



Sun Cluster Software Installation Guide for Solaris OS



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

Part No: 819-2970-10
December 2006, Revision A

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

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

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

This distribution may include materials developed by third parties.

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

Sun, Sun Microsystems, the Sun logo, the Solaris logo, the Java Coffee Cup logo, docs.sun.com, JumpStart, Sun Enterprise, Sun Fire, SunPlex, Sun StorEdge, Sun StorageTek, Java, and Solaris are trademarks or registered trademarks of Sun Microsystems, Inc. in the U.S. and other countries. All SPARC trademarks are used under license and are trademarks or registered trademarks of SPARC International, Inc. in the U.S. and other countries. Products bearing SPARC trademarks are based upon an architecture developed by Sun Microsystems, Inc. ORACLE is a registered trademark of Oracle Corporation.

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

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

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

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

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

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

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

Sun, Sun Microsystems, le logo Sun, le logo Solaris, le logo Java Coffee Cup, docs.sun.com, JumpStart, Sun Enterprise, Sun Fire, SunPlex, Sun StorEdge, Sun StorageTek, Java et Solaris sont des marques de fabrique ou des marques déposées de Sun Microsystems, Inc. aux États-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 États-Unis et dans d'autres pays. Les produits portant les marques SPARC sont basés sur une architecture développée par Sun Microsystems, Inc. ORACLE est une marque déposée registre de Oracle Corporation.

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

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

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

Contents

Preface	9
1 Planning the Sun Cluster Configuration	13
Finding Sun Cluster Installation Tasks	13
Planning the Solaris OS	14
Guidelines for Selecting Your Solaris Installation Method	15
Solaris OS Feature Restrictions	15
Solaris Software Group Considerations	16
System Disk Partitions	16
Guidelines for Non-Global Zones in a Cluster	19
Planning the Sun Cluster Environment	20
Licensing	20
Software Patches	20
Public Network IP Addresses	21
Console-Access Devices	21
Logical Addresses	22
Public Networks	22
Guidelines for NFS	23
Service Restrictions	24
Sun Cluster Configurable Components	25
Planning the Global Devices, Device Groups, and Cluster File Systems	31
Global Devices	31
Device Groups	31
Cluster File Systems	32
Choosing Mount Options for Cluster File Systems	33
Mount Information for Cluster File Systems	35
Planning Volume Management	36
Guidelines for Volume-Manager Software	37
Guidelines for Solaris Volume Manager Software	38

Guidelines for VERITAS Volume Manager Software	39
File-System Logging	41
Mirroring Guidelines	42
2 Installing Software on the Cluster	45
Installing the Software	45
▼ How to Prepare for Cluster Software Installation	46
▼ How to Install Cluster Control Panel Software on an Administrative Console	47
▼ How to Install Solaris Software	50
▼ How to Configure Internal Disk Mirroring	54
▼ How to Install Sun Multipathing Software	55
▼ SPARC: How to Install VERITAS File System Software	57
▼ How to Install Sun Cluster Framework and Data-Service Software Packages	58
▼ How to Set Up the Root Environment	61
3 Establishing the Cluster	63
Establishing a New Cluster or New Cluster Node	63
▼ How to Configure Sun Cluster Software on All Nodes (scinstall)	64
▼ How to Configure Sun Cluster Software on All Nodes (XML)	72
▼ How to Install Solaris and Sun Cluster Software (JumpStart)	79
▼ How to Prepare the Cluster for Additional Cluster Nodes	95
▼ How to Change the Private Network Configuration When Adding Nodes or Private Networks	98
▼ How to Configure Sun Cluster Software on Additional Cluster Nodes (scinstall)	103
▼ How to Configure Sun Cluster Software on Additional Cluster Nodes (XML)	110
▼ How to Update Quorum Devices After Adding a Node to a Cluster	114
▼ How to Configure Quorum Devices	117
▼ How to Verify the Quorum Configuration and Installation Mode	121
▼ How to Change Private Hostnames	122
▼ How to Configure Network Time Protocol (NTP)	123
4 Configuring Solaris Volume Manager Software	127
Configuring Solaris Volume Manager Software	127
▼ SPARC: How to Set the Number of Volume Names and Disk Sets	128
▼ How to Create State Database Replicas	129
Mirroring the Root Disk	130

