and Cables <<<
```

```
...
```

```
For node "node1",
```

```
What is the name of the first cluster transport adapter? adapter
```

#### 21. Specify the connection endpoint of the first adapter.

- If the cluster does not use transport junctions, specify the name of the adapter on the second node to which this adapter connects.

```
...
```

```
Name of adapter on "node2" to which "adapter" is connected? adapter
```

- If the cluster uses transport junctions, specify the name of the first transport junction and its port.

```
...
```

```
For node "node1",
```

```
Name of the junction to which "adapter" is connected? switch
```

```
...
```

```
For node "node1",
```

```
Use the default port name for the "adapter" connection (yes/no) [yes]?
```

---

**Note** – If your configuration uses SCI adapters, do not accept the default when you are prompted for the adapter connection (the port name). Instead, provide the port name (0, 1, 2, or 3) found on the Dolphin switch itself, to which the node is *physically* cabled. The following example shows the prompts and responses for declining the default port name and specifying the Dolphin switch port name 0.

```
...
 Use the default port name for the "adapter" connection
(yes/no) [yes]? no
 What is the name of the port you want to use? 0
```

---

## 22. Specify the second cluster interconnect transport adapter of the first node.

```
...
For node "node1",
 What is the name of the second cluster transport adapter? adapter
```

## 23. Specify the connection endpoint of the second adapter.

- If the cluster does not use transport junctions, specify the name of the adapter on the second node to which this adapter connects.

```
...
 Name of adapter on "node2" to which "adapter" is connected? adapter
```

- If the cluster uses transport junctions, specify the name of the second transport junction and its port.

```
...
For node "node1",
 Name of the junction to which "adapter" is connected? switch
...
For node "node1",
 Use the default port name for the "adapter" connection (yes/no) [yes]?
```

---

**Note** – If your configuration uses SCI adapters, do not accept the default when you are prompted for the adapter connection (the port name). Instead, provide the port name (0, 1, 2, or 3) found on the Dolphin switch itself, to which the node is *physically* cabled. The following example shows the prompts and responses for declining the default port name and specifying the Dolphin switch port name 0.

```
...
 Use the default port name for the "adapter" connection
(yes/no) [yes]? no
 What is the name of the port you want to use? 0
```

---

## 24. Does this cluster use transport junctions?



- If yes, repeat Step 20 through Step 23 for each additional cluster node.
- If no, go to Step 25.

**25. Specify the global devices file system name for each cluster node.**

```
>>> Global Devices File System <<<
...
 The default is to use /globaldevices.

For node "node1",
 Is it okay to use this default (yes/no) [yes]?

For node "node2",
 Is it okay to use this default (yes/no) [yes]?
```

**26. Accept or decline the generated `scinstall` commands.**

The `scinstall` command generated from your input is displayed for confirmation.

```
>>> Confirmation <<<

 Your responses indicate the following options to scinstall:

For node "node1",
 scinstall -c jumpstart-dir -h node1 \
...
 Are these the options you want to use (yes/no) [yes]?

For node "node2",
 scinstall -c jumpstart-dir -h node2 \
...
 Are these the options you want to use (yes/no) [yes]?

 Do you want to continue with JumpStart set up (yes/no) [yes]?
```

If you do not accept the generated commands, the `scinstall` utility returns you to the Main Menu. From there you can rerun menu option 3 and provide different answers. Your previous answers display as the defaults.

**27. If necessary, make adjustments to the default `class` file, or profile, created by `scinstall`.**

The `scinstall` command creates the following `autosinstall.class` default class file in the `jumpstart-dir/autosinstall.d/3.1` directory.

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

The default `class` file installs the End User System Support software group (`SUNWCuser`) of Solaris software. If your configuration has additional Solaris software requirements, change the `class` file accordingly. See “Solaris Software Group Considerations” on page 13 for more information.

You can change the profile in one of the following ways.

- Edit the `autoscinstall.class` file directly. These changes are applied to all nodes in all clusters that use this custom JumpStart directory.
- Update the `rules` file to point to other profiles, then run the `check` utility to validate the `rules` file.

As long as the Solaris operating environment install profile meets minimum Sun Cluster file system allocation requirements, there are no restrictions on other changes to the install profile. See “System Disk Partitions” on page 14 for partitioning guidelines and requirements to support Sun Cluster 3.1 software. For more information about JumpStart profiles, see “Preparing Custom JumpStart Installations” in *Solaris 8 Advanced Installation Guide* or “Preparing Custom JumpStart Installations (Tasks)” in *Solaris 9 Installation Guide*.

#### 28. Do you intend to use the Remote Shared Memory Application Programming Interface (RSMAPI) or use SCI-PCI adapters for the interconnect transport?

- If yes **and** you install the End User System Support software group, add the following entries to the default `class` file as described in Step 27.

```
package SUNWrsm add
package SUNWrsmx add
package SUNWrsmo add
package SUNWrsmox add
```

In addition, you must create or modify a post-installation finish script at Step 33 to install the Sun Cluster packages to support the RSMAPI and SCI-PCI adapters. If you install a higher software group than End User System Support, the `SUNWrsm*` packages are installed with the Solaris software and do not need to be added to the `class` file.

- If no, go to Step 29.

#### 29. Do you intend to use SunPlex Manager?

- If yes **and** you install the End User System Support software group, add the following entries to the default `class` file as described in Step 27.

```
package SUNWapchr add
package SUNWapchu add
```

If you install a higher software group than End User System Support, the `SUNWrsm*` packages are installed with the Solaris software and do not need to be added to the `class` file.

- If no, go to Step 30.

#### 30. Set up Solaris patch directories.

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

Also place copies of any hardware-related patches that must be installed after Solaris software is installed into each of these directories.

**31. Set up files to contain the necessary hostname information locally on each node.**

- a. **On the JumpStart install server, create files 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. This could be the JumpStart install server or another machine.
- IP address and hostname of each node in the cluster.

**32. Do you intend to use the Remote Shared Memory Application Programming Interface (RSMAPI) or use SCI-PCI adapters for the interconnect transport?**

- If yes, follow instructions in Step 33 to set up a post-installation finish script to install the following additional packages. Install the appropriate packages from the `/cdrom/suncluster_3_1/SunCluster_3.1/Sol_ver/Packages` directory of the Sun Cluster 3.1 CD-ROM in the order given in the following table.

---

**Note** – In the CD-ROM path, replace *ver* with 8 (for Solaris 8) or 9 (for Solaris 9)

---

**TABLE 2-9** Sun Cluster 3.1 Packages to Support the RSMAPI and SCI-PCI Adapters

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

- If no, go to Step 33 if you intend to add your own post-installation finish script. Otherwise, skip to Step 34.

### 33. (Optional) Add your own post-installation finish script.

---

**Note** – If you intend to use the Remote Shared Memory Application Programming Interface (RSM API) or use SCI-PCI adapters for the interconnect transport, you must modify the finish script into install the Sun Cluster SUNWscriif software package. This package is not automatically installed by `scinstall`.

---

You can add your own finish script, which is run after the standard finish script installed by the `scinstall` command. See “Preparing Custom JumpStart Installations” in *Solaris 8 Advanced Installation Guide* or “Preparing Custom JumpStart Installations (Tasks)” in *Solaris 9 Installation Guide* for information about creating a JumpStart finish script.

- a. **Name your finish script `finish`.**
- b. **Copy your finish script to the `jumpstart-dir/autoscinstall.d/nodes/node` directory, one directory for each node in the cluster.**  
Alternately, use this naming convention to create symbolic links to a shared finish script.

### 34. If you use an administrative console, display a console screen for each node in the cluster.

If `cconsole(1M)` is installed and configured on your administrative console, you can use it to display the individual console screens. Otherwise, you must connect to the consoles of each node individually.

### 35. From the `ok PROM` prompt on the console of each node, type the `boot net - install` command to begin the network JumpStart installation of each node.

```
ok boot net - install
```

---

**Note** – The dash (-) in the command must be surrounded by a space on each side.

---

Sun Cluster installation output is logged in the `/var/cluster/logs/install/scinstall.log.pid` file, where `pid` is the process ID number of the `scinstall` instance.

---

**Note** – Unless you have installed your own `/etc/inet/ntp.conf` file, the `scinstall` command installs a default `ntp.conf` file for you. Because the default file is shipped with references to eight nodes, 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 100 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.

---

**Note** – The Solaris interface groups feature is disabled by default during Solaris software installation. Interface groups are not supported in a Sun Cluster configuration and should not be reenabled. See the `ifconfig(1M)` man page for more information about Solaris interface groups.

---

### 36. Are you installing a new node to an existing cluster?

- If no, go to Step 37.
- If yes, create mount points on the new node for all existing cluster file systems.

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

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

b. On the node 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 returned by the `mount` command is `/global/dg-schost-1`, run `mkdir -p /global/dg-schost-1` on the node being added to the cluster.

---

**Note** – The mount points become active after you reboot the cluster in Step 39.

---

c. Is VERITAS Volume Manager (VxVM) installed on any nodes that are already in the cluster?

- If yes, ensure that the same `vxio` number is used on the VxVM-installed nodes and 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 by changing the `/etc/name_to_major` entry to use a different number.

- If no, go to Step 37.

### 37. Install any Sun Cluster software patches.

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

### 38. Do you intend to use dynamic reconfiguration on Sun Enterprise 10000 servers?

- If yes, on each node add the following entry to the `/etc/system` file.

```
set kernel_cage_enable=1
```

This entry becomes effective after the next system reboot. See the *Sun Cluster 3.1 System Administration Guide* for procedures to perform dynamic reconfiguration tasks in a Sun Cluster configuration. See your server documentation for more information about dynamic reconfiguration.

- If no, go to Step 39.

### 39. Did you add a new node to an existing cluster, or install Sun Cluster software patches that require you to reboot the entire cluster, or both?

- If no, reboot the individual node if any patches you installed require a node reboot or if any other changes you made require a reboot to become active.
- If yes, perform a reconfiguration reboot as instructed in the following steps.

#### 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.

---

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

```
ok boot
```

Until cluster install mode is disabled, only the first-installed node, which established the cluster, has a quorum vote. In an established cluster that is still in install mode, if the cluster is not shut down before the first-installed node is rebooted, the remaining cluster nodes cannot obtain quorum and the entire cluster shuts down. Cluster nodes remain in install mode until the first time you run the `scsetup(1M)` command, during the procedure “How to Perform Post-Installation Setup” on page 90.

### 40. Set up the name service look-up order.

Go to “How to Configure the Name Service Switch” on page 87.

## ▼ 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 lookup order of the following entries.**
  - For the `hosts` and `netmasks` database entries, place `files` after `cluster`.  
For Sun Cluster HA for NFS, insert `[SUCCESS=return]` after `files` and before name services. This ensures that, if the node succeeds in resolving a name locally, it does not contact the listed name service(s) but instead returns success immediately.

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

- For all other database entries, place `files` first in 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 it in, in which case it is ignored.

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 88.

## ▼ How to Set Up the Root Environment

Perform these tasks on each node in the cluster.

---

**Note** – In a Sun Cluster configuration, user initialization files for the various shells must verify that they are run from an interactive shell 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.

---

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

  - a. **Set the `PATH` to include `/usr/sbin` and `/usr/cluster/bin`.**
    - For VERITAS Volume Manager (VxVM), also set your `PATH` to include `/etc/vx/bin`.
    - If you will install the `VRTSvmsa` package, also add `/opt/VRTSvmsa/bin` to your `PATH` for VxVM 3.2 or earlier, or add `/opt/VRTSob/bin` to your `PATH` for VxVM 3.5.
    - For VERITAS File System (VxFS), also set your `PATH` to include `/opt/VRTSvxfs/sbin`, `/usr/lib/fs/vxfs/bin`, and `/etc/fs/vxfs`.
  - b. **Set the setting the `MANPATH` to include `/usr/cluster/man`. Also include the volume manager-specific paths:**
    - For Solstice DiskSuite/Solaris Volume Manager software, set your `MANPATH` to include `/usr/share/man`.
    - For VxVM, set your `MANPATH` to include the following path.
      - For VxVM 3.1 and earlier, use `/opt/VRTSvxvm/man`.
      - For VxVM 3.1.1 and later, use `/opt/VRTS/man`.
    - For VxVM 3.2 or earlier, if you will install the `VRTSvmsa` package, also add `/opt/VRTSvmsa/man` to your `MANPATH`.
    - For VxFS, also set your `MANPATH` to include `/opt/VRTS/man`.

3. **(Optional) For ease of administration, set the same root password on each node, if you have not already done so.**
4. **Repeat Step 1 through Step 3 on each remaining cluster node.**



5. **Install data service software packages.**

Go to “How to Install Data Service Software Packages” on page 89.

## ▼ How to Install Data Service Software Packages

Perform this task on each cluster node.

---

**Note** – If you used SunPlex Manager to install Sun Cluster HA for NFS or Sun Cluster HA for Apache, or both, and you do not intend to install any other data services, you do not need to perform this procedure. Instead, go to “How to Perform Post-Installation Setup” on page 90.

---

1. **Become superuser on a cluster node.**

2. **If you install from the CD-ROM, insert the Sun Cluster 3.1 Agents CD-ROM into the CD-ROM drive on the node.**

3. **Start the `scinstall(1M)` utility.**

```
scinstall
```

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

- Interactive `scinstall` enables you to type ahead. Therefore, do not press Return 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.

4. **To add data services, type 4 (Add support for a new data service to this cluster node).**

Follow the prompts to select all data services to install.

---

**Note** – You must install the same set of data service packages on each node, even if a node is not expected to host resources for an installed data service.

---

5. **If you installed from a CD-ROM, eject the CD-ROM.**

6. **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, first shut down the cluster by using the `scshutdown(1M)`

command, then reboot the each node in the cluster.

---

**Note** – Until cluster install mode is disabled, only the first-installed node, which established the cluster, has a quorum vote. In an established cluster which is still in install mode, if the cluster is not shut down before the first-installed node is rebooted, the remaining cluster nodes cannot obtain quorum and the entire cluster shuts down. Cluster nodes remain in install mode until the first time you run the `scsetup(1M)` command, during the procedure “How to Perform Post-Installation Setup” on page 90.

---

7. Repeat Step 1 through Step 6 on each remaining cluster node.
8. Perform post-installation setup and assign quorum votes.  
Go to “How to Perform Post-Installation Setup” on page 90.

## ▼ How to Perform Post-Installation Setup

Perform this procedure one time only, after the cluster is fully formed.

Verify that all nodes have joined the cluster.

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. Determine the global device ID (DID) of each shared disk you will configure as a quorum device.**

Use the `scdidadm` output from Step 2 to identify the DID name of each shared disk you will configure as a quorum device. For example, the output in the previous substep shows that global device `d2` is shared by `phys-schost-1` and `phys-schost-2`. You will use this information in Step 8. See “Quorum Devices” on page 22 for further information about planning quorum devices.

**4. Are you adding a new node to an existing cluster?**

- If yes, you might need to update the quorum configuration to accommodate your cluster’s new configuration. See the *Sun Cluster 3.1 Concepts Guide* for information about quorum. To change the quorum configuration, follow procedures in “Administering Quorum” in *Sun Cluster 3.1 System Administration Guide*. When the quorum configuration is satisfactory, go to Step 12.
- If no, go to Step 6.

**5. Did you use SunPlex Manager to install Sun Cluster software?**

- If yes, skip to Step 11. During Sun Cluster installation, SunPlex Manager assigns quorum votes and removes the cluster from install mode for you.
- If no, go to Step 6.

**6. Become superuser on one node of the cluster.**

**7. 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 11.

---

If the quorum setup process is interrupted or fails to complete successfully, rerun `scsetup`.

**8. At the prompt `Do you want to add any quorum disks?`, configure at least one shared quorum device if your cluster is a two-node cluster.**

A two-node cluster remains in install mode until a shared quorum device is configured. After the `scsetup` utility configures the quorum device, the message `Command completed successfully` is displayed. If your cluster has three or more nodes, quorum device configuration is optional.

**9. At the prompt `Is it okay to reset "installmode"?`, answer `Yes`.**

After the `scsetup` utility sets quorum configurations and vote counts for the cluster, the message `Cluster initialization is complete` is displayed and the utility returns you to the Main Menu.

**10. From any node, verify the device and node quorum configurations.**

```
% scstat -q
```

**11. From any node, verify that cluster install mode is disabled.**

You do not need to be superuser to run this command.