▼ How to Mirror the Root (/) File System	131
▼ How to Mirror the Global Devices Namespace	134
▼ How to Mirror File Systems Other Than Root (/) That Cannot Be Unmounted	138
▼ How to Mirror File Systems That Can Be Unmounted	142
Creating Disk Sets in a Cluster	146
▼ How to Create a Disk Set	147
Adding Drives to a Disk Set	150
▼ How to Add Drives to a Disk Set	150
▼ How to Repartition Drives in a Disk Set	152
▼ How to Create an md . tab File	152
▼ How to Activate Volumes	154
Configuring Dual-String Mediators	155
Requirements for Dual-String Mediators	156
▼ How to Add Mediator Hosts	156
▼ How to Check the Status of Mediator Data	157
▼ How to Fix Bad Mediator Data	157
5 Installing and Configuring VERITAS Volume Manager	159
Installing and Configuring VxVM Software	159
Setting Up a Root Disk Group Overview	160
▼ How to Install VERITAS Volume Manager Software	160
▼ SPARC: How to Encapsulate the Root Disk	162
▼ How to Create a Root Disk Group on a Nonroot Disk	163
▼ How to Mirror the Encapsulated Root Disk	164
Creating Disk Groups in a Cluster	166
▼ How to Create a Disk Group	167
▼ How to Register a Disk Group	168
▼ How to Assign a New Minor Number to a Device Group	169
▼ How to Verify the Disk Group Configuration	170
Unencapsulating the Root Disk	172
▼ How to Unencapsulate the Root Disk	172
6 Creating Cluster File Systems and Non-Global Zones	175
Creating Cluster File Systems	175
▼ How to Create Cluster File Systems	175
Configuring a Non-Global Zone on a Cluster Node	179

- ▼ How to Create a Non-Global Zone on a Cluster Node 179

- 7 Installing and Upgrading the Sun Cluster Module to Sun Management Center 183**
 - SPARC: Installing the Sun Cluster Module for Sun Management Center 183
 - SPARC: Installation Requirements for Sun Cluster Monitoring 184
 - ▼ SPARC: How to Install the Sun Cluster Module for Sun Management Center 185
 - ▼ SPARC: How to Start Sun Management Center 186
 - ▼ SPARC: How to Add a Cluster Node as a Sun Management Center Agent Host Object 186
 - ▼ SPARC: How to Load the Sun Cluster Module 187
 - SPARC: Upgrading Sun Management Center Software 188
 - ▼ SPARC: How to Upgrade Sun Cluster Module Software for Sun Management Center 189
 - ▼ SPARC: How to Upgrade Sun Management Center Software 190

- 8 Upgrading Sun Cluster Software 193**
 - Upgrade Requirements and Software Support Guidelines 193
 - Choosing a Sun Cluster Upgrade Method 195
 - Performing a Standard Upgrade to Sun Cluster 3.2 Software 196
 - ▼ How to Prepare the Cluster for Upgrade (Standard) 197
 - ▼ How to Upgrade the Solaris OS and Volume Manager Software (Standard) 202
 - ▼ How to Upgrade Sun Cluster 3.2 Software (Standard) 207
 - Performing a Dual-Partition Upgrade to Sun Cluster 3.2 Software 214
 - ▼ How to Prepare the Cluster for Upgrade (Dual-Partition) 215
 - ▼ How to Upgrade the Solaris OS and Volume Manager Software (Dual-Partition) 224
 - ▼ How to Upgrade Sun Cluster 3.2 Software (Dual-Partition) 229
 - Performing a Live Upgrade to Sun Cluster 3.2 Software 236
 - ▼ How to Prepare the Cluster for Upgrade (Live Upgrade) 237
 - ▼ How to Upgrade the Solaris OS and Sun Cluster 3.2 Software (Live Upgrade) 239
 - Completing the Upgrade 249
 - ▼ How to Verify Upgrade of Sun Cluster 3.2 Software 249
 - ▼ How to Finish Upgrade to Sun Cluster 3.2 Software 251
 - Recovering From an Incomplete Upgrade 257
 - ▼ How to Recover from a Failed Dual-Partition Upgrade 258
 - Recovering From Storage Configuration Changes During Upgrade 260
 - ▼ How to Handle Storage Reconfiguration During an Upgrade 260
 - ▼ How to Resolve Mistaken Storage Changes During an Upgrade 261

9	Uninstalling Software From the Cluster	263
	Uninstalling the Software	263
	▼ How to Unconfigure Sun Cluster Software to Correct Installation Problems	263
	▼ How to Remove Sun Cluster Information From a JumpStart Install Server	267
	▼ How to Uninstall SunPlex Manager Software	268
	▼ How to Uninstall the SUNWscrdt Package	269
	▼ How to Unload the RSMRDT Driver Manually	270
A	Sun Cluster Installation and Configuration Worksheets	273
	Installation and Configuration Worksheets	274
	Local File System Layout Worksheet	275
	Local Devices Worksheets	277
	Device Group Configurations Worksheet	279
	Volume-Manager Configurations Worksheet	281
	Volumes Worksheet (Solaris Volume Manager)	283
	Index	285

Preface

The *Sun Cluster Software Installation Guide for Solaris OS* contains guidelines for planning a Sun™ Cluster configuration, and provides procedures for installing, configuring, and upgrading the Sun Cluster software on both SPARC® based systems and x86 based systems.