```
% scconf -p | grep "Cluster install mode:"
Cluster install mode: disabled
```

**12. Install volume management software.**

- To install Solstice DiskSuite software or configure Solaris Volume Manager software, go to “Installing and Configuring Solstice DiskSuite/Solaris Volume Manager Software” on page 126.
- To install VERITAS Volume Manager software, go to “Installing and Configuring VxVM Software” on page 159.

## ▼ 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, the transport adapters.

---

**Note** – If the node has already joined the cluster and is no longer in install mode (see Step 11 of “How to Perform Post-Installation Setup” on page 90), 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 3.1 System Administration Guide*.

---

**1. Attempt to reinstall the node.**

Certain failed installations can be corrected simply by repeating the 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 you will uninstall.**

**3. From the active cluster member, add the node you intend to uninstall to the cluster’s node authentication list.**

```
/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 3.1 System Administration Guide* for procedures.

**4. Become superuser on the node you intend to uninstall.**

**5. Reboot the node into non-cluster mode.**

```
shutdown -g0 -y -i0
ok boot -x
```

**6. Uninstall the node.**

Run the `scinstall` command from a directory that does not contain any files delivered by the Sun Cluster packages.

```
cd /
/usr/cluster/bin/scinstall -r
```

See the `scinstall(1M)` man page for more information.

**7. Reinstall 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 them.

---

## 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 34
- Volume manager installation and configuration as described in “Installing and Configuring Solstice DiskSuite/Solaris Volume Manager Software” on page 126 or “Installing and Configuring VxVM Software” on page 159

**TABLE 2–10** Task Map: Configuring the Cluster

| Task                                         | For Instructions, Go To ...                                                      |
|----------------------------------------------|----------------------------------------------------------------------------------|
| Create and mount cluster file systems.       | “How to Add Cluster File Systems” on page 94                                     |
| Configure IP Network Multipathing groups.    | “How to Configure Internet Protocol (IP) Network Multipathing Groups” on page 98 |
| (Optional) Change a node’s private hostname. | “How to Change Private Hostnames” on page 99                                     |

**TABLE 2-10** Task Map: Configuring the Cluster (Continued)

| Task                                                                                                 | For Instructions, Go To ...                                                                                                                                                                                                    |
|------------------------------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Create or modify the NTP configuration file.                                                         | "How to Configure Network Time Protocol (NTP)" on page 100                                                                                                                                                                     |
| (Optional) Install the Sun Cluster module to Sun Management Center software.                         | "Installing the Sun Cluster Module for Sun Management Center" on page 102<br>Sun Management Center documentation                                                                                                               |
| Install third-party applications and configure the applications, data services, and resource groups. | <i>Sun Cluster 3.1 Data Service Planning and Administration Guide</i><br>"Sun Cluster Installation and Configuration Worksheets" in <i>Sun Cluster 3.1 Data Service Release Notes</i><br>Third-party application documentation |

## ▼ How to Add Cluster File Systems

Perform this procedure for each cluster file system you add.



**Caution** – Any data on the disks is destroyed when you create a file system. Be sure you specify the correct disk device name. If you specify the wrong device name, you will erase data that you might not intend to delete.

If you used SunPlex Manager to install data services, one or more cluster file systems already exist if there were sufficient shared disks on which to create the cluster file systems.

### 1. Ensure that volume manager software is installed and configured.

For volume manager installation procedures, see "Installing and Configuring Solstice DiskSuite/Solaris Volume Manager Software" on page 126 or "Installing and Configuring VxVM Software" on page 159.

### 2. Do you intend to install VERITAS File System (VxFS) software?

- If no, go to Step 3.
- If yes, perform the following steps.
  - a. Follow the procedures in your VxFS installation documentation to install VxFS software on each node of the cluster.
  - b. In the `/etc/system` file on each node, change the setting value for the following entry from `0x4000` to `0x6000`.

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

Sun Cluster software requires a minimum default stack size setting of 0x6000. Because VxFS installation changes this setting to 0x4000, you must manually change it back to 0x6000 after VxFS installation is complete.

### 3. Become superuser on any node in the cluster.

---

**Tip** – For faster file system creation, become superuser on the current primary of the global device you create a file system for.

---

### 4. Create a file system.

- For a VxFS file system, follow procedures provided in your VxFS documentation.
- 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.

**TABLE 2-11** Examples of Raw Disk Device Names

| Volume Manager                            | Sample Disk Device Name   | Description                                       |
|-------------------------------------------|---------------------------|---------------------------------------------------|
| Solstice DiskSuite/Solaris Volume Manager | /dev/md/oracle/rdisk/d1   | Raw disk device d1 within the oracle diskset      |
| VERITAS Volume Manager                    | /dev/vx/rdisk/oradg/vol01 | Raw disk device vol01 within the oradg disk group |
| None                                      | /dev/global/rdisk/d1s3    | Raw disk device d1s3                              |

### 5. 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 will not be 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

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

a. Use the following required mount options.

---

**Note** – Logging is required for all cluster file systems.

---

- **Solaris UFS logging** – Use the `global`, `logging` mount options. See the `mount_ufs(1M)` man page for more information about UFS mount options.

---

**Note** – The `syncdir` mount option is not required for UFS cluster file systems.

- If you specify `syncdir`, you are guaranteed POSIX-compliant file system behavior for the `write()` system call, in the sense that if a `write()` succeeds you are guaranteed that there is space on disk.
- If you do not specify `syncdir`, you will have the same behavior that is seen with UFS file systems. When you do not specify `syncdir`, performance of writes that allocate disk blocks, such as when appending data to a file, can significantly improve. However, in some cases, without `syncdir` you would not discover an out-of-space condition (`ENOSPC`) until you close a file. The case in which you will see `ENOSPC` on close is only during a very short time period after a failover. With `syncdir` (and POSIX behavior), the out-of-space condition would be discovered before the close.

- 
- **Solstice DiskSuite trans metadvice or Solaris Volume Manager transactional volume** – Use the `global` mount option only. Do **not** use the `logging` mount option.

---

**Note** – Transactional volumes are scheduled to be removed from the Solaris operating environment in an upcoming Solaris release. Solaris UFS logging, available since the Solaris 8 release, provides the same capabilities but superior performance, as well as lower system administration requirements and overhead.

---

See your Solstice DiskSuite documentation for information about setting up trans metadvice or see your Solaris Volume Manager documentation for information about setting up transactional volumes.

- **VxFS logging** – Use the `global`, `log` mount options. See the `VxFS mount_vxfs` man page 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`.

See the `vfstab(4)` man page for details.

- 7. On any node in the cluster, verify that mount points exist and `/etc/vfstab` file entries are correct on all nodes of the cluster.

```
sccheck
```

If no errors occur, nothing is returned.

- 8. From any node in the cluster, mount the cluster file system.

```
mount /global/device-group/mountpoint
```

---

**Note** – For VERITAS File System (VxFS), mount the 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.

---

- 9. 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.

---

**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.

---

- 10. Configure IP Network Multipathing groups.

Go to “How to Configure Internet Protocol (IP) Network Multipathing Groups” on page 98.

## Example—Creating a Cluster File System

The following example creates a UFS cluster file system on the Solstice DiskSuite metadvice /dev/md/oracle/rdsk/d1.

```
newfs /dev/md/oracle/rdsk/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/rdsk/d1 /global/oracle/d1 ufs 2 yes global,logging
(save and exit)

(on one node)
sccheck
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.

---

**Note** – All public network adapters *must* belong to an IP Network Multipathing group.

---

1. **Have available your completed “Public Networks Worksheet” in *Sun Cluster 3.1 Release Notes*.**

2. **Configure IP Network Multipathing groups.**

Perform procedures for IPv4 addresses in “Deploying Network Multipathing” in *IP Network Multipathing Administration Guide* or “Administering Network Multipathing (Task)” in *System Administration Guide: IP Services* Also follow these guidelines for IP Network Multipathing groups in a Sun Cluster configuration:

- Each public network adapter must belong to a multipathing group.
- The `local-mac-address?` variable must have a value of `true` for Ethernet adapters. This is a change from the requirement for Sun Cluster 3.0 software.
- You must configure a test IP address for each multipathing group adapter.

- 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 they are not highly available.
  - There are no requirements or restrictions for the name of a multipathing group.
3. **Do you intend to change any private hostnames?**
    - If yes, go to “How to Change Private Hostnames” on page 99.
    - If no, go to Step 4.
  4. **Did you install your own `/etc/inet/ntp.conf` file before you installed Sun Cluster software?**
    - If no, go to “How to Configure Network Time Protocol (NTP)” on page 100 to install or create the NTP configuration file.
    - If yes, go to Step 5.
  5. **Do you intend to use Sun Management Center to configure resource groups or monitor the cluster?**
    - If yes, go to “Installing the Sun Cluster Module for Sun Management Center” on page 102.
    - If no, install third-party applications, register resource types, set up resource groups, and configure data services. See the documentation supplied with the application software and the *Sun Cluster 3.1 Data Service Planning and Administration Guide*.

## ▼ How to Change Private Hostnames

Perform this task if you do not want to use the default private hostnames (`clusternodenodeid-priv`) assigned during Sun Cluster software installation.

---

**Note** – Do *not* perform this procedure after applications and data services have been configured and started. Otherwise, an application or data service might continue to use the old private hostname after it is renamed, which would cause hostname conflicts. If any applications or data services are running, stop them before you perform this procedure.

---

1. **Become superuser on a node in the cluster.**
2. **Start the `scsetup(1M)` utility.**

```
scsetup
```

3. To work with private hostnames, type 5 (Private hostnames).

4. To change a private hostname, type 1 (Change a private hostname).

Follow the prompts to change the private hostname. Repeat for each private hostname to change.

5. Verify the new private hostnames.

```
scconf -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
```

6. Did you install your own `/etc/inet/ntp.conf` file before you installed Sun Cluster software?

- If yes, go to Step 7.
- If no, go to “How to Configure Network Time Protocol (NTP)” on page 100 to install or create the NTP configuration file.

7. Do you intend to use Sun Management Center to monitor the cluster?

- If yes, go to “Installing the Sun Cluster Module for Sun Management Center” on page 102.
- If no, install third-party applications, register resource types, set up resource groups, and configure data services. See the documentation supplied with the application software and the *Sun Cluster 3.1 Data Service Planning and Administration Guide*.

## ▼ 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 and when you change the private hostname of a node in the cluster.

---

**Note** – 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, as long as this basic requirement for synchronization is met.

See the *Sun Cluster 3.1 Concepts Guide* 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. **Did you install your own `/etc/inet/ntp.conf` file before you installed Sun Cluster software?**
  - If yes, you do not need to modify your `ntp.conf` file. Skip to Step 8.
  - If no, go to Step 2.
2. **Become superuser on a cluster node.**
3. **Do you have your own `/etc/inet/ntp.conf` file to install on the cluster nodes?**
  - If yes, copy your `/etc/inet/ntp.conf` file to **each** node of the cluster, then skip to Step 6.

---

**Note** – All cluster nodes must be synchronized to the same time.

---

- If no, go to Step 4 to edit the `/etc/inet/ntp.conf.cluster` file. Sun Cluster software creates this file as the NTP configuration file if an `/etc/inet/ntp.conf` file is not found during Sun Cluster installation. Do not rename the `ntp.conf.cluster` file as `ntp.conf`.
4. **On one node of the cluster, edit the private hostnames in the `/etc/inet/ntp.conf.cluster` file.**

If `/etc/inet/ntp.conf.cluster` does not exist on the node, you might have an `/etc/inet/ntp.conf` file from an earlier installation of Sun Cluster software. If so, perform the following edits on that `ntp.conf` file.

    - a. **Ensure that an entry exists for the private hostname of each cluster node.**
    - b. **Remove any unused private hostnames.**

If the `ntp.conf.cluster` file contains nonexistent private hostnames, when a node is rebooted the system will generate error messages when the node attempts to contact those nonexistent private hostnames.
    - c. **If you changed a node's private hostname, ensure that the NTP configuration file contains the new private hostname.**
    - 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.conf.cluster` 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.**

- For `ntp.conf.cluster`, 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 `ntp.conf` does not exist but `ntp.conf.cluster` does exist, the NTP daemon is started using `ntp.conf.cluster` as the NTP configuration file.

- For `ntp.conf`, run the following command.

```
/etc/init.d/xntpd start
```

**8. Do you intend to use Sun Management Center to monitor the cluster?**

- If yes, go to “Installing the Sun Cluster Module for Sun Management Center” on page 102.
- If no, install third-party applications, register resource types, set up resource groups, and configure data services. See the documentation supplied with the application software and the *Sun Cluster 3.1 Data Service Planning and Administration Guide*.

---

## Installing the Sun Cluster Module for Sun Management Center

The following table lists the tasks to perform to install the Sun Cluster module software for Sun Management Center.

**TABLE 2-12** Task Map: Installing the Sun Cluster Module for Sun Management Center

| Task                                                                            | For Instructions, Go To...                                                                                |
|---------------------------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------|
| Install Sun Management Center server, help server, agent, and console packages. | Sun Management Center documentation<br>“Installation Requirements for Sun Cluster Monitoring” on page 103 |
| Install Sun Cluster module packages.                                            | “How to Install the Sun Cluster Module for Sun Management Center” on page 103                             |
| Start Sun Management Center server, console, and agent processes.               | “How to Start Sun Management Center” on page 104                                                          |

**TABLE 2-12** Task Map: Installing the Sun Cluster Module for Sun Management Center  
(Continued)

| Task                                                                | For Instructions, Go To...                                                           |
|---------------------------------------------------------------------|--------------------------------------------------------------------------------------|
| Add each cluster node as a Sun Management Center agent host object. | "How to Add a Cluster Node as a Sun Management Center Agent Host Object" on page 105 |
| Load the Sun Cluster module to begin to monitor the cluster.        | "How to Load the Sun Cluster Module" on page 106                                     |

## 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 non-cluster nodes. If you have an administrative console or other dedicated machine, you can realize improved performance if you run the console process on the administrative console and the server process on a separate machine. You must install the Sun Management Center agent package on each cluster node.

Follow procedures in the Sun Management Center documentation to install the Sun Management Center packages.

- **Simple Network Management Protocol (SNMP) port** – When you install Sun Management Center on the agent, 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 you choose for reference later when you configure the cluster for monitoring.

### ▼ 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) were 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 CD-ROM into the CD-ROM drive.**

**c. Change to the `/cdrom/suncluster_3_1/SunCluster_3.1/Sol_ver/Packages` directory, where *ver* is 8 (for Solaris 8) or 9 (for Solaris 9) .**

The following example uses the path to the Solaris 8 version of Sun Cluster software.

```
cd /cdrom/suncluster_3_1/SunCluster_3.1/Sol_8/Packages
```

**d. Install the Sun Cluster module server package.**

```
pkgadd -d . SUNWscssv
```

**e. Change to the CD-ROM root directory and eject the CD-ROM.**

**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 the *Sun Cluster 3.1 Release Notes* for the location of patches and installation instructions.

**5. Start Sun Management Center.**

Go to “How to Start Sun Management Center” on page 104.

## ▼ 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 and click Login.

6. Add cluster nodes as monitored host objects.

Go to “How to Add a Cluster Node as a Sun Management Center Agent Host Object” on page 105.

## ▼ 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 will contain the Sun Management Center agent host object 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 one.

See your Sun Management Center documentation for information about how to create Sun Management Center domains.

2. **Select Edit>Create an Object from the pull-down menu.**
3. **Select 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 you chose when you installed the Sun Management Center agent.**
7. **Click OK.**  
A Sun Management Center agent host object is created in the domain.
8. **Load the Sun Cluster module.**  
Go to "How to Load the Sun Cluster Module" on page 106.

## ▼ How to Load the Sun Cluster Module

Perform this procedure to start cluster monitoring.

1. **From the Sun Management Center main window, point to an icon of a cluster node, right-click to display the pull-down menu, and select Load Module.**  
The Load Module window lists each available Sun Management Center module and whether it is currently loaded.
2. **Select Sun Cluster: Not loaded, which is usually at the bottom of the list, and click OK.**  
The Module Loader window shows the current parameter information for the selected module.
3. **Click OK.**  
After a few moments the module is loaded and a Sun Cluster icon is displayed in the Details window.
4. **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.
5. **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, 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, 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 URL file:/opt/SUNWsymon/lib/locale/C/help/main.top.html.

---

**Note** – The Help button in the Sun Management Center browser accesses Sun Management Center online help, not 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.

Install third-party applications, register resource types, set up resource groups, and configure data services. See the documentation supplied with the application software and the *Sun Cluster 3.1 Data Service Planning and Administration Guide*.



## Upgrading Sun Cluster Software

---

This chapter provides step-by-step procedures to upgrade a Sun Cluster 3.0 configuration to Sun Cluster 3.1 software.

---

**Note** – Upgrade from Solaris 8 to Solaris 9 software is not supported in a Sun Cluster configuration.

---

The following step-by-step instructions are in this chapter.

- “How to Prepare the Cluster for Upgrade” on page 111
- “How to Upgrade the Solaris Operating Environment” on page 113
- “How to Upgrade Cluster Software” on page 116
- “How to Finish Upgrading Cluster Software” on page 119
- “How to Upgrade Sun Management Center Software” on page 122

For overview information about planning your Sun Cluster 3.1 configuration, see Chapter 1. For a high-level description of the related procedures for Sun Cluster software upgrade, see “Upgrading From Sun Cluster 3.0 to Sun Cluster 3.1 Software” on page 109.

---

## Upgrading From Sun Cluster 3.0 to Sun Cluster 3.1 Software

Perform the following tasks to upgrade your cluster from Sun Cluster 3.0 to Sun Cluster 3.1 software.

**TABLE 3-1** Task Map: Upgrading From Sun Cluster 3.0 to Sun Cluster 3.1 Software

| Task                                                                                                                                                           | For Instructions, Go To ...                                                          |
|----------------------------------------------------------------------------------------------------------------------------------------------------------------|--------------------------------------------------------------------------------------|
| Read the upgrade guidelines.                                                                                                                                   | "Overview of Upgrading From Sun Cluster 3.0 to Sun Cluster 3.1 Software" on page 110 |
| Take the cluster out of production, disable resources, and back up shared data and system disks.                                                               | "How to Prepare the Cluster for Upgrade" on page 111                                 |
| Upgrade the Solaris operating environment, if necessary. Optionally, upgrade VERITAS Volume Manager (VxVM).                                                    | "How to Upgrade the Solaris Operating Environment" on page 113                       |
| Upgrade to Sun Cluster 3.1 framework and data services software. If necessary, upgrade applications.                                                           | "How to Upgrade Cluster Software" on page 116                                        |
| Finish the cluster software upgrade. If you upgraded VxVM, upgrade disk groups. Reregister resource types, enable resources, and bring resource groups online. | "How to Finish Upgrading Cluster Software" on page 119                               |

## Overview of Upgrading From Sun Cluster 3.0 to Sun Cluster 3.1 Software

This section provides guidelines for upgrading from Sun Cluster 3.0 to Sun Cluster 3.1 software.

### Upgrade Guidelines

Follow these guidelines to upgrade from Sun Cluster 3.0 to Sun Cluster 3.1 software:

- Have available the test IP addresses your public network adapters will use. The `scinstall` upgrade utility will prompt you for a test IP address for each public network adapter in the cluster. See the *IP Network Multipathing Administration Guide* (Solaris 8) or *System Administration Guide: IP Services* (Solaris 9) for information about test IP addresses.
- You must upgrade all software to a version that is supported by Sun Cluster 3.1 software. For example, you must upgrade a data service agent that is supported on Sun Cluster 3.0 software to the version supported on Sun Cluster 3.1 software. If the related application is not also supported on Sun Cluster 3.1 software, you must also upgrade the application to a supported release.
- The `scinstall` upgrade utility only upgrades those data services that are provided with Sun Cluster 3.1 software. You must manually upgrade any custom or third-party data services.

- Upgrade from Solaris 8 to Solaris 9 software is not supported in a Sun Cluster configuration.
- Sun Cluster 3.1 software does not support direct upgrade from Sun Cluster 2.x software.
- Sun Cluster 3.1 software does not support downgrade back to Sun Cluster 3.0 software.

## ▼ How to Prepare the Cluster for Upgrade

Before you upgrade the software, perform the following steps to take the cluster out of production.

**1. 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 operating environment
- Sun Cluster 3.1 framework
- Sun Cluster 3.1 data services (agents)
- Applications managed by Sun Cluster 3.1 data-service agents
- VERITAS Volume Manager

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

**2. Have available your list of test IP addresses, one for each public network adapter in the cluster.**

A test IP address is required for each public network adapter in the cluster, regardless of whether the adapter is the active or the backup adapter in the group. The test IP addresses will be 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 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.

- 3. Notify users that cluster services will be unavailable during upgrade.**
- 4. 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 that node for unresolved error or warning messages.

**5. List the names of all resource types installed on your cluster before upgrade.**

```
% scrgadm -p | grep "Res Type"
```

Keep this list for reference when you reregister the resource types at the end of the upgrade process.

**6. Become superuser on a node of the cluster.**

**7. Disable all resources in the cluster.**

The disabling of resources during upgrade prevents the cluster from bringing the resources online automatically if a node is mistakenly rebooted into cluster mode.

**a. From any node, list all enabled resources in the cluster.**

```
scrgadm -pv | grep "Res enabled"
```

**b. Use the `scswitch` command to disable each enabled resource in the cluster.**

```
scswitch -n -j resource
-n Disables
-j resource Specifies the resource
```

---

**Note** – If you are upgrading from the Sun Cluster 3.0 5/02 release, you can use the `scsetup(1M)` utility instead of the command line. From the Main Menu, choose Resource Groups, then choose Enable/Disable Resources.

---

**8. Switch each resource group offline.**

```
scswitch -F -g resource-group
-F Switches a resource group offline
-g resource-group Specifies the name of the resource group to take offline
```

**9. Move each resource group into the unmanaged state.**

```
scswitch -u -g resource-group
-u Moves the specified resource group in the unmanaged state
-g resource-group Specifies the name of the resource group to move into the unmanaged state
```



10. Verify that all resources on all nodes are disabled and that all resource groups are in the unmanaged state.

You do not need to be superuser to run this command.

```
scstat -g
```

11. Stop all databases that are running on each node of the cluster.

12. Back up all shared data.

13. From one node, shut down the cluster.

```
scshutdown
ok
```

14. On each node, boot the node into non-cluster mode.

```
ok boot -x
```

15. Back up each system disk.

16. Do you intend to upgrade the Solaris software?

You must upgrade the Solaris operating environment if Sun Cluster 3.1 software does not support the update release of Solaris software that you currently run on your cluster.

---

**Note** – Upgrade from Solaris 8 to Solaris 9 software is not supported in a Sun Cluster configuration.

---

- If no, go to “How to Upgrade the Solaris Operating Environment” on page 113.
- If yes, go to “How to Upgrade Cluster Software” on page 116.

## ▼ How to Upgrade the Solaris Operating Environment

Perform this procedure on each node in the cluster to upgrade the Solaris operating environment to support Sun Cluster 3.1 software. However, if the cluster is already installed with a version of Solaris software supported by Sun Cluster 3.1 software and you do not intend to upgrade to a later compatible update of the Solaris software, you do not otherwise need to perform this procedure. Instead, go to “How to Upgrade Cluster Software” on page 116.

---

**Note** – The cluster must already run on, or be upgraded to, at least the minimum required level of the Solaris 8 or Solaris 9 operating environment to support Sun Cluster 3.1 software. See “Supported Products” in *Sun Cluster 3.1 Release Notes* for more information.

Upgrade from Solaris 8 to Solaris 9 software is not supported in a Sun Cluster configuration.

---

**1. Become superuser on the cluster node to upgrade.**

**2. Does your cluster use VxVM?**

- If no, go to Step 3.
- If yes, follow upgrade procedures in your VxVM installation documentation, then return to Step 3 when the VxVM procedure says to upgrade Solaris software.

---

**Note** – When instructed to boot into single-user mode, use the following command to reboot into single-user non-cluster mode:

```
boot -- -xs
```

---

**3. 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 contain an uppercase K or S in the file name, no further action is necessary concerning these links. If these links do not exist, or if these links exist but contain lowercase k or s in the file name, you will move aside these links in Step 6.

**4. Upgrade the Solaris software.**

---

**Note** – Sun Cluster software does not support upgrade from Solaris 8 to Solaris 9 software.

---

**a. Do you intend to use the Maintenance Update method to upgrade the Solaris software?**

- If yes, go to Step b.
- If no, make a record of all entries that are already commented out for later reference, then temporarily comment out all global device entries in the `/etc/vfstab` file.

Do this to prevent the Solaris upgrade from attempting to mount the global devices. This step is not necessary if you use the Maintenance Update method to upgrade Solaris software.

**b. Follow upgrade instructions for the version of Solaris to which you will upgrade your cluster.**

---

**Note** – To reboot a node during Solaris software upgrade, always add the `-x` option to the command. This ensures that the node reboots into non-cluster mode. The following two commands boot a node into single-user non-cluster mode:

```
reboot -- -xs
ok boot -xs
```

---

**c. Did you comment out the global device entries in the `/etc/vfstab` file in Step a?**

- If no, go to Step 6.
- If yes, uncomment all global device entries that you commented out in the `/a/etc/vfstab` file.

---

**Note** – Leave commented out all entries that were already commented out before you performed Step a.

---

**5. Did you perform VxVM procedures in Step 2 to upgrade Solaris software?**

- If no, go to Step 6.
- If yes, perform the remaining steps in the VxVM upgrade procedure you used, but do not reboot except into non-cluster mode.

```
reboot -- -x
```

**6. If the Apache links in Step 3 did not already exist or they contained a lowercase `k` or `s` in the file names before you upgraded the Solaris software, move aside the restored Apache links.**

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
```

**7. Reboot the node in non-cluster mode, if you have not already done so.**

Include the double dashes (--) in the command.

```
reboot -- -x
```

**8. Install any Solaris software patches and hardware-related patches, and download any needed firmware contained in the hardware patches. For Solstice DiskSuite software, install any Solstice DiskSuite software patches.**

---

**Note** – Do not reboot after adding patches. You will reboot the node 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.

**9. Upgrade to Sun Cluster 3.1 software.**

Go to “How to Upgrade Cluster Software” on page 116.

## ▼ How to Upgrade Cluster Software

Perform this procedure on each node. You can perform this procedure on more than one node at a time.

**1. Become superuser on a node of the cluster.**

**2. Insert the Sun Cluster 3.1 CD-ROM into the CD-ROM drive on the node.**

If the Volume Management daemon `vold(1M)` is running and configured to manage CD-ROM devices, it automatically mounts the CD-ROM on the `/cdrom/suncluster_3_1` directory.

**3. Upgrade the node to Sun Cluster 3.1 software.**

**a. Change to the**

`/cdrom/suncluster_3_1/SunCluster_3.1/Sol_ver/Tools` directory, where *ver* is 8 (for Solaris 8) or 9 (for Solaris 9).

The following example uses the path to the Solaris 8 version of Sun Cluster software.

```
cd /cdrom/suncluster_3_1/SunCluster_3.1/Sol_8/Tools
```

**b. Upgrade the cluster software framework.**

```
./scinstall -u update -S interact
```

`-S` Specifies the test IP addresses to use to convert NAFO groups to IP Network Multipathing groups

`interact` Specifies that `scinstall` will prompt the user for each test IP address needed

If upgrade processing is interrupted, ensure that the node is in non-cluster mode, then restart the `scinstall` command.

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.

During upgrade processing, `scinstall` makes the following configuration changes:

- Converts NAFO groups to IP Network Multipathing groups but keeps the original NAFO-group name
- Renames the `ntp.conf` file to `ntp.conf.cluster`, if `ntp.conf.cluster` does not already exist on the node
- Sets the `local-mac-address?` variable to `true`, if the variable is not already set to that value

**c. Change to the CD-ROM root directory and eject the CD-ROM.**

**d. Install any Sun Cluster patches.**

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

---

**Note** – Do not reboot the node at this time.

---

**4. Upgrade applications and apply application patches as needed.**

See your application documentation for installation instructions. Follow these guidelines to upgrade applications in a Sun Cluster 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 must reboot the node during application upgrade, add the `-x` option to the boot command. This ensures that the node reboots into non-cluster mode. The following two commands boot a node into single-user non-cluster mode:

```
reboot -- -xs
ok boot -xs
```

**5. Upgrade data services to the Sun Cluster 3.1 software versions.**

---

**Note** – Only those data services provided on the Sun Cluster 3.1 Agents CD-ROM are automatically upgraded by `scinstall(1M)`. You must manually upgrade any custom or third-party data services.

---

a. **Insert the Sun Cluster 3.1 Agents CD-ROM into the CD-ROM drive on the node to upgrade.**

b. **Upgrade the data service software.**

```
scinstall -u update -s all -d /cdrom/cdrom0
```

```
-u update Specifies upgrade
```

```
-s all Updates all Sun Cluster data services
```

If upgrade processing is interrupted, ensure that the node is in non-cluster mode, then restart the `scinstall` command.

c. **Eject the CD-ROM.**

d. **As needed, manually upgrade any custom data services that are not supplied on the Sun Cluster 3.1 Agents CD-ROM.**

e. **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.

f. **Repeat Step b through Step e on each node of the cluster.**

6. **After all nodes are upgraded, reboot each node into the cluster.**

```
reboot
```

7. **From any node, verify that all of the nodes are cluster members.**

```
scstat
```

```
-- Cluster Nodes --
```

|               | Node name     | Status |
|---------------|---------------|--------|
|               | -----         | -----  |
| Cluster node: | phys-schost-1 | Online |
| Cluster node: | phys-schost-2 | Online |

See the `scstat(1M)` man page for more information about displaying cluster status.