Note – This Sun Cluster release supports systems that use the SPARC and x86 families of processor architectures: UltraSPARC, SPARC64, and AMD64. In this document, the label x86 refers to systems that use the AMD64 family of processor architectures.

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

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

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

Using UNIX Commands

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

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

- Online documentation for the Solaris OS
- Other software documentation that you received with your system
- Solaris OS man pages

Typographic Conventions

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

TABLE P-1 Typographic Conventions

Typeface	Meaning	Example
AaBbCc123	The names of commands, files, and directories, and onscreen computer output	Edit your <code>.login</code> file. Use <code>ls -a</code> to list all files. <code>machine_name%</code> you have mail.
AaBbCc123	What you type, contrasted with onscreen computer output	<code>machine_name%</code> su Password:
<i>aabcc123</i>	Placeholder: replace with a real name or value	The command to remove a file is <i>rm filename</i> .
<i>AaBbCc123</i>	Book titles, new terms, and terms to be emphasized	Read Chapter 6 in the <i>User's Guide</i> . <i>A cache</i> is a copy that is stored locally. Do <i>not</i> save the file. Note: Some emphasized items appear bold online.

Shell Prompts in Command Examples

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

TABLE P-2 Shell Prompts

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

Related Documentation

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

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

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

Related Third-Party Web Site References

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

Documentation, Support, and Training

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

- [Documentation](http://www.sun.com/documentation/) (<http://www.sun.com/documentation/>)
- [Support](http://www.sun.com/support/) (<http://www.sun.com/support/>)
- [Training](http://www.sun.com/training/) (<http://www.sun.com/training/>)

Getting Help

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

- Your name and email address (if available)
- Your company name, address, and phone number
- The model number and serial number of your systems
- The release number of the Solaris OS (for example, Solaris 10)
- The release number of Sun Cluster (for example, Sun Cluster 3.2)

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

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

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

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:

- “Finding Sun Cluster Installation Tasks” on page 13
- “Planning the Solaris OS” on page 14
- “Planning the Sun Cluster Environment” on page 20
- “Planning the Global Devices, Device Groups, and Cluster File Systems” on page 31
- “Planning Volume Management” on page 36

Finding Sun Cluster Installation Tasks

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

TABLE 1-1 Sun Cluster Software Installation Task Information

Task	Instructions
Set up cluster hardware.	<i>Sun Cluster 3.1 - 3.2 Hardware Administration Manual for Solaris OS</i> Documentation that shipped with your server and storage devices
Plan cluster software installation.	Chapter 1 “Installation and Configuration Worksheets” on page 274
Install software packages. Optionally, install and configure Sun StorEdge™ QFS software.	“Installing the Software” on page 45 <i>Sun StorEdge QFS Installation and Upgrade Guide</i>
Establish a new cluster or a new cluster node.	“Establishing a New Cluster or New Cluster Node” on page 63

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

Task	Instructions
Configure Solaris Volume Manager software.	“Configuring Solaris Volume Manager Software” on page 127 Solaris Volume Manager documentation
Install and configure VERITAS Volume Manager (VxVM) software.	“Installing and Configuring VxVM Software” on page 159 VxVM documentation
Configure cluster file systems, if used.	“How to Create Cluster File Systems” on page 175
<i>(Optional)</i> On the Solaris 10 OS, create non-global zones.	“How to Create a Non-Global Zone on a Cluster Node” on page 179
<i>(Optional)</i> SPARC: Install and configure the Sun Cluster module to Sun Management Center.	“SPARC: Installing the Sun Cluster Module for Sun Management Center” on page 183 Sun Management Center documentation
Plan, install, and configure resource groups and data services. Create highly available local file systems, if used.	<i>Sun Cluster Data Services Planning and Administration Guide for Solaris OS</i>
Develop custom data services.	<i>Sun Cluster Data Services Developer’s Guide for Solaris OS</i>
Upgrade to Sun Cluster 3.2 software.	Chapter 8 “Configuring Solaris Volume Manager Software” on page 127 or “Installing and Configuring VxVM Software” on page 159 Volume manager documentation “SPARC: How to Upgrade Sun Cluster Module Software for Sun Management Center” on page 189

Planning the Solaris OS

This section provides the following guidelines for planning Solaris software installation in a cluster configuration.

- “Guidelines for Selecting Your Solaris Installation Method” on page 15
- “Solaris OS Feature Restrictions” on page 15
- “Solaris Software Group Considerations” on page 16
- “System Disk Partitions” on page 16
- “Guidelines for Non-Global Zones in a Cluster” on page 19

For more information about Solaris software, see your Solaris installation documentation.