8. **Finish the upgrade.**

Go to “How to Finish Upgrading Cluster Software” on page 119.

## Example—Upgrading From Sun Cluster 3.0 to Sun Cluster 3.1 Software

The following example shows the process of upgrading a two-node cluster from Sun Cluster 3.0 to Sun Cluster 3.1 software on the Solaris 8 operating environment. The cluster node names are `phys-schost-1` and `phys-schost-2`.

*(On the first node, upgrade framework software from the Sun Cluster 3.1 CD-ROM)*

```
phys-schost-1# cd /cdrom/suncluster_3_1/SunCluster_3.1/Sol_8/Tools
phys-schost-1# ./scinstall -u update -S interact
```

*(On the first node, upgrade data services from the Sun Cluster 3.1 Agents CD-ROM)*

```
phys-schost-1# ./scinstall -u update -s all -d /cdrom/cdrom0
```

*(On the second node, upgrade framework software from the Sun Cluster 3.1 CD-ROM)*

```
phys-schost-2# cd /cdrom/suncluster_3_1/SunCluster_3.1/Sol_8/Tools
phys-schost-2# ./scinstall -u update -S interact
```

*(On the second node, upgrade data services from the Sun Cluster 3.1 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 cluster membership)*

```
scstat
```

## ▼ How to Finish Upgrading Cluster Software

Perform these steps after you have upgraded software on each node of the cluster and all nodes are booted into the cluster.

### 1. Did you upgrade VxVM?

- If no, go to Step 2.
- If yes, 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:

```
vxdg upgrade dgname
```

See your VxVM administration documentation for more information about upgrading disk groups.

### 2. Are you using Sun Management Center to monitor your Sun Cluster configuration?

- If no, skip to Step 3.

- If yes, install the Sun Cluster 3.1 module packages for Sun Management Center on the server machine and help server machine.
- a. **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.
  - b. **Become superuser on the Sun Management Center server machine**
  - c. **Insert the Sun Cluster 3.1 CD-ROM into the CD-ROM drive.**
  - d. **Change to the**  
`/cdrom/suncluster_3_1/SunCluster_3.1/Sol_ver/Packages`  
**directory, where *ver* is 8 (for Solaris 8) or 9 (for Solaris 9) .**  
The following example uses the path to the Solaris 8 version of Sun Cluster software.  

```
cd /cdrom/suncluster_3_1/SunCluster_3.1/Sol_8/Packages
```
  - e. **Install the Sun Cluster module server package SUNWscssv.**  

```
pkgadd -d . SUNWscssv
```
  - f. **Change to the CD-ROM root directory and eject the CD-ROM.**
  - g. **Become superuser on the Sun Management Center help server machine**
  - h. **Install the Sun Cluster module help server package SUNWscsh1.**  
Use the same procedure as in Step c through Step f.
3. **Reregister all resource types which received a new version from the upgrade, then re-enable resources and bring resource groups online.**
    - a. **From any node, start the `scsetup(1M)` utility.**  

```
scsetup
```
    - b. **To work with resource groups, type 2 (Resource groups).**
    - c. **To register resource types, type 4 (Resource type registration).**  
Type **yes** when prompted to continue.
    - d. **Type 1 (Register all resource types which are not yet registered).**  
The `scsetup` utility displays all resource types which are not registered. Type **yes** to continue to register these resource types.
    - e. **Type 8 (Change properties of a resource).**  
Type **yes** to continue.



- f. **Type 3 (Manage resource versioning).**  
Type **yes** to continue.
- g. **Type 1 (Show versioning status).**  
The `scsetup` utility displays what resources you can upgrade to new versions of the same resource type, and the state that the resource should be in before the upgrade can begin.  
Type **yes** to continue.
- h. **Type 4 (Re-version all eligible resources).**  
Type **yes** to continue when prompted.
- i. **Return to the Resource Group Menu.**
- j. **To enable resources, type 6 (Enable/Disable a resource).**  
Type **yes** to continue when prompted.
- k. **Select a resource to enable and follow the prompts. Repeat for each disabled resource.**
- l. **When all resources are re-enabled, type q to return to the Resource Group Menu.**
- m. **Type 5 (Online/Offline or Switchover a resource group).**  
Type **yes** to continue when prompted.
- n. **Follow the prompts to bring each resource group online.**
- o. **Exit the `scsetup` utility.**  
Type **q** to back out of each submenu, or press Ctrl-C.

---

**Note** – To upgrade future versions of resource types, see “Upgrading a Resource Type” in *Sun Cluster 3.1 Data Service Planning and Administration Guide*.

---

The cluster upgrade is complete. You can now return the cluster to production.

---

# Sun Management Center Software Upgrade

This section describes how to upgrade from Sun Management Center 2.1.1 to Sun Management Center 3.0 software or from Sun Management Center 2.1.1 to Sun Management Center 3.5 software on a Sun Cluster 3.1 configuration.

## ▼ How to Upgrade Sun Management Center Software

Perform this procedure to upgrade from Sun Management Center 2.1.1 to Sun Management Center 3.0 software or from Sun Management Center 2.1.1 to Sun Management Center 3.0 software on a Sun Cluster 3.1 configuration.

### 1. Have available the following items.

- **Sun Cluster 3.1 CD-ROM or the path to the CD-ROM image.** You will use the CD-ROM to reinstall 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, select File>Exit from the menu bar.
- 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 from the locations listed in Table 3-2.

You must remove all Sun Cluster module packages from all locations. Otherwise, the Sun Management Center software upgrade might fail because of package

dependency problems. After you upgrade Sun Management Center software, you will reinstall these packages in Step 5.

```
pkgrm module-package
```

**TABLE 3-2** Locations to Remove Sun Cluster Module Packages

| Location                                  | Package to Remove    |
|-------------------------------------------|----------------------|
| Each cluster node                         | SUNWscsam, SUNWscsa1 |
| Sun Management Center console machine     | SUNWscscn            |
| Sun Management Center server machine      | SUNWscssv            |
| Sun Management Center help server machine | SUNWscsh1            |

**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 to the locations listed in Table 3-3.**

For Sun Management Center software, you install the help server package SUNWscsh1 on the console machine as well as on the help server machine.

```
cd /cdrom/suncluster_3_1/SunCluster_3.1/Packages
pkgadd module-package
```

**TABLE 3-3** Locations to Install Sun Cluster Module Packages

| Location                                  | Package to Install   |
|-------------------------------------------|----------------------|
| Each cluster node                         | SUNWscsam, SUNWscsa1 |
| Sun Management Center console machine     | SUNWscsh1            |
| Sun Management Center server machine      | SUNWscssv            |
| Sun Management Center help server machine | SUNWscsh1            |

**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 on all involved machines.**

Follow procedures in “How to Start Sun Management Center” on page 104.

**8. Load the Sun Cluster module.**

Follow procedures in “How to Load the Sun Cluster Module” on page 106.

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,

from the console's Details window select Module>Unload Module.

# Installing and Configuring Solstice DiskSuite/Solaris Volume Manager Software

---

Install and configure your local and multihost disks for Solstice DiskSuite/Solaris Volume Manager software by using the procedures in this appendix, along with the planning information in “Planning Volume Management” on page 24. See your Solstice DiskSuite/Solaris Volume Manager documentation for additional details.

The following procedures are in this appendix.

- “How to Install Solstice DiskSuite Software” on page 129
- “How to Set the Number of Metadevice or Volume Names and Disksets” on page 130
- “How to Create State Database Replicas” on page 132
- “How to Mirror the Root (/) File System” on page 133
- “How to Mirror the Global Namespace” on page 137
- “How to Mirror File Systems Other than Root (/) That Cannot Be Unmounted” on page 140
- “How to Mirror File Systems That Can Be Unmounted” on page 144
- “How to Create a Diskset” on page 147
- “How to Add Drives to a Diskset” on page 150
- “How to Repartition Drives in a Diskset” on page 151
- “How to Create an `md.tab` File” on page 152
- “How to Activate Metadevices or Volumes” on page 154
- “How to Add Mediator Hosts” on page 156
- “How to Check the Status of Mediator Data” on page 156
- “How to Fix Bad Mediator Data” on page 157

---

# Installing and Configuring Solstice DiskSuite/Solaris Volume Manager Software

Before you begin, have available the following information.

- Mappings of your storage disk drives.
- The following completed configuration planning worksheets. See “Planning Volume Management” on page 24 for planning guidelines.
  - “Local File Systems With Mirrored Root Worksheet” in *Sun Cluster 3.1 Release Notes* or “Local File Systems with Non-Mirrored Root Worksheet” in *Sun Cluster 3.1 Release Notes*
  - “Disk Device Groups Worksheet” in *Sun Cluster 3.1 Release Notes*
  - “Volume Manager Configurations Worksheet” in *Sun Cluster 3.1 Release Notes*
  - “Metadevices Worksheet (Solstice DiskSuite/Solaris Volume Manager)” in *Sun Cluster 3.1 Release Notes*

The following table lists the tasks to perform to install and configure Solstice DiskSuite/Solaris Volume Manager software for Sun Cluster configurations.

---

**Note** – If you used SunPlex Manager to install Solstice DiskSuite software (Solaris 8), the procedures “How to Install Solstice DiskSuite Software” on page 129 through “How to Create State Database Replicas” on page 132 are already completed. Go to “Mirroring the Root Disk” on page 133 or “How to Create a Diskset” on page 147 to continue to configure Solstice DiskSuite software.

If you installed Solaris 9, Solaris Volume Manager is already installed and you can start at “How to Set the Number of Metadevice or Volume Names and Disksets” on page 130.

---

**TABLE A-1** Task Map: Installing and Configuring Solstice DiskSuite/Solaris Volume Manager Software

| Task                                                                             | For Instructions, Go To ...                                                                                              |
|----------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------|
| Plan the layout of your Solstice DiskSuite/Solaris Volume Manager configuration. | “Planning Volume Management” on page 24<br>“Solstice DiskSuite/Solaris Volume Manager Configuration Example” on page 127 |

**TABLE A-1** Task Map: Installing and Configuring Solstice DiskSuite/Solaris Volume Manager Software (Continued)

| Task                                                                                                                                      | For Instructions, Go To ...                                                                                                                                                         |
|-------------------------------------------------------------------------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| For Solaris 8, install Solstice DiskSuite software.                                                                                       | "How to Install Solstice DiskSuite Software" on page 129                                                                                                                            |
| Calculate the number of metadvice names and disksets needed for your configuration, and modify the <code>/kernel/drv/md.conf</code> file. | "How to Set the Number of Metadvice or Volume Names and Disksets" on page 130                                                                                                       |
| Create state database replicas on the local disks.                                                                                        | "How to Create State Database Replicas" on page 132                                                                                                                                 |
| (Optional) Mirror file systems on the root disk.                                                                                          | "Mirroring the Root Disk" on page 133                                                                                                                                               |
| Create disksets by using the <code>metaset</code> command.                                                                                | "How to Create a Diskset" on page 147                                                                                                                                               |
| Add disk drives to the disksets.                                                                                                          | "How to Add Drives to a Diskset" on page 150                                                                                                                                        |
| Repartition drives in a diskset to allocate space to slices 1 through 6.                                                                  | "How to Repartition Drives in a Diskset" on page 151                                                                                                                                |
| List device ID 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 152                                                                                                                             |
| Initialize the <code>md.tab</code> files.                                                                                                 | "How to Activate Metadevices or Volumes" on page 154                                                                                                                                |
| For dual-string configurations, configure mediator hosts, check the status of mediator data, and, if necessary, fix bad mediator data.    | "Mediators Overview" on page 155<br>"How to Add Mediator Hosts" on page 156<br>"How to Check the Status of Mediator Data" on page 156<br>"How to Fix Bad Mediator Data" on page 157 |
| Configure the cluster.                                                                                                                    | "Configuring the Cluster" on page 93                                                                                                                                                |

## Solstice DiskSuite/Solaris Volume Manager Configuration Example

The following example helps to explain the process for determining the number of disks to place in each diskset when you use Solstice DiskSuite/Solaris Volume Manager software. In this example three storage devices are used, and existing applications 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 used to determine the number of drives needed in the sample configuration. If you have three storage devices, you would need 28 drives that 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 Gbyte of disk space because the number of disks needed was rounded up.

**TABLE A-2** Determining 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             |
| oracle1 | 5 Gbytes  | 3x2.1 Gbyte disks * 2 (Mirror) | 6             |
| oracle2 | 10 Gbytes | 5x2.1 Gbyte disks * 2 (Mirror) | 10            |

The following table shows the allocation of drives among the two disksets and four data services.

**TABLE A-3** Division of Disksets

| Diskset     | Data Services    | Disks | Storage Device 1 | Storage Device 2 | Storage Device 3 |
|-------------|------------------|-------|------------------|------------------|------------------|
| dg-schost-1 | nfs1,<br>oracle1 | 12    | 4                | 4                | 4                |
| dg-schost-2 | nfs2,<br>oracle2 | 16    | 5                | 6                | 5                |

Initially, four disks on each storage device (a total of 12 disks) are assigned to dg-schost-1, and five or six disks on each (a total of 16) are assigned to dg-schost-2.

No hot spare disks are assigned to either diskset. A minimum of one hot spare disk per storage device per diskset enables one drive to be hot spared, which restores full two-way mirroring).



## ▼ How to Install Solstice DiskSuite Software

---

**Note** – If you used SunPlex Manager to install Solstice DiskSuite software, do not perform this procedure. Instead, go to “Mirroring the Root Disk” on page 133.

If you installed Solaris 9 software, do not perform this procedure. Solaris Volume Manager software is installed with Solaris 9 software. Instead, go to “How to Set the Number of Metadevice or Volume Names and Disksets” on page 130.

---

Perform this task on each node in the cluster.

1. **Become superuser on the cluster node.**
2. **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.

3. **Install the Solstice DiskSuite software packages.**

---

**Note** – If you have Solstice DiskSuite software patches to install, do not reboot after you install the Solstice DiskSuite software.

---

Install software packages in the order 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
```

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.

4. **If you installed from a CD-ROM, eject the CD-ROM.**
5. **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.
6. **Repeat Step 1 through Step 5 on the other nodes of the cluster.**
7. **From one node of the cluster, manually populate the global device namespace for Solstice DiskSuite.**

```
scgdevs
```

**8. Set the number of metadvice names and disksets expected in the cluster.**

Go to “How to Set the Number of Metadvice or Volume Names and Disksets” on page 130.

## ▼ How to Set the Number of Metadvice or Volume Names and Disksets

---

**Note** – If you used SunPlex Manager to install Solstice DiskSuite software, do not perform this procedure. Instead, go to “Mirroring the Root Disk” on page 133.

---

This procedure describes how to determine the number of Solstice DiskSuite metadvice or Solaris Volume Manager volume names and disksets 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 metadvice or volume names per diskset 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 `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 metadvice or volumes you will use.

---

**1. Have available the following completed worksheets:**

- “Disk Device Groups Worksheet” in *Sun Cluster 3.1 Release Notes*

**2. Determine the total number of disksets you expect to need in the cluster, then add one for private disk management.**

The cluster can have a maximum of 32 disksets, 31 disksets for general use plus one diskset for private disk management. The default number of disksets is 4. You will supply this value for the `md_nsets` field in Step 4.

**3. Determine the largest metadvice or volume name you expect to need for any diskset in the cluster.**

Each diskset can have a maximum of 8192 metadvice or volume names. You will supply this value for the `nmd` field in Step 4.

- a. Determine the quantity of metadvice or volume names you expect to need for each diskset.**

If you use local metadevices or volumes, ensure that each local metadvice or volume name is unique throughout the cluster and does not use the same name as any device ID (DID) in the cluster.

---

**Tip** – Choose a range of numbers to use exclusively for DID names and a range for each node to use exclusively for its local metadvice or volume names. For example, DIDs might use names in the range from d1 to d100, local metadevices or volumes on node 1 might use names in the range from d100 to d199, local metadevices or volumes on node 2 might use d200 to d299, and so on.

---

**b. Determine the highest of the metadvice or volume names you expect to use in any diskset.**

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/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 disksets served by each node. Failure to follow this guideline can result in serious Solstice DiskSuite/Solaris Volume Manager errors and possible loss of data.

---

**a. Set the `md_nsets` field to the value determined in Step 2.**

**b. Set the `nmd` field to the value 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 replicas.**

Go to “How to Create State Database Replicas” on page 132.

## ▼ How to Create State Database Replicas

---

**Note** – If you used SunPlex Manager to install Solstice DiskSuite software, do not perform this procedure. Instead, go to “Mirroring the Root Disk” on page 133.

---

Perform this procedure on each node in the cluster.

1. **Become superuser on the cluster node.**
2. **Create replicas on one or more local disks for each cluster node by using the `metadb` command.**

```
metadb -af slice-1 slice-2 slice-3
```

---

**Tip** – To provide protection of state data, which is necessary to run Solstice DiskSuite/Solaris Volume Manager software, create at least three replicas for each node. Also, you can place replicas on more than one disk to provide protection if one of the disks fails.

---

See the `metadb(1M)` man page and your Solstice DiskSuite/Solaris Volume Manager documentation for details.

3. **Verify the replicas.**

```
metadb
```

The `metadb` command displays the list of replicas.

4. **Do you intend to mirror file systems on the root disk?**

- If yes, go to “Mirroring the Root Disk” on page 133.
- If no, go to “How to Create a Diskset” on page 147 to create Solstice DiskSuite/Solaris Volume Manager disksets.

### Example—Creating State Database Replicas

The following example shows three Solstice DiskSuite state database replicas, each created on a different disk. For Solaris Volume Manager, the replica size would be larger.

```
metadb -af c0t0d0s7 c0t1d0s7 c1t0d0s7
metadb
flags first blk block count /dev/dsk/c0t0d0s7
a u 16 1034
a u 1050 1034
a u 2084 1034
a u 2084 1034
```

## 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 133
- “How to Mirror the Global Namespace” on page 137
- “How to Mirror File Systems Other than Root (/) That Cannot Be Unmounted” on page 140
- “How to Mirror File Systems That Can Be Unmounted” on page 144

---

**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 metadevice
```

---



---

**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 a node of the cluster.**
2. **Use the `metainit(1M)` command to put the root slice in a single-slice (one-way) concatenation.**

Use 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. Evacuate any resource groups or device groups from the node.**

```
scswitch -S -h node
```

`-S` Evacuates all resource groups and device groups

`-h node` Specifies the name of the node from which to evacuate 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. Is the disk that is used to mirror the root disk physically connected to more than one node (multiported)?**

- If no, go to Step 11.
- If yes, enable the `localonly` property of the raw disk device group for the disk 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 `sccidadm(1M) -L` command to display the full device ID (DID) pseudo-driver 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 DID pseudo-driver name.

```
sccidadm -L
...
```

```
1 phys-schost-3:/dev/rdisk/c1t1d0 /dev/did/rdsk/d2
scconf -c -D name=dsk/d2,localonly=true
```

For more information about the `localonly` property, see the `scconf_dg_rawdisk(1M)` man page.

**b. View the node list of the raw disk device group.**

Output will look similar to the following, where *N* is the DID number.

```
scconf -pvv | grep dsk/dN
Device group name: dsk/d2
...
(dsk/d2) Device group node list: phys-schost-1, phys-schost-3
...
```

**c. Does the node list contain more than one node name?**

- If yes, go to Step d.
- If no, go to Step e.

**d. Remove all nodes from the node list for the raw disk device group except the node whose root disk you mirrored.**

Only the node whose root disk you mirrored should remain in the node list.

```
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
```

**e. 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 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
```

**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. Do you intend to mirror the global namespace,  
/global/.devices/node@nodeid?**

- If yes, go to “How to Mirror the Global Namespace” on page 137.
- If no, go to Step 14.

**14. Do you intend to mirror file systems than cannot be unmounted?**

- If yes, go to “How to Mirror File Systems Other than Root (/) That Cannot Be Unmounted” on page 140.
- If no, go to Step 15.

**15. Do you intend to mirror user-defined file systems?**

- If yes, go to “How to Mirror File Systems That Can Be Unmounted” on page 144.
- If no, go to “How to Create a Diskset” on page 147 to create a diskset.

## Example—Mirroring the Root (/) File System

The following example shows creation of mirror d0 on the node `phys-schost-1`, which consists of submirror d10 on partition `c0t0d0s0` and submirror d20 on partition `c2t2d0s0`. Disk `c2t2d0` is a multiported 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

(Reboot the node)
scswitch -S -h phys-schost-1
shutdown -g0 -y -i6

(Attach the second submirror)
metattach d0 d20
d0: Submirror d20 is attached

(Display the node list of the mirror disk's raw disk device group)
sconfg -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 for the raw disk device group)
```



```
scconf -r -D name=dsk/d2,nodelist=phys-schost-3

(Enable the localonly property of the mirrored disk's raw disk device group)
scconf -c -D name=dsk/d2,localonly=true

(Record the alternate boot path)
ls -l /dev/rdisk/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 metadvice or volume name for the mirror *must* be unique throughout the cluster.

---

5. **Attach the second submirror to the mirror.**

This attachment starts a sync 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/rdsk/mirror /global/.devices/node@nodeid ufs 2 no global
```

**7. Repeat Step 1 through Step 6 on each remaining node of the cluster.**

Ensure that each metadevice or volume name for a mirror is unique throughout the cluster.

**8. Wait for the sync of the mirrors, started in Step 5, to complete.**

Use the `metastat(1M)` command to view mirror status.

```
metastat mirror
```

**9. Is the disk that is used to mirror the global namespace physically connected to more than one node (multiported)?**

- If no, go to Step 10.
- If yes, enable the `localonly` property of the raw disk device group for the disk 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 (DID) pseudo-driver 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 DID pseudo-driver name.

```
scdidadm -L
...
1 phys-schost-3:/dev/rdisk/clt1d0 /dev/did/rdisk/d2
scconf -c -D name=dsk/d2,localonly=true
```

For more information about the `localonly` property, see the `scconf_dg_rawdisk(1M)` man page.

**b. View the node list of the raw disk device group.**

Output will look similar to the following, where *N* is the DID number.

```
scconf -pvv | grep dsk/dN
Device group name: dsk/d2
...
(dsk/d2) Device group node list: phys-schost-1, phys-schost-3
...
```

**c. Does the node list contain more than one node name?**

- If yes, go to Step d.
- If no, go to Step e.

**d. Remove all nodes from the node list for the raw disk device group except the node whose root disk you mirrored.**

Only the node whose root disk you mirrored should remain in the node list.

```
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

**e. 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 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

**10. Do you intend to mirror file systems other than root (/) that cannot be unmounted?**

- If yes, go to “How to Mirror File Systems Other than Root (/) That Cannot Be Unmounted” on page 140.
- If no, go to Step 11.

**11. Do you intend to mirror user-defined file systems?**

- If yes, go to “How to Mirror File Systems That Can Be Unmounted” on page 144.
- If no, go to “How to Create a Diskset” on page 147 to create a diskset.

## 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`. Disk `c2t2d0` is a multiported 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 DID 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 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 for the raw disk device group)
scconf -r -D name=dsk/d2,nodelist=phys-schost-3

(Enable the localonly property of the mirrored disk's raw disk device group)
scconf -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.**  
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 does *not* need to be unique throughout the cluster.

---

**5. Repeat Step 1 through Step 4 for each unmountable file system to be mirrored.**

**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. Evacuate any resource groups or device groups from the node.**

```
scswitch -S -h node

-S Evacuates all resource groups and device groups
-h node Specifies the name of the node from which to evacuate resource
 or device groups
```

**8. Reboot the node.**

```
shutdown -g0 -y -i6
```

**9. Attach the second submirror to each mirror.**

This attachment starts a sync of the submirrors.

```
metattach mirror submirror2
```

**10. Wait for the sync of the mirrors, started in Step 9, to complete.**

Use the `metastat(1M)` command to view mirror status.

```
metastat mirror
```

**11. Is the disk that is used to mirror the unmountable file system physically connected to more than one node (multiported)?**

- If no, go to Step 12.
- If yes, enable the `localonly` property of the raw disk device group for the disk 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 (DID) pseudo-driver 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 DID pseudo-driver name.

```
scdidadm -L
...
1 phys-schost-3:/dev/rdisk/c1t1d0 /dev/did/rdisk/d2
sccnf -c -D name=dsk/d2,localonly=true
```

For more information about the `localonly` property, see the `sccnf_dg_rawdisk(1M)` man page.

- b. View the node list of the raw disk device group.

Output will look similar to the following, where *N* is the DID number.

```
sccnf -pvv | grep dsk/dN
Device group name: dsk/d2
...
(dsk/d2) Device group node list: phys-schost-1, phys-schost-3
...
```

- c. Does the node list contain more than one node name?

- If yes, go to Step d.
- If no, go to Step e.

- d. Remove all nodes from the node list for the raw disk device group except the node whose root disk you mirrored.

Only the node whose root disk you mirrored should remain in the node list.

```
sccnf -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
```

- e. Use the `sccnf(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 prevents unintentional fencing of the node from its boot device if the boot device is connected to multiple nodes.

```
sccnf -c -D name=rawdisk-groupname,localonly=true

-D name=rawdisk-groupname Specifies the name of the raw disk device group
```

## 12. Do you intend to mirror user-defined file systems?

- If yes, go to “How to Mirror File Systems That Can Be Unmounted” on page 144.
- If no, go to “How to Create a Diskset” on page 147 to create a diskset.

## 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`. Disk `c2t2d0` is a multiported 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

(Reboot the node)
scswitch -S -h phys-schost-1
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 DID name of the mirrored disk's raw disk device group)
scdidadm -L
...
1 phys-schost-3:/dev/rdsk/c2t2d0 /dev/did/rdsk/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 for the raw disk device group)
scconf -r -D name=dsk/d2,nodelist=phys-schost-3

(Enable the localonly property of the mirrored disk's raw disk device group)
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, nodes do not need to be rebooted.

1. **Become superuser on a node of the cluster.**
2. **Put in a single slice (one-way) concatenation the slice on which a user-defined file system that can be unmounted resides.**

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 metadvice or volume name for the mirror does *not* need to be unique throughout the cluster.

---

5. **Repeat Step 1 through Step 4 for each mountable file system to be mirrored.**
6. **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
```

7. **Attach the second submirror to the mirror.**

This attachment starts a sync of the submirrors.



```
metattach mirror submirror2
```

**8. Wait for the sync of the mirrors, started in Step 7, to complete.**

Use the `metastat(1M)` command to view mirror status.

```
metastat mirror
```

**9. Is the disk that is used to mirror the user-defined file system physically connected to more than one node (multiported)?**

- If no, go to Step 10.
- If yes, enable the `localonly` property of the raw disk device group for the disk 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 (DID) pseudo-driver 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 DID pseudo-driver name.

```
scdidadm -L
...
1 phys-schost-3:/dev/rdisk/ctl1d0 /dev/did/rdisk/d2
sconfig -c -D name=dsk/d2,localonly=true
```

For more information about the `localonly` property, see the `sconfig_dg_rawdisk(1M)` man page.

**b. View the node list of the raw disk device group.**

Output will look similar to the following, where *N* is the DID number.

```
sconfig -pvv | grep dsk/dN
Device group name: dsk/d2
...
(dsk/d2) Device group node list: phys-schost-1, phys-schost-3
...
```

**c. Does the node list contain more than one node name?**

- If yes, go to Step d.
- If no, go to Step e.

**d. Remove all nodes from the node list for the raw disk device group except the node whose root disk you mirrored.**

Only the node whose root disk you mirrored should remain in the node list.

```
sconfig -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

**e. 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 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

**10. Create a diskset.**

Go to “How to Create a Diskset” on page 147.

## 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`. Disk `c2t2d0` is a multiported disk, so the `localonly` property is enabled.

```
(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 fck mount mount
#to mount to fck 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 DID 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 node list of the mirror disk's raw disk device group)
 # 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 for the raw disk device group)
 # sccconf -r -D name=dsk/d2,nodelist=phys-schost-3

 (Enable the localonly property of the mirrored disk's raw disk device group)
 # sccconf -c -D name=dsk/d2,localonly=true

```

## ▼ How to Create a Diskset

Perform this procedure for each diskset you create.

---

**Note** – If you used SunPlex Manager to install Solstice DiskSuite, one to three disksets might already exist. See “Using SunPlex Manager to Install Sun Cluster Software” on page 62 for information about the metaset created by SunPlex Manager.

---

1. **Do you intend to create more than three disksets in the cluster?**
  - If yes, go to Step 2 to prepare the cluster for more than three disksets. Do this regardless of whether you are installing disksets for the first time or you are adding more disksets to a fully configured cluster.
  - If no, go to Step 6.
2. **Ensure that the value of the `md_nsets` variable is set high enough to accommodate the total number of disksets you intend to create in the 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 disksets in the cluster will be greater than the existing value of `md_nsets` minus one, on each node increase the value of `md_nsets` to the desired value.**

The maximum permissible number of disksets is one less than the value of `md_nsets`. The maximum possible value of `md_nsets` is 32.

- 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/Solaris Volume Manager errors and possible loss of data.

---

- d. **From one node, shut down the cluster.**

```
scshutdown -g0 -y
```

- e. **Reboot each node of the cluster.**

```
ok> boot
```

3. **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.

4. **From one node of the cluster, run the `scgdevs(1M)` command.**

5. **On each node, verify that the `scgdevs` command has completed before you attempt to create any disksets.**

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
```

6. **Ensure that the diskset you intend to create meets one of the following requirements.**

- If configured with exactly two disk strings, the diskset must connect to exactly two nodes and use exactly two mediator hosts, which must be the same two hosts used for the diskset. See “Mediators Overview” on page 155 for details on how to set up mediators.
- If configured with more than two disk strings, ensure that for any two disk strings S1 and S2, the sum of the number of disks on those strings exceeds the number of disks on the third string S3. Stated as a formula, the requirement is that  $\text{count}(S1) + \text{count}(S2) > \text{count}(S3)$ .

7. **Ensure that the local state database replicas exist.**

For instructions, see “How to Create State Database Replicas” on page 132.