Guidelines for Selecting Your Solaris Installation Method

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

See “[How to Install Solaris and Sun Cluster Software \(JumpStart\)](#)” on page 79 for details about the `scinstall` JumpStart installation method. See your Solaris installation documentation for details about standard Solaris installation methods.

Solaris OS Feature Restrictions

Consider the following points when you plan the use of the Solaris OS in a Sun Cluster configuration:

- **Solaris 10 Zones** - Install Sun Cluster 3.2 framework software *only* in the global zone. To determine whether you can install a Sun Cluster data service directly in a non-global zone, see the documentation for that data service.

If you configure non-global zones on a cluster node, the loopback file system (LOFS) must be enabled. See the information for LOFS for additional considerations.

- **Loopback file system (LOFS)** - During cluster creation with the Solaris 9 version of Sun Cluster software, LOFS capability is disabled by default. During cluster creation with the Solaris 10 version of Sun Cluster software, LOFS capability is not disabled by default.

If the cluster meets both of the following conditions, you must disable LOFS to avoid switchover problems or other failures:

- Sun Cluster HA for NFS is configured on a highly available local file system.
- The `automountd` daemon is running.

If the cluster meets only one of these conditions, you can safely enable LOFS.

If you require both LOFS and the `automountd` daemon to be enabled, exclude from the automounter map all files that are part of the highly available local file system that is exported by Sun Cluster HA for NFS.

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

- **IP Filter** - Sun Cluster software does not support filtering with Solaris IP Filter. The use of the STREAMS autopush(1M) mechanism by Solaris IP Filter conflicts with Sun Cluster software's use of the mechanism.

Solaris Software Group Considerations

Sun Cluster 3.2 software requires at least the End User Solaris Software Group. However, other components of your cluster configuration might have their own Solaris software requirements as well. Consider the following information when you decide which Solaris software group you are installing.

- **Servers** - Check your server documentation for any Solaris software requirements. For example, Sun Enterprise™ 10000 servers require the Entire Solaris Software Group Plus OEM Support.
- **SCI-PCI adapters** - To use SCI-PCI adapters, which are available for use in SPARC based clusters only, or the Remote Shared Memory Application Programming Interface (RSMAPI), ensure that you install the RSMAPI software packages, which are SUNWrsm and SUNWrsmo, and for the Solaris 9 OS on SPARC based platforms also SUNWrsmx and SUNWrsmox. The RSMAPI software packages are included only in some Solaris software groups. For example, the Developer Solaris Software Group includes the RSMAPI software packages but the End User Solaris Software Group does not.

If the software group that you install does not include the RSMAPI software packages, install the RSMAPI software packages manually before you install Sun Cluster software. Use the pkgadd(1M) command to manually install the software packages. See the Section (3RSM) man pages for information about using the RSMAPI.

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

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

System Disk Partitions

Add this information to the appropriate [“Local File System Layout Worksheet”](#) on page 275.

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

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

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

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

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

See the following guidelines for additional partition planning information:

- “Guidelines for the Root (/) File System” on page 17
- “Guidelines for the `/globaldevices` File System” on page 18
- “Volume Manager Requirements” on page 18

Guidelines for the Root (/) File System

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

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

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

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

Guidelines for the `/globaldevices` File System

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

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

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

Volume Manager Requirements

If you use 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 Solaris Volume Manager software to function properly. See your Solaris Volume Manager documentation for more information.

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

Example - Sample File-System Allocations

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

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

TABLE 1-2 Example File-System Allocation

Slice	Contents	Size Allocation	Description
0	/	6.75GB	Remaining free space on the disk after allocating space to slices 1 through 7. Used for the Solaris OS, Sun Cluster software, data-services software, volume-manager software, Sun Management Center agent and Sun Cluster module agent packages, root file systems, and database and application software.
1	swap	1GB	512 Mbytes for the Solaris OS. 512 Mbytes for Sun Cluster software.
2	overlap	8.43GB	The entire disk.
3	/globaldevices	512MB	The Sun Cluster software later assigns this slice a different mount point and mounts the slice as a cluster file system.
4	unused	-	Available as a free slice for encapsulating the root disk under VxVM.
5	unused	-	-
6	unused	-	-
7	volume manager	20MB	Used by Solaris Volume Manager software for the state database replica, or used by VxVM for installation after you free the slice.

Guidelines for Non-Global Zones in a Cluster

For information about the purpose and function of Solaris 10 zones in a cluster, see “Support for Solaris Zones on Sun Cluster Nodes” in *Sun Cluster Concepts Guide for Solaris OS*.

Consider the following points when you create a Solaris 10 non-global zone, simply referred to as a zone, on a cluster node.