8. **Become superuser on the cluster node that will master the diskset.**

9. **Create the diskset.**

This command also registers the diskset as a Sun Cluster disk device group.

```
metaset -s setname -a -h node1 node2
```

|                   |                                                                |
|-------------------|----------------------------------------------------------------|
| -s <i>setname</i> | Specifies the diskset name                                     |
| -a                | Adds (creates) the diskset                                     |
| -h <i>node1</i>   | Specifies the name of the primary node to master the diskset   |
| <i>node2</i>      | Specifies the name of the secondary node to master the diskset |

---

**Note** – Running the `metaset` command to set up a Solstice DiskSuite/Solaris Volume Manager device group on a cluster creates one secondary node by default, regardless of the number of nodes that are included in that device group. You can change the desired number of secondary nodes by using the `scsetup(1M)` utility after the device group is created. Refer to “Administering Disk Device Groups” in *Sun Cluster 3.1 System Administration Guide* for more information about how to change the `numsecondaries` property.

---

#### 10. Verify the status of the new diskset.

```
metaset -s setname
```

#### 11. Add drives to the diskset.

Go to “Adding Drives to a Diskset” on page 149.

## Example—Creating a Diskset

The following command creates two disksets, `dg-schost-1` and `dg-schost-2`, with the nodes `phys-schost-1` and `phys-schost-2` assigned 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 Diskset

When you add a disk drive to a diskset, Solstice DiskSuite/Solaris Volume Manager repartitions it as follows so that the state database for the diskset can be placed on the drive.

- A small portion of each drive is reserved in slice 7 for use by Solstice DiskSuite/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 diskset only if slice 7 is not set up correctly.

- Any existing data on the disks is lost by the repartitioning.
- If slice 7 starts at cylinder 0, and the disk is large enough to contain a state database replica, the disk is not repartitioned.

## ▼ How to Add Drives to a Diskset

### 1. Become superuser on the node.

### 2. Ensure that the diskset has been created.

For instructions, see “How to Create a Diskset” on page 147.

### 3. List the device ID (DID) mappings.

```
scdidadm -L
```

- Choose drives that are shared by the cluster nodes that will master or potentially master the diskset.
- Use the full DID pseudo-driver names when you add drives to a diskset.

The first column of output is the DID instance number, the second column is the full path (physical path), and the third column is the full DID pseudo-driver 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 DID 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 diskset.

```
metaset -s setname -t
```

`-s setname`            Specifies the diskset name

`-t`                      Takes ownership of the diskset

### 5. Add the drives to the diskset.

Use the full DID pseudo-driver name.

```
metaset -s setname -a DIDname
```

`-a`                      Adds the disk drive to the diskset

`DIDname`                Device ID (DID) name of the shared disk

---

**Note** – Do *not* use the lower-level device name (*cNtXdY*) when you add a drive to a diskset. 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 diskset and drives.

```
metaset -s setname
```

7. Do you intend to repartition drives for use in metadevices or volumes?

- If yes, go to “How to Repartition Drives in a Diskset” on page 151.
- If no, go to “How to Create an `md.tab` File” on page 152 to define metadevices or volumes by using an `md.tab` file.

## Example—Adding Drives to a Diskset

The `metaset` command adds the disk drives `/dev/did/dsk/d1` and `/dev/did/dsk/d2` to the diskset `dg-schost-1`.

```
metaset -s dg-schost-1 -a /dev/did/dsk/d1 /dev/did/dsk/d2
```

## ▼ How to Repartition Drives in a Diskset

The `metaset(1M)` command repartitions drives in a diskset so that a small portion of each drive is reserved in slice 7 for use by Solstice DiskSuite/Solaris Volume Manager software. The remainder of the space on each drive is placed into slice 0. To make more effective use of the disk, 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 diskset.

When you repartition a drive, you must meet the following conditions to prevent the `metaset(1M)` command from repartitioning the disk.

- Create a partition 7 starting at cylinder 0 that is large enough to hold a state database replica (approximately 2 Mbytes).
- 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 disk.

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 152.

## ▼ 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 disksets you created.

---

**Note** – If you are using local metadevices or volumes, ensure that local metadevices or volumes names are distinct from the device ID (DID) names used to form disksets. For example, if the DID name `/dev/did/dsk/d3` is used in a diskset, 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, for node 2 use d200-d299, and so on.

---

### 1. Become superuser on the cluster node.

### 2. List the DID mappings for reference when you create your `md.tab` file.

Use the full DID pseudo-driver 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 path (physical path), and the third column is the full DID pseudo-driver 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/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 disks that will be used for the submirrors, you must back up the data before metadvice or volume setup and restore it onto the mirror.

---

4. Activate the metadevices or volumes defined in the `md.tab` files.

Go to “How to Activate Metadevices or Volumes” on page 154.

## Example—Sample `md.tab` File

The following sample `md.tab` file defines the metadevices for the diskset named `dg-schost-1`. The ordering of lines in the `md.tab` file is not important.

```
dg-schost-1/d0 -t dg-schost-1/d1 dg-schost-1/d4
dg-schost-1/d1 -m dg-schost-1/d2
dg-schost-1/d2 1 1 /dev/did/rdisk/d1s4
dg-schost-1/d3 1 1 /dev/did/rdisk/d5s4
dg-schost-1/d4 -m dg-schost-1/d5
dg-schost-1/d5 1 1 /dev/did/rdisk/d3s5
dg-schost-1/d6 1 1 /dev/did/rdisk/d57s5
```

The sample `md.tab` file is constructed as follows.

---

**Note** – The following example uses Solstice DiskSuite terminology. For Solaris Volume Manager, a trans metadvice is instead called a *transactional volume* and a metadvice is instead called a *volume*. Otherwise, the following process is valid for both volume managers.

---

- The first line defines the trans metadvice `d0` to consist of a master (UFS) metadvice `d1` and a log device `d4`. The `-t` signifies this is a trans metadvice. The master and log devices are specified by their position after the `-t` flag.

```
dg-schost-1/d0 -t dg-schost-1/d1 dg-schost-1/d4
```

- The second line defines the master device as a mirror of the metadevices. The `-m` in this definition signifies a mirror device, and one of the submirrors, `d2`, is associated with the mirror device, `d1`.

```
dg-schost-1/d1 -m dg-schost-1/d2
```

- The fifth line similarly defines the log device, `d4`, as a mirror of metadevices.

```
dg-schost-1/d4 -m dg-schost-1/d5
```

- The third line defines the first submirror of the master device, `d2`, as a one-way stripe.

```
dg-schost-1/d2 1 1 /dev/did/rdisk/d1s4
```

- The fourth line defines the second submirror of the master device, `d3`.

```
dg-schost-1/d3 1 1 /dev/did/rdisk/d55s4
```

- Finally, the log device submirrors, `d5` and `d6`, are defined. In this example, simple metadevices for each submirror are created.

```
dg-schost-1/d5 1 1 /dev/did/rdisk/d3s5
dg-schost-1/d6 1 1 /dev/did/rdisk/d57s5
```

## ▼ How to Activate Metadevices or Volumes

Perform this procedure to activate Solstice DiskSuite metadevices or Solaris Volume Manager volumes defined in the `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 diskset on the node where the command will be executed.**

4. **Take ownership of the diskset.**

```
metaset -s setname -t
-s setname Specifies the diskset name
-t Takes ownership of the diskset
```

5. **Activate the diskset'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 diskset in the cluster.**

If necessary, run the `metainit(1M)` command from another node that has connectivity to the disks. This step is required for cluster-pair topologies, where the disks 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. Does your cluster contain disksets configured with exactly two disk enclosures and two nodes?

- If yes, those disksets require mediators. Go to “Mediators Overview” on page 155 to add mediator hosts.
- If no, go to “How to Add Cluster File Systems” on page 94 to create a cluster file system.

## Example—Activating Metadevices or Volumes in the `md.tab` File

In the following example, all metadevices defined in the `md.tab` file for diskset `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
```

## Mediators Overview

A 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.

Mediators are required for all Solstice DiskSuite/Solaris Volume Manager disksets configured with exactly two disk strings and two cluster nodes. A *disk string* consists of a disk enclosure, its physical disks, 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.

- Disksets must be configured with exactly two mediator hosts, and those two mediator hosts must be the same two cluster nodes used for the diskset.
- A diskset cannot have more than two mediator hosts.
- Mediators cannot be configured for disksets that do not meet the two-string and two-host criteria.

These rules do not require that the entire cluster have exactly two nodes. Rather, they only require that those disksets 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 mediators.

1. **Become superuser on the node that currently masters the diskset you intend to add mediator hosts to.**
2. **Run the `metaset(1M)` command to add each node with connectivity to the diskset as a mediator host for that diskset.**

```
metaset -s setname -a -m mediator-host-list
-s setname Specifies the diskset name
-a Adds to the diskset
-m mediator-host-list Specifies the name of the node to add as a mediator host
 for the diskset
```

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 156.

## Example—Adding Mediator Hosts

The following example adds the nodes `phys-schost-1` and `phys-schost-2` as mediator hosts for the diskset `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 156.**
2. **Run the `medstat` command.**

```
medstat -s setname
-s setname Specifies the diskset name
```

See the `medstat(1M)` man page for more information.

**3. Is Bad the value in the Status field?**

- If yes, go to “How to Fix Bad Mediator Data” on page 157 to repair the affected mediator host.
- If no, go to “How to Add Cluster File Systems” on page 94 to create a cluster file system.

## ▼ How to Fix Bad Mediator Data

Perform this procedure to repair bad mediator data.

1. **Identify the mediator host(s) with bad mediator data as described in the procedure “How to Check the Status of Mediator Data” on page 156.**
2. **Become superuser on the node that owns the affected diskset.**
3. **Remove the mediator host(s) with bad mediator data from all affected disksets.**

```
metaset -s setname -d -m mediator-host-list
-s setname Specifies the diskset name
-d Deletes from the diskset
-m mediator-host-list Specifies the name of the node to remove as a mediator
 host for the diskset
```

**4. Restore the mediator host.**

```
metaset -s setname -a -m mediator-host-list
-a Adds to the diskset
-m Specifies the name of the node to add as a mediator host for the
mediator-host-list diskset
```

See the `mediator(7D)` man page for details about mediator-specific options to the `metaset` command.

**5. Create a cluster file system.**

Go to “How to Add Cluster File Systems” on page 94.



## 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 appendix, along with the planning information in “Planning Volume Management” on page 24. See your VxVM documentation for additional details.

The following procedures are in this appendix.

- “How to Install VERITAS Volume Manager Software and Encapsulate the Root Disk” on page 161
- “How to Mirror the Encapsulated Root Disk” on page 165
- “How to Install VERITAS Volume Manager Software Only” on page 167
- “How to Create a rootdg Disk Group on a Non-Root Disk” on page 170
- “How to Create and Register a Disk Group” on page 171
- “How to Assign a New Minor Number to a Disk Device Group” on page 173
- “How to Verify the Disk Group Configuration” on page 174
- “How to Unencapsulate the Root Disk” on page 174

---

## Installing and Configuring VxVM Software

Before you begin, have available the following information.

- Mappings of your storage disk drives.
- The following completed configuration planning worksheets. See “Planning Volume Management” on page 24 for planning guidelines.
  - “Local File Systems With Mirrored Root Worksheet” in *Sun Cluster 3.1 Release Notes* or “Local File Systems with Non-Mirrored Root Worksheet” in *Sun Cluster 3.1 Release Notes*

- “Disk Device Groups Worksheet” in *Sun Cluster 3.1 Release Notes*
- “Volume Manager Configurations Worksheet” in *Sun Cluster 3.1 Release Notes*

The following table lists the tasks to perform to install and configure VxVM software for Sun Cluster configurations.

**TABLE B-1** Task Map: Installing and Configuring VxVM Software

| Task                                                                                                                                                                                                                       | For Instructions, Go To ...                                                                                                                          |
|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------|
| Plan the layout of your VxVM configuration.                                                                                                                                                                                | “Planning Volume Management” on page 24                                                                                                              |
| Determine how you will create the <code>rootdg</code> disk group on each node.                                                                                                                                             | “Setting Up a <code>rootdg</code> Disk Group Overview” on page 160                                                                                   |
| Install VxVM software and create the <code>rootdg</code> disk group:                                                                                                                                                       |                                                                                                                                                      |
| <ul style="list-style-type: none"> <li>■ <b>Method 1</b> – Install VxVM software and encapsulate the root disk by using the <code>scvxinstall</code> command, and optionally mirror the encapsulated root disk.</li> </ul> | “How to Install VERITAS Volume Manager Software and Encapsulate the Root Disk” on page 161<br>“How to Mirror the Encapsulated Root Disk” on page 165 |
| <ul style="list-style-type: none"> <li>■ <b>Method 2</b> – Install VxVM software and create <code>rootdg</code> on local, non-root disks.</li> </ul>                                                                       | “How to Install VERITAS Volume Manager Software Only” on page 167<br>“How to Create a <code>rootdg</code> Disk Group on a Non-Root Disk” on page 170 |
| Create disk groups and volumes.                                                                                                                                                                                            | “How to Create and Register a Disk Group” on page 171                                                                                                |
| Resolve any minor number conflicts between disk device groups by assigning a new minor number.                                                                                                                             | “How to Assign a New Minor Number to a Disk Device Group” on page 173                                                                                |
| Verify the disk groups and volumes.                                                                                                                                                                                        | “How to Verify the Disk Group Configuration” on page 174                                                                                             |
| Configure the cluster.                                                                                                                                                                                                     | “Configuring the Cluster” on page 93                                                                                                                 |

## Setting Up a `rootdg` Disk Group Overview

Each cluster node requires the creation of a `rootdg` 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 `rootdg` disk group must be restricted to only that node.
- Remote nodes must never access data stored in another node’s `rootdg`.
- Do not use the `scconf(1M)` command to register the `rootdg` disk group as a disk device group.



- Whenever possible, configure the `rootdg` for each node on a non-shared disk.

Sun Cluster software supports the following methods to configure the `rootdg` 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 non-root 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 operating environment 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 non-root disks for use as `rootdg`.

A `rootdg` disk group that is created on local non-root disks is local to that node, neither globally accessible nor highly available. As with the root disk, to encapsulate a non-root 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.

## Where to Go From Here

Install VxVM by using one of the following installation methods, depending on how you intend to create the `rootdg` disk group.

- If you intend to encapsulate the root disk, go to “How to Install VERITAS Volume Manager Software and Encapsulate the Root Disk” on page 161.
- If you intend to create the `rootdg` disk group on local non-root disks, go to “How to Install VERITAS Volume Manager Software Only” on page 167.

## ▼ 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 `rootdg` disk group on local, non-root disks, go instead to “How to Install VERITAS Volume Manager Software Only” on page 167.

---

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 device(s) VxVM will manage.

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. **Become superuser on a node you intend to install with VxVM.**
3. **Add all nodes in the cluster to the cluster node authentication list.**
  - a. **Start the `scsetup(1M)` utility.**

```
scsetup
```

The Main Menu is displayed.
  - b. **To access the New Nodes Menu, type 6 at the Main Menu.**
  - c. **To add a node to the authorized list, type 3 at the New Nodes Menu.**
  - d. **Specify the name of a machine which may add itself.**

Follow the prompts to add the node's name to the cluster. You will be asked for the name of the node to be added.
  - e. **Verify that the task has been performed successfully.**

The `scsetup` utility prints a Command completed successfully message if it completes the task without error.
  - f. **Repeat Step c through Step e for each node of the cluster until all cluster nodes are added to the node authentication list.**
  - g. **Quit the `scsetup` utility.**
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:

- Disables Dynamic Multipathing (DMP)

---

**Note** – Although the `scvxinstall` utility disables Dynamic Multipathing (DMP) at the start of installation processing, DMP is automatically re-enabled by VxVM version 3.1.1 or later when the `VRTSvxvm` package is installed. Earlier versions of VxVM must still run with DMP disabled.

---

- 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 `rootdg` 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** – During installation there are two automatic reboots. 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. Do you intend to install VxVM on another node?**

- If yes, repeat Step 2 through Step 13.
- If no, go to Step 15.

**15. Are there one or more nodes that you do *not* intend to install with VxVM?**

---

**Note** – If you intend to enable the VxVM cluster feature, you *must* install VxVM on all nodes of the cluster.

---

- If yes, go to Step 16.
- If no, go to Step 17.

**16. 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.

**17. Prevent any new machines from being added to the cluster.**

- a. Start the `scsetup(1M)` utility.

```
scsetup
The Main Menu is displayed.
```

- b. To access the New Nodes Menu, type 6 at the Main Menu.

- c. Type 1 at the New Nodes Menu.

Follow the `scsetup` prompts. This option tells the cluster to ignore all requests coming in over the public network from any new machine that tries to add itself to the cluster.

- d. Quit the `scsetup` utility.

**18. Do you intend to mirror the encapsulated root disk?**

- If yes, go to “How to Mirror the Encapsulated Root Disk” on page 165.
- If no, go to “How to Create and Register a Disk Group” on page 171.

---

**Note** – If you later need to unencapsulate the root disk, follow the procedures in “How to Unencapsulate the Root Disk” on page 174.

---

## ▼ 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 following the procedures in your VxVM documentation.

For maximum availability and simplified administration, use a local disk for the mirror. See “Mirroring the Root Disk” on page 29 for additional guidelines.




---

**Caution** – Do not use a quorum device to mirror a root disk. Doing so might prevent the node from booting from the root disk mirror under certain circumstances.

---

### 2. Display the device ID (DID) mappings.

```
sddidadm -L
```

### 3. From the DID mappings, locate the disk used to mirror the root disk.

### 4. Extract the raw disk device group name from the DID 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/rdsk/dN
```

### 5. View the node list of the raw disk device group.

Output will look similar to the following.

```
sconfig -pvv | grep dsk/dN
Device group name: dsk/dN
...
(dsk/dN) Device group node list: phys-schost-1, phys-schost-3
...
```

### 6. Does the node list contain more than one node name?

- If yes, go to Step 7.

- If no, go to Step 9.

**7. Remove from the node list for the raw disk device group all nodes except the node whose root disk you mirrored.**

Only the node whose root disk you mirrored should remain in the node list.

```
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

**8. 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 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.

**9. Repeat this procedure for each node in the cluster whose encapsulated root disk you intend to mirror.**

**10. Create disk groups.**

Go to “How to Create and Register a Disk Group” on page 171.

## 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 multiported 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 for the raw disk device group)
scconf -r -D name=dsk/d2,nodelist=phys-schost-3

(Enable the localonly property for the mirror disk's raw disk device group)
scconf -c -D name=dsk/d2,localonly=true
```

## ▼ 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 `rootdg` disk group by encapsulating the root disk, do not use this procedure. Instead, go to “How to Install VERITAS Volume Manager Software and Encapsulate the Root Disk” on page 161 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 device(s) VxVM will manage.

1. **Ensure that all nodes in the cluster are running in cluster mode.**
2. **Become superuser on a cluster node you intend to install with VxVM.**
3. **Add all nodes in the cluster to the cluster node authentication list.**
  - a. **Start the `scsetup(1M)` utility.**

```
scsetup
```

The Main Menu is displayed.
  - b. **To access the New Nodes Menu, type 6 at the Main Menu.**
  - c. **To add a node to the authorized list, type 3 at the New Nodes Menu.**
  - d. **Specify the name of a machine which may add itself.**

Follow the prompts to add the node's name to the cluster. You will be asked for the name of the node to be added.
  - e. **Verify that the task has been performed successfully.**

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

f. Repeat Step c through Step e for each node of the cluster until all cluster nodes are added to the node authentication list.

g. Quit the `scsetup` utility.

4. Insert the VxVM CD-ROM into the CD-ROM drive on the node.

5. Start `scvxinstall` in non-interactive installation mode.

```
scvxinstall -i
```

The `scvxinstall` command automatically performs the following tasks.

- Disables Dynamic Multipathing (DMP)

---

**Note** – Although the `scvxinstall` utility disables Dynamic Multipathing (DMP) at the start of installation processing, DMP is automatically re-enabled by VxVM version 3.1.1 or later when the `VRTSvxvm` package is installed. Earlier versions of VxVM must still run with DMP disabled.

---

- 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

---

**Note** – You will add VxVM licenses during the next procedure, “How to Create a `rootdg` Disk Group on a Non-Root Disk” on page 170.

---

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 VRTSvmman
```

10. Do you intend to install VxVM on another node?

- If yes, repeat Step 2 through Step 9.



- If no, go to Step 11.

**11. Are there one or more nodes that you do *not* intend to install with VxVM?**

---

**Note** – If you intend to enable the VxVM cluster feature, you *must* install VxVM on all nodes of the cluster.

---

- If yes, go to Step 12.
- If no, go to Step 13.

**12. 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 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.

**13. Prevent any new machines from being added to the cluster.**

- a. Start the `scsetup(1M)` utility.

```
scsetup
```

The Main Menu is displayed.

- b. To access the New Nodes Menu, type 6 at the Main Menu.

- c. Type 1 at the New Nodes Menu.

Follow the `scsetup` prompts. This option tells the cluster to ignore all requests coming in over the public network from any new machine that tries to add itself to the cluster.

- d. Quit the `scsetup` utility.

**14. Create a `rootdg` disk group.**

Go to “How to Create a rootdg Disk Group on a Non-Root Disk” on page 170.

## ▼ How to Create a rootdg Disk Group on a Non-Root Disk

Use this procedure to create a rootdg disk group by encapsulating or initializing local disks other than the root disk.

1. **Have available the VERITAS Volume Manager (VxVM) license keys.**
2. **Become superuser on the node.**
3. **(Optional) If the disks will 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.

4. **Start the `vxinstall` utility.**

```
vxinstall
```

When prompted, make the following choices or entries.

- Supply the VxVM license key.
- If you intend to enable the VxVM cluster feature, supply the cluster feature license key.
- Choose Custom Installation.
- Do not encapsulate the root disk.
- Choose any disks to add to the rootdg disk group.
- Do not accept automatic reboot.

5. **Evacuate any resource groups or device groups from the node.**

```
scswitch -S -h node
```

-S                   Evacuates all resource groups and device groups

-h *node*            Specifies the name of the node from which to evacuate resource or device groups

6. **Reboot the node.**

```
shutdown -g0 -y -i6
```

7. **Use the `vxdiskadm` command to add multiple disks to the rootdg disk group.**

The rootdg disk group becomes tolerant of a disk failure when it contains multiple disks. See VxVM documentation for procedures.

## 8. Create disk groups.

Go to “How to Create and Register a Disk Group” on page 171.

# ▼ 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 3.1 System Administration Guide* 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 up the disk group you add.

### 1. Have available the following information.

- Mappings of your storage disk drives. See the appropriate manual of the *Sun Cluster 3.x Hardware Administration Manual* to perform an initial installation of your storage device.
- The following completed configuration planning worksheets.
  - “Local File Systems With Mirrored Root Worksheet” in *Sun Cluster 3.1 Release Notes* or “Local File Systems with Non-Mirrored Root Worksheet” in *Sun Cluster 3.1 Release Notes*
  - “Disk Device Groups Worksheet” in *Sun Cluster 3.1 Release Notes*
  - “Volume Manager Configurations Worksheet” in *Sun Cluster 3.1 Release Notes*

See “Planning Volume Management” on page 24 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 Parallel Server/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 documented in the VxVM documentation.

---

**Note** – You can use Dirty Region Logging (DRL) to decrease volume recovery time in the event of a node failure. However, DRL might decrease I/O throughput.

---

**4. Is the VxVM cluster feature enabled?**

- If yes, go to Step 7. If the VxVM cluster feature is enabled, do *not* register a shared disk group as a Sun Cluster disk device group.
- If no, go to Step 5.

**5. Register the disk group as a Sun Cluster disk device group.**

**a. Start the `scsetup(1M)` utility.**

```
scsetup
```

**b. To work with disk device groups, type 4 (Device groups and volumes).**

**c. To register a disk device group, type 1 (Register a VxVM disk group).**

Follow the instructions and type the VxVM disk device group to be registered as a Sun Cluster disk device group.

**d. If you encounter the following error when you attempt to register the disk device group, renumber the disk device group.**

```
scconf: Failed to add device group - in use
```

To renumber the disk device group, use the procedure “How to Assign a New Minor Number to a Disk Device Group” on page 173. This procedure enables you to assign a new minor number that does not conflict with a minor number used by existing disk device groups.

**e. When finished, type `q` (Quit) to leave the `scsetup` utility.**

**6. Verify that the disk device group is registered.**

Look for the disk device information for the new disk displayed by the following command.

```
scstat -D
```

---

**Note** – If you change any configuration information for a VxVM disk group or volume, you must reregister the Sun Cluster disk device group by using `scsetup`. Such configuration changes include adding or removing volumes, as well as changing the group, owner, or permissions of existing volumes. Reregistration after configuration changes ensures that the global namespace is in the correct state. See “Administering Disk Device Groups” in *Sun Cluster 3.1 System Administration Guide* for procedures for how to reregister a disk device group.

---

7. Verify the configuration of your VxVM disk groups and volumes.  
Go to “How to Verify the Disk Group Configuration” on page 174.

## ▼ 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.  

```
ls -l /global/.devices/node@1/dev/vx/dsk/*
```
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 Step 5 of “How to Create and Register a Disk Group” on page 171 to register the disk group as a Sun Cluster disk device group.

## 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 `vxdg renumber` 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
vxdg renumber dg3 5000
```

## ▼ 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 (rootdg), and disk groups are imported on the current primary node only.**

```
vxdisk list
```

2. **Verify that all volumes have been started.**

```
vxprint
```

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 93.

## ▼ How to Unencapsulate the Root Disk

Perform this procedure to unencapsulate the root disk.

1. **Ensure that only Solaris root file systems—root (/), swap, the global devices namespace, /usr, /var, /opt, and /home—are present on the root disk.**

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 you intend to unencapsulate.**

3. **Evacuate all resource groups and device groups from the node.**

```
scswitch -S -h node
```

-S                      Evacuates all resource groups and device groups

-h *node*                Specifies the name of the node from which to evacuate resource or device groups

4. **Determine the node ID number of the node.**

```
clinfo -n
N
```

5. **Unmount the global-devices file system for this node, where *N* is the node ID number 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 fsck mount mount
#to mount to fsck point type pass at boot 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 `rootdg` disk group.

```
vxedit -rf rm rootdiskxNvol
```



---

**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 related to global devices entries is restored after the root disk is unencapsulated.