- **Unique zone name** - The zone name must be unique in the node. Do not specify the same name to more than one zone on the same node.
- **Reusing a zone name on multiple nodes** - To simplify cluster administration, you can use the same name for a zone on each node where resource groups are to be brought online in that zone.
- **Private IP addresses** - Do not attempt to use more private IP addresses than are available in the cluster.
- **Mounts** - Do not include global mounts in zone definitions. Include only loopback mounts.
- **Failover services** - In multiple-node clusters, while Sun Cluster software permits you to specify different zones on the same node in a failover resource group’s node list, doing so is useful only during testing. If a single node hosts all zones in the node list, the node becomes a single point of failure for the resource group. For highest availability, zones in a failover resource group’s node list should be on different nodes.

In single-node clusters, there is no functional risk if you specify multiple zones in a failover resource group's node list.

- **Scalable services** - Do not create non-global zones for use in the same scalable service on the same node. Each instance of the scalable service must run on a different cluster node.
- **LOFS** - Solaris Zones requires that the loopback file system (LOFS) be enabled. However, the Sun Cluster HA for NFS data service requires that LOFS be disabled, to avoid switchover problems or other failures. If you configure both non-global zones and Sun Cluster HA for NFS in your cluster, do one of the following to prevent possible problems in the data service:
 - Disable the automountd daemon.
 - Exclude from the automounter map all files that are part of the highly available local file system that is exported by Sun Cluster HA for NFS.

Planning the Sun Cluster Environment

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

- [“Licensing” on page 20](#)
- [“Software Patches” on page 20](#)
- [“Public Network IP Addresses” on page 21](#)
- [“Console-Access Devices” on page 21](#)
- [“Logical Addresses” on page 22](#)
- [“Public Networks” on page 22](#)
- [“Guidelines for NFS” on page 23](#)
- [“Service Restrictions” on page 24](#)
- [“Sun Cluster Configurable Components” on page 25](#)

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

Licensing

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

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

Software Patches

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

- For information about current required patches, see “Patches and Required Firmware Levels” in *Sun Cluster 3.2 Release Notes for Solaris OS* or consult your Sun service provider.
- For general guidelines and procedures for applying patches, see Chapter 10, “Patching Sun Cluster Software and Firmware,” in *Sun Cluster System Administration Guide for Solaris OS*.

Public Network IP Addresses

For information about the use of public networks by the cluster, see “Public Network Adapters and Internet Protocol (IP) Network Multipathing” in *Sun Cluster Concepts Guide for Solaris OS*.

You must set up a number of public network 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 public network IP addresses assigned. Add these IP addresses to the following locations:

- Any naming services that are used
- The local `/etc/inet/hosts` file on each cluster node, after you install Solaris software
- For Solaris 10, the local `/etc/inet/ipnodes` file on each cluster node, after you install Solaris software

TABLE 1-3 Sun Cluster Components That Use Public Network IP Addresses

Component	Number of IP Addresses Needed
Administrative console	1 IP address per subnet.
Cluster nodes	1 IP address per node, per subnet.
Domain console network interface (Sun Fire™ 15000)	1 IP address per domain.
(Optional) Non-global zones	1 IP address per subnet.
Console-access device	1 IP address.
Logical addresses	1 IP address per logical host resource, per subnet.

For more information about planning IP addresses, see *System Administration Guide: IP Services (Solaris 9 or Solaris 10)*.

Console-Access Devices

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

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

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

Alternatively, if you connect an administrative console directly to cluster nodes or through a management network, you instead provide the hostname of each cluster node and its serial port number that is used to connect to the administrative console or the management network.

Logical Addresses

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

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

Public Networks

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

- **Separation of public and private network** - Public networks and the private network (cluster interconnect) must use separate adapters, or you must configure tagged VLAN on tagged-VLAN capable adapters and VLAN-capable switches to use the same adapter for both the private interconnect and the public network.
- **Minimum** - All cluster nodes must be connected to at least one public network. Public-network connections can use different subnets for different nodes.
- **Maximum** - You can have as many additional public-network connections as your hardware configuration allows.
- **Scalable services** - All nodes that run a scalable service must either use the same subnet or set of subnets or use different subnets that are routable among themselves.
- **IPv4** - Sun Cluster software supports IPv4 addresses on the public network.
- **IPv6** - Sun Cluster software supports IPv6 addresses on the public network under the following conditions or restrictions:
 - Sun Cluster software does not support IPv6 addresses on the public network if the private interconnect uses SCI adapters.
 - Sun Cluster software supports IPv6 addresses for both failover and scalable data services.

- **IPMP groups** - Each public network adapter that is used for data-service traffic must belong to an IP network multipathing (IPMP) group. If a public-network adapter is not used for data-service traffic, you do not have to configure it in an IPMP group.