---

8. Unencapsulate the root disk.

```
/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 you created in Step 9.

```
newfs /dev/rdisk/cNtXdYsZ
```

11. Determine the device ID (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 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 and Solstice DiskSuite/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.**



# Index

---

## A

### adapters

#### Ethernet

local-mac-address? variable changes during upgrade, 117

local-mac-address? variable configuration, 40

local-mac-address? variable requirements, 21

#### IP Network Multipathing group

requirements, 22

removing bad configuration information, 46

#### SCI-PCI

configuring on additional nodes, 58

configuring on the first node, 49

installing Solaris packages, 43

installing Solaris packages with JumpStart, 82

installing Sun Cluster packages, 45, 54, 69, 84

installing Sun Cluster packages with JumpStart, 83

package requirements, 13

port names, 50, 80

test IP addresses, 18

### adding

*See also* installing

disk drives to a diskset, 149

mediator hosts, 156

nodes to the Sun Cluster module to Sun Management Center, 105

### administrative console

IP addresses, 18

### administrative console (Continued)

MANPATH, 38

PATH, 38

alternate boot path, displaying, 135

### Apache

installing packages, 64

modifying links during upgrade, 114

### authentication list

adding nodes during Sun Cluster

installation, 92

adding nodes during VxVM installation, 162

autoscinstall.class file, 81

## B

### boot devices, determining the alternate boot

path, 135

## C

cconsole command, 39

CCP, *See* Cluster Control Panel (CCP)

ccp command, 38

.cdtoc file, requirements for SunPlex Manager installation, 72

checking, *See* verifying

class file, modifying, 81

### Cluster Control Panel (CCP)

installing, 37

starting, 39

- cluster file systems
  - caution notice, 94
  - checking the configuration, 97
  - configuring, 94
  - planning, 23
  - required mount options, 96
- cluster interconnects
  - planning, 20
  - removing bad configuration information, 46
- cluster name, 19
- cluster nodes
  - adding to the Sun Cluster module to Sun Management Center, 105
  - determining the ID number, 174
  - installing
    - additional nodes, 53
    - by using JumpStart, 74
    - by using SunPlex Manager, 68
    - first node, 44
  - planning, 19
  - upgrading, 110
- clusters file, administrative console, 38
- configuring
  - cluster file systems, 94
  - disksets, 147
  - IP Network Multipathing groups, 98
  - md.tab file, 152
  - name service switch, 87
  - Network Time Protocol (NTP), 100
  - ntp.conf file, 100
  - quorum devices, 91
  - Solaris Volume Manager, 126
  - Solstice DiskSuite, 126
  - state database replicas, 132
  - user accounts for SunPlex Manager, 67
  - user work environment, 88
  - VERITAS Volume Manager (VxVM), 159
- console-access devices
  - IP addresses, 18
  - planning, 18
  - serial port numbers, 38
- creating, *See* configuring

## D

- data services
  - installing by using `scinstall`, 89

- data services (Continued)
  - installing by using SunPlex Manager, 64
  - upgrading, 117
- deporting disk device groups, 171
- device groups, *See* disk device groups
- device IDs (DID)
  - determining names, 90
  - displaying names, 142
- DID, *See* device IDs (DID)
- Dirty Region Logging (DRL), planning, 28
- disabling install mode, 91
- disk device groups
  - evacuating, 170
  - importing and deporting, 171
  - node lists
    - removing nodes from, 166
    - viewing, 165
  - planning, 21
  - registering disk groups as, 172
  - reminoring, 173
  - verifying the configuration, 174
- disk drives
  - adding to disksets, 149
  - repartitioning, 151
- disk groups
  - See also* disk device groups
  - configuring, 171
- disk strings, mediator requirements, 155
- disksets
  - adding disk drives, 149
  - configuring, 147
  - configuring the maximum allowed, 130
  - planning the maximum number, 26
  - repartitioning disk drives, 151
- domain console network interfaces, IP addresses, 18
- drives, *See* disk drives
- DRL, *See* Dirty Region Logging (DRL)
- Dynamic Multipathing (DMP), 163, 168

## E

- enabling the kernel cage, 44
- encapsulated root disks
  - configuring, 161
  - mirroring, 165
  - planning, 27

- Enclosure-Based Naming, planning, 27
- error messages
  - sccheck command, 47, 56, 73
  - xntpd daemon, 51
- /etc/clusters file, administrative console, 38
- /etc/inet/hosts file, 18, 43
  - JumpStart, 83
- /etc/inet/ntp.conf file
  - changes during upgrade, 117
  - configuring, 100
  - default, 51
- /etc/init.d/xntpd.cluster start command, 102
- /etc/init.d/xntpd start command, 102
- /etc/lvm/md.tab file, 152
- /etc/name\_to\_major file
  - on non-VxVM nodes, 43, 164, 169
  - on VxVM-installed nodes, 163, 168
- /etc/nsswitch.conf file, 87
- /etc/release file, 43
- /etc/serialports file, 38
- /etc/system file
  - kernel\_cage\_enable variable, 44
  - stack size setting, 94
- /etc/user\_attr file, adding users to, 67
- /etc/vfstab file
  - adding mount points, 96
  - checking the configuration, 97
  - modifying during upgrade, 115
- Ethernet adapters
  - local-mac-address? variable
    - changes during upgrade, 117
    - configuring, 40
    - requirements, 21

## F

- file-system logging, planning, 28
- finish script, JumpStart, 84

## G

- global devices
  - caution notice, 175

- global devices (Continued)
  - /global/.devices directory
    - mirroring, 137
    - node@nodeid file system, 25
  - /globaldevices partition
    - creating, 41
    - planning, 14
  - planning, 23
  - updating the namespace, 148
  - /global directory, 24
- global file systems, *See* cluster file systems

## H

- help, 9
- hosts file, 18, 43
  - JumpStart, 83
- hot spare disks, planning, 25

## I

- importing disk device groups, 171
- initialization files, configuring the user work environment, 88
- install mode
  - disabling after installation, 91
  - verifying, 92
- installing
  - See also* adding
  - Apache packages, 64
  - caution when rebooting the first-installed node, 55
  - Cluster Control Panel (CCP), 37
  - data services
    - by using scinstall, 89
    - by using SunPlex Manager, 64
  - RSMAPI
    - Solaris packages, 43
    - Solaris packages with JumpStart, 82
    - Sun Cluster packages, 45, 54, 69
    - Sun Cluster packages with JumpStart, 83, 84
  - Solaris
    - alone, 39
    - with Sun Cluster, 74
  - Solstice DiskSuite, 126

- installing, Solstice DiskSuite (Continued)
  - by using SunPlex Manager, 64
  - from the Solaris CD-ROM, 129
- Sun Cluster, 34
  - additional nodes, 53
  - by using JumpStart, 74
  - by using SunPlex Manager, 64
  - first node, 44
  - log files, 51
  - status, 73
  - verifying, 90
- Sun Cluster module to Sun Management Center, 102, 103
- SunPlex Manager, 64
  - log files, 73
- VERITAS File System (VxFS), 94
- VERITAS Volume Manager (VxVM), 159
  - and encapsulating the root disk, 161
  - without encapsulating the root disk, 167
- interconnect, *See* cluster interconnect
- IP addresses, planning, 17
- IP Network Multipathing groups
  - configuring, 98
  - planning, 22
  - test IP address requirements, 22
- IPMP, *See* IP Network Multipathing groups

## J

- JumpStart
  - class file, 81
  - /etc/inet/hosts file, 83
  - finish script, 84
  - installing Solaris and Sun Cluster, 74
- junctions
  - See* transport junctions

## K

- kernel\_cage\_enable variable, 44
- /kernel/drv/md.conf file, 26
  - caution notice, 27, 131
  - configuring, 130

## L

- licenses, planning, 17
- loading the Sun Cluster module to Sun Management Center, 106
- local-mac-address? variable
  - changes during upgrade, 117
  - configuring, 40
  - requirements, 21
- localonly property, enabling, 166
- log files
  - Sun Cluster installation, 51
  - SunPlex Manager installation, 73
- logging
  - configuring, 96
  - planning, 28
- logical addresses, planning, 18

## M

- MANPATH
  - administrative console, 38
  - cluster nodes, 88
- md.conf file
  - caution notice, 131
  - configuring, 130
  - planning, 26
- md\_nsets field
  - configuring, 130
  - planning, 26
- md.tab file, configuring, 152
- mediators
  - adding hosts, 156
  - checking status, 156
  - overview, 155
  - planning, 26
  - repairing data, 157
- metadevices
  - activating, 154
  - configuring the maximum allowed, 130
  - planning the maximum number, 26
- minor number conflicts, repairing, 173
- mirroring
  - multihost disks, 29
  - planning, 29
  - root disks, 133
    - caution notice, 165
    - planning, 29

mount options, requirements for cluster file systems, 96

mount points

cluster file systems, 23

modifying the `/etc/vfstab` file, 96

multihost disks, mirroring, 29

multiported disks, planning, 25

## N

NAFO groups

*See also* IP Network Multipathing groups

upgrade to IP Network Multipathing groups, 117

name service switch, configuring, 87

`name_to_major` file

on non-VxVM nodes, 43, 164, 169

on VxVM-installed nodes, 168

Network Time Protocol (NTP), configuring, 100

`nmd` field

configuring, 130

planning, 26

node lists

planning, 25

removing nodes from, 166

viewing, 165

nodes, *See* cluster nodes

non-cluster mode

rebooting into, 93

rebooting into single-user, 115

`nsswitch.conf` file, configuring, 87

NTP, *See* Network Time Protocol (NTP)

`ntp.conf` file

changes during upgrade, 117

configuring, 100

default, 51

## O

online help, Sun Cluster module to Sun Management Center, 107

`/opt/SUNWcluster/bin/ccp` command, 38

`/opt/SUNWcluster/bin` directory, 38

`/opt/SUNWcluster/man` directory, 38

## P

partitions

`/globaldevices`, 14, 41

repartitioning disk drives, 151

`/sds`, 41

volume manager, 14

patches, planning, 17

`PATH`

administrative console, 38

cluster nodes, 88

PCI adapters, *See* PCI-PCI adapters

ports, *See* serial ports

private hostnames

changing, 99

planning, 20

verifying, 100

private network, planning, 19

profile, JumpStart, 81

public networks, planning, 21

## Q

quorum devices

and mirroring, 30

caution notice, 165

initial setup, 91

planning, 22

## R

RBAC authorization, adding to user

accounts, 66

rebooting

caution notice, 55

into non-cluster mode, 93

into single-user non-cluster mode, 115

registering, VxVM disk device groups, 172

release file, 43

Remote Shared Memory Application

Programming Interface (RSM-API)

installing

Solaris packages, 43

Solaris packages with JumpStart, 82

Sun Cluster packages, 45, 54, 69

Sun Cluster packages with JumpStart, 83, 84

- Remote Shared Memory Application Programming Interface (RSMAPI) (Continued)
  - package requirements, 13
- removing Sun Cluster, 92
- repairing
  - mediator data, 157
  - minor number conflicts, 173
- replicas, configuring, 132
- resource groups, evacuating, 170
- resource types, reregistering after upgrade, 120
- root (/) file systems, mirroring, 133
- root disk groups (`rootdg`)
  - configuring
    - on encapsulated root disks, 161
    - on non-root disks, 170
  - planning, 27
- root disks
  - encapsulating, 161
  - mirroring, 133
    - caution notice, 165
    - planning, 29
  - unencapsulating, 174
- root environment, configuring, 88
- `rpcmod` settings, correcting after VxFS installation, 94
- RSMAPI, *See* Remote Shared Memory Application Programming Interface (RSMAPI)

## S

- `sccheck` command
  - preinstallation check, 47, 56, 73
  - `vfstab` file check, 97
- `scconf` command
  - adding nodes to the authentication list, 92
  - enabling the `localonly` property, 134, 138, 141, 145
  - removing bad interconnect information, 46
  - removing nodes from a node list, 135, 138, 142, 145
  - verifying install mode, 92
  - viewing private hostnames, 100
- `scdidadm` command
  - determining device ID (DID) names, 90
  - displaying device ID (DID) names, 142

- `scgdevs` command, updating the `global-devices` namespace, 148
- SCI-PCI adapters
  - configuring
    - on additional nodes, 58
    - on the first node, 49
  - installing
    - Solaris packages, 43
    - Solaris packages with JumpStart, 82
    - Sun Cluster packages, 45, 54, 69
    - Sun Cluster packages with JumpStart, 83, 84
  - package requirements
    - Solaris, 13
  - port names, 50, 80
- `scinstall` command
  - installing additional nodes, 53
  - installing data services, 89
  - installing the first node, 44
  - JumpStart installation, 74
  - uninstalling Sun Cluster, 92
  - upgrading, 116
- `scinstall.log.pid` file, 51
- `scsetup` command
  - changing private hostnames, 99
  - post-installation setup, 91
  - registering disk device groups, 172
- `scstat` command, verifying disk group configurations, 174
- `scswitch` command, evacuating resource groups and device groups, 170
- `scvxinstall` command
  - installing VxVM and encapsulating the root disk, 161
  - installing VxVM only, 167
- `/sds` partition, 41
- secondary root disks, 30
- serial ports
  - configuring on the administrative console, 38
  - Simple Network Management Protocol (SNMP), 103
- `serialports` file, 38
- Simple Network Management Protocol (SNMP), port for Sun Management Center, 103
- single-user non-cluster mode, rebooting into, 115
- SNMP, port for Sun Management Center, 103

- Solaris
  - installing
    - alone, 39
    - with Sun Cluster, 74
  - planning, 13
    - /globaldevices file system, 15
    - partitions, 14
    - root (/) file system, 14
    - software groups, 13
    - swap, 15
    - volume managers, 15
  - upgrading, 113
  - version, 43
- Solaris Volume Manager
  - coexistence with VERITAS Volume Manager, 164, 169
  - configuring, 126
  - disksets
    - adding disk drives, 149
    - configuring, 147
    - configuring the maximum allowed, 130
    - repartitioning disk drives, 151
  - MANPATH, 88
  - md.tab file, 152
  - mediators
    - adding hosts, 156
    - checking status, 156
    - overview, 155
    - repairing bad data, 157
  - mirroring
    - global namespace, 137
    - root (/) file system, 133
    - root disks, 133
  - planning, 26
  - state database replicas, 132
  - transactional-volume logging
    - planning, 28
  - volumes
    - activating, 154
    - configuring the maximum allowed, 130
    - planning the maximum number, 26
- Solstice DiskSuite
  - coexistence with VERITAS Volume Manager, 164, 169
  - configuring, 126
  - disksets
    - adding disk drives, 149
    - configuring, 147
- Solstice DiskSuite, disksets (Continued)
  - configuring the maximum allowed, 130
  - repartitioning disk drives, 151
  - installing, 126
    - by using SunPlex Manager, 64
    - from the Solaris CD-ROM, 129
  - MANPATH, 88
  - md.tab file, 152
  - mediators
    - adding hosts, 156
    - checking status, 156
    - overview, 155
    - repairing bad data, 157
  - metadevices
    - activating, 154
    - configuring the maximum allowed, 130
    - planning the maximum number, 26
  - mirroring
    - global namespace, 137
    - root (/) file system, 133
    - root disks, 133
  - planning, 26
  - state database replicas, 132
  - trans-metadevice logging
    - planning, 28
  - spm directory, log files, 73
  - SSP, *See* console-access devices
  - stack size setting, correcting after VxFS
    - installation, 94
  - starting Sun Management Center, 104
  - state database replicas, configuring, 132
  - status
    - mediators, 156
    - Sun Cluster
      - installation logs, 51
      - SunPlex Manager installation log, 73
      - verifying, 90
- Sun Cluster module to Sun Management Center, 102
  - adding nodes, 105
  - installing, 103
  - loading, 106
  - requirements, 103
- Sun Enterprise 10000 servers
  - dynamic reconfiguration support, 44
  - kernel\_cage\_enable variable, 44
  - serialports file, 38

- Sun Fire 15000 servers
  - IP addresses, 18
  - serial port numbers, 38
- Sun Management Center
  - requirements, 103
  - starting, 104
- Sun Cluster module, 102
  - adding nodes, 105
  - installing, 103
  - loading, 106
  - online help, 107
  - upgrading, 120
- upgrading, 122
- SunPlex Manager, 62
  - adding RBAC authorization, 66
  - configuring new user accounts, 67
  - installing, 64
  - using to install software, 68
- swap, 14
- system controllers (SC), *See* console-access devices
- system file
  - kernel\_cage\_enable variable, 44
  - stack size setting, 94
- System Service Processor (SSP), *See* console-access devices

## T

- technical support, 9
- telnet command, serial port numbers, 38
- terminal concentrators (TC), *See* console-access devices
- test IP addresses, requirements, 22
- three-way mirroring, 29
- transport adapters, *See* adapters
- transport junctions
  - planning, 20
  - port names, 50, 80
  - planning, 20
  - removing bad configuration information, 46

## U

- UFS logging
  - configuring, 96

- UFS logging (Continued)
  - planning, 28
- unencapsulating the root disk, 174
- uninstalling Sun Cluster, 92
- upgrading, 110
  - data services, 117
  - guidelines for, 110
  - resource types, 120
  - Solaris, 113
  - Sun Cluster module to Sun Management Center, 120
  - Sun Management Center, 122
  - verifying, 118
  - VERITAS Volume Manager (VxVM), 114
- user accounts, configuring for SunPlex Manager, 67
- user\_attr file, adding users to, 67
- user work environment, modifying user initialization files, 88
  - /usr/cluster/bin directory, 88
  - /usr/cluster/bin/scconf command, adding nodes to the authentication list, 92
  - /usr/cluster/man directory, 88

## V

- /var/cluster/logs/install/scinstall.log.pid file, 51
- /var/cluster/spm directory, 73
- verifying
  - install mode, 92
  - preconfiguration requirements, 47
  - private hostnames, 100
  - upgrade, 118
- VERITAS File System (VxFS)
  - administering, 97
  - installing, 94
  - logging
    - configuring, 96
    - planning, 28
  - MANPATH, 88
  - mounting cluster file systems, 24, 97
  - PATH, 88
- VERITAS Volume Manager (VxVM)
  - and Dynamic Multipathing (DMP), 168
  - cluster feature, 171



VERITAS Volume Manager (VxVM)  
 (Continued)

- configuring
  - disk groups, 171
  - non-VxVM nodes, 164, 169
  - volumes, 171
- disk device groups
  - importing and deporting, 171
  - reminoring, 173
- disk groups
  - registering as disk device groups, 172
- Dynamic Multipathing (DMP), 163
- encapsulation, 27
- Enclosure-Based Naming, 27
- installing, 159
  - and encapsulating the root disk, 161
  - VxVM only, 167
- MANPATH, 88
- mirroring the encapsulated root disk, 165
- PATH, 88
- planning, 15, 27
- removing man pages, 163, 168
- root disk groups (`rootdg`)
  - configuring on non-root disks, 170
  - planning, 27, 160
- unencapsulating the root disk, 174
  - caution notice, 175
- upgrading, 114
- verifying disk group configurations, 174
- `vfstab` file
  - adding mount points, 96
  - checking the configuration, 97
  - modifying during upgrade, 115
- `vold` daemon, 64
- volume managers
  - partitions, 14
  - planning
    - general, 24
    - Solaris Volume Manager, 26
    - Solstice DiskSuite, 26
    - VERITAS Volume Manager, 27
- volumes
  - Solaris Volume Manager
    - activating, 154
    - configuring the maximum allowed, 130
    - planning the maximum number, 26
  - VERITAS Volume Manager (VxVM)
    - configuring, 171

VxFS, *See* VERITAS File System (VxFS)

`vxio` driver major number
 

- configuring on non-VxVM nodes, 164, 169
- configuring on VxVM-installed nodes, 163, 168

VxVM, *See* VERITAS Volume Manager (VxVM)

## X

`xntpd.cluster` start command, 102

`xntpd` daemon, error messages, 51

`xntpd` start command, 102