In the Sun Cluster 3.2 release, the `scinstall` utility no longer automatically configures a single-adapter IPMP group on each unconfigured public-network adapter during Sun Cluster creation. Instead, the `scinstall` utility automatically configures a multiple-adapter IPMP group for each set of public-network adapters in the cluster that uses the same subnet. On the Solaris 10 OS, these groups are probe based. However, the `scinstall` utility ignores adapters that are already configured in an IPMP group. If any adapter in an IPMP group that the `scinstall` utility configures will not be used for data-service traffic, you can remove that adapter from the group.

For guidelines and instructions to configure IPMP groups, follow the procedures in Part VI, “IPMP,” in *System Administration Guide: IP Services*. To modify IPMP groups after cluster installation, follow the guidelines in “How to Administer IP Network Multipathing Groups in a Cluster” in *Sun Cluster System Administration Guide for Solaris OS* and procedures in “Administering IPMP (Tasks)” in *System Administration Guide: IP Services* (Solaris 9 or Solaris 10).

- **Local MAC address support** - All public network adapters must use network interface cards (NICs) that support local MAC address assignment. Local MAC address assignment is a requirement of IPMP.
- **local-mac-address setting** - The `local-mac-address?` variable must use the default value `true` for Ethernet adapters. Sun Cluster software does not support a `local-mac-address?` value of `false` for Ethernet adapters. This requirement is a change from Sun Cluster 3.0, which did require a `local-mac-address?` value of `false`.
- **IP Filter** - Sun Cluster configurations do not support filtering with Solaris IP Filter.

For more information about public-network interfaces, see *Sun Cluster Concepts Guide for Solaris OS*.

Guidelines for NFS

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

- **NFS client** - No Sun Cluster node can be an NFS client of a Sun Cluster HA for NFS-exported file system that is being mastered on a node in the same cluster. Such cross-mounting of Sun Cluster HA for NFS is prohibited. Use the cluster file system to share files among cluster nodes.
- **NFSv3 protocol** - If you are mounting file systems on the cluster nodes from external NFS servers, such as NAS filers, and you are using the NFSv3 protocol, you cannot run NFS client mounts and the Sun Cluster HA for NFS data service on the same cluster node. If you do, certain Sun Cluster HA for NFS data-service activities might cause the NFS daemons to stop and restart, interrupting NFS services. However, you can safely run the Sun Cluster HA for NFS data service if you use the NFSv4 protocol to mount external NFS file systems on the cluster nodes.

- **Locking** - Applications that run locally on the cluster must not lock files on a file system that is exported through NFS. Otherwise, local blocking (for example, `flock(3UCB)` or `fcntl(2)`) might interfere with the ability to restart the lock manager (`lockd(1M)`). During restart, a blocked local process might be granted a lock which might be intended to be reclaimed by a remote client. This would cause unpredictable behavior.
- **NFS security features** - Sun Cluster software does not support the following options of the `share_nfs(1M)` command:
 - `secure`
 - `sec=dh`

However, Sun Cluster software does support the following security features for NFS:

- The use of secure ports for NFS. You enable secure ports for NFS by adding the entry `set nfsrv:nfs_portmon=1` to the `/etc/system` file on cluster nodes.
- The use of Kerberos with NFS. For more information, see “Securing Sun Cluster HA for NFS With Kerberos V5” in *Sun Cluster Data Service for NFS Guide for Solaris OS*.

Service Restrictions

Observe the following service restrictions for Sun Cluster configurations:

- **Routers** - Do not configure cluster nodes as routers (gateways). If the system goes down, the clients cannot find an alternate router and cannot recover.
- **NIS+ servers** - Do not configure cluster nodes as NIS or NIS+ servers. There is no data service available for NIS or NIS+. However, cluster nodes can be NIS or NIS+ clients.
- **Boot and install servers** - Do not use a Sun Cluster configuration to provide a highly available boot or installation service on client systems.
- **RARP** - Do not use a Sun Cluster configuration to provide an `rarpd` service.
- **RPC program numbers** - If you install an RPC service on the cluster, the service must not use any of the following program numbers:
 - 100141
 - 100142
 - 100248

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

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

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

- Processes that run in the real-time scheduling class

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

Sun Cluster Configurable Components

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

- “Cluster Name” on page 25
- “Node Names” on page 25
- “Zone Names” on page 25
- “Private Network” on page 26
- “Private Hostnames” on page 27
- “Cluster Interconnect” on page 27
- “Quorum Devices” on page 29

Add this information to the appropriate configuration planning worksheet.

Cluster Name

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

Node Names

The cluster node name is the same name that you assign to the machine when you install it with the Solaris OS. See the `hosts(4)` man page for information about naming requirements.

In single-node cluster installations, the default cluster name is the name of the node.

During Sun Cluster configuration, you specify the names of all nodes that you are installing in the cluster.

Zone Names

On the Solaris 10 OS, use the naming convention `nodename:zonename` to specify a non-global zone to a Sun Cluster command.

- The *nodename* is the name of the cluster node.
- The *zonename* is the name that you assign to the non-global zone when you create the zone on the node. The zone name must be unique on the node. However, you can use the same zone name on different nodes, because the different node name in `nodename:zonename` makes the complete non-global zone name unique in the cluster.

To specify the global zone, you only need to specify the node name.

Private Network

Note – You do not need to configure a private network for a single-node cluster. The `scinstall` utility automatically assigns the default private-network address and netmask, even though a private network is not used by the cluster.

Sun Cluster software uses the private network for internal communication among nodes and among non-global zones that are managed by Sun Cluster software. A Sun Cluster configuration requires at least two connections to the cluster interconnect on the private network. When you configure Sun Cluster software on the first node of the cluster, you specify the private-network address and netmask in one of the following ways:

- Accept the default private-network address (172.16.0.0) and netmask (255.255.248.0). This IP address range supports a combined maximum of 64 nodes and non-global zones and a maximum of 10 private networks.

Note – The maximum number of nodes that an IP address range can support does not reflect the maximum number of nodes that the hardware configuration can support.

- Specify a different allowable private-network address and accept the default netmask.
- Accept the default private-network address and specify a different netmask.
- Specify both a different private-network address and a different netmask.

If you choose to specify a different netmask, the `scinstall` utility prompts you for the number of nodes and the number of private networks that you want the IP address range to support. The number of nodes that you specify should also include the expected number of non-global zones that will use the private network.

The utility calculates the netmask for the minimum IP address range that will support the number of nodes and private networks that you specified. The calculated netmask might support more than the supplied number of nodes, including non-global zones, and private networks. The `scinstall` utility also calculates a second netmask that would be the minimum to support twice the number of nodes and private networks. This second netmask would enable the cluster to accommodate future growth without the need to reconfigure the IP address range.

The utility then asks you what netmask to choose. You can specify either of the calculated netmasks or provide a different one. The netmask that you specify must minimally support the number of nodes and private networks that you specified to the utility.

Note – To change the private-network address and netmask after the cluster is established, see “How to Change the Private Network Address or Address Range of an Existing Cluster” in *Sun Cluster System Administration Guide for Solaris OS*. You must bring down the cluster to make these changes.

Changing the cluster private IP address range might be necessary to support the addition of nodes, non-global zones, or private networks.

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

- **Address and netmask sizes** - The private network address cannot be smaller than the netmask. For example, you can use a private network address of 172 . 16 . 10 . 0 with a netmask of 255 . 255 . 255 . 0. But you cannot use a private network address of 172 . 16 . 10 . 0 with a netmask of 255 . 255 . 0 . 0.
- **Acceptable addresses** The address must be included in the block of addresses that RFC 1918 reserves for use in private networks. You can contact the InterNIC to obtain copies of RFCs or view RFCs online at <http://www.rfcs.org>.
- **Use in multiple clusters** - You can use the same private network address in more than one cluster. Private IP network addresses are not accessible from outside the cluster.
- **IPv6** - Sun Cluster software does not support IPv6 addresses for the private interconnect. The system does configure IPv6 addresses on the private network adapters to support scalable services that use IPv6 addresses. But internode communication on the private network does not use these IPv6 addresses.

See “Planning Your TCP/IP Network (Tasks),” in *System Administration Guide: IP Services* (Solaris 9 or Solaris 10) for more information about private networks.

Private Hostnames

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

For the Solaris 10 OS, the creation of a private hostname for a non-global zone is optional. There is no required naming convention for the private hostname of a non-global zone.

Cluster Interconnect

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

- Between two transport adapters

- Between a transport adapter and a transport switch, also called a transport junction

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

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

During Sun Cluster configuration, you specify configuration information for two cluster interconnects. You can configure additional private-network connections after the cluster is established by using the `clsetup(1CL)` utility.

For guidelines about cluster interconnect hardware, see “Interconnect Requirements and Restrictions” in *Sun Cluster 3.1 - 3.2 Hardware Administration Manual for Solaris OS*. For general information about the cluster interconnect, see “Cluster-Interconnect Components” in *Sun Cluster Overview for Solaris OS* and *Sun Cluster Concepts Guide for Solaris OS*.

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 switch.

Consider the following guidelines and restrictions:

- **IPv6** - Sun Cluster software does not support IPv6 communications over the private interconnects.
- **Local MAC address assignment** - All private network adapters must use network interface cards (NICs) that support local MAC address assignment. Link-local IPv6 addresses, which are required on private network adapters to support IPv6 public network addresses, are derived from the local MAC addresses.
- **Tagged VLAN adapters** – Sun Cluster software supports tagged Virtual Local Area Networks (VLANs) to share an adapter between the private interconnect and the public network. To configure a tagged VLAN adapter for the private interconnect, specify the adapter name and its VLAN ID (VID) in one of the following ways:
 - Specify the usual adapter name, which is the device name plus the instance number or physical point of attachment (PPA). For example, the name of instance 2 of a Cassini Gigabit Ethernet adapter would be `ce2`. If the `scinstall` utility asks whether the adapter is part of a shared virtual LAN, answer **yes** and specify the adapter’s VID number.
 - Specify the adapter by its VLAN virtual device name. This name is composed of the adapter name plus the VLAN instance number. The VLAN instance number is derived from the formula $(1000 * V) + N$, where V is the VID number and N is the PPA.

As an example, for VID 73 on adapter `ce2`, the VLAN instance number would be calculated as $(1000 * 73) + 2$. You would therefore specify the adapter name as `ce73002` to indicate that it is part of a shared virtual LAN.

For information about configuring VLAN in a cluster, see “Configuring VLANs as Private Interconnect Networks” in *Sun Cluster 3.1 - 3.2 Hardware Administration Manual for Solaris OS*. For general information about VLAN, see *Solaris 9 9/05 Sun Hardware Platform Guide*.

- **SBus SCI adapters** – The SBus Scalable Coherent Interface (SCI) is not supported as a cluster interconnect. However, the SCI-PCI interface is supported.
- **Logical network interfaces** - Logical network interfaces are reserved for use by Sun Cluster software.

See the `scconf_trans_adap_*(1M)` family of man pages for information about a specific transport adapter.

Transport Switches

If you use transport switches, such as a network switch, specify a transport switch name for each interconnect. You can use the default name `switchN`, where *N* is a number that is automatically assigned during configuration, or create another name.

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

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

If your two-node cluster is direct connected, you can still specify a transport switch for the interconnect.

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

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

During Sun Cluster installation of a two-node cluster, you can choose to let the `scinstall` utility automatically configure a SCSI quorum device. This quorum device is chosen from the available shared SCSI storage disks. The `scinstall` utility assumes that all available shared SCSI storage disks are supported to be quorum devices.

If you want to use a quorum server or a Network Appliance NAS device as the quorum device, you configure it after `scinstall` processing is completed.

After installation, you can also configure additional quorum devices by using the `clsetup(1CL)` utility.

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

If your cluster configuration includes third-party shared storage devices that are not supported for use as quorum devices, you must use the `clsetup` utility to configure quorum manually.

Consider the following points when you plan quorum devices.

- **Minimum** – A two-node cluster must have at least one quorum device, which can be a shared SCSI disk, a quorum server, or a Network Appliance NAS device. For other topologies, quorum devices are optional.
- **Odd-number rule** – If more than one quorum device is configured in a two-node cluster, or in a pair of nodes directly connected to the quorum device, configure an odd number of quorum devices. This configuration ensures that the quorum devices have completely independent failure pathways.
- **Distribution of quorum votes** – For highest availability of the cluster, ensure that the total number of votes that are contributed by quorum devices is less than the total number of votes that are contributed by nodes. Otherwise, the nodes cannot form a cluster if all quorum devices are unavailable, even if all nodes are functioning.
- **Connection** – You must connect a quorum device to at least two nodes.
- **SCSI fencing protocol** – When a SCSI quorum device is configured, its SCSI protocol is automatically set to SCSI-2 in a two-node cluster or SCSI-3 in cluster with three or more nodes. You cannot change the SCSI protocol of a device after it is configured as a quorum device.
- **Replicated devices** – Sun Cluster software does not support replicated devices as quorum devices.
- **ZFS storage pools** – Do not add a configured quorum device to a Zettabyte File System (ZFS) storage pool. When a configured quorum device is added to a ZFS storage pool, the disk is relabeled as an EFI disk and quorum configuration information is lost. The disk can then no longer provide a quorum vote to the cluster.

Once a disk is in a storage pool, you can configure that disk as a quorum device. Or, you can unconfigure the quorum device, add it to the storage pool, then reconfigure the disk as a quorum device.

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

Planning the Global Devices, Device Groups, and Cluster File Systems

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

- “Global Devices” on page 31
- “Device Groups” on page 31
- “Cluster File Systems” on page 32
- “Choosing Mount Options for Cluster File Systems” on page 33
- “Mount Information for Cluster File Systems” on page 35

Global Devices

For information about the purpose and function of global devices, see “Devices” in *Sun Cluster Overview for Solaris OS* and “Global Devices” in *Sun Cluster Concepts Guide for Solaris OS*.

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

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

Device Groups

For information about the purpose and function of device groups, see “Devices” in *Sun Cluster Overview for Solaris OS* and “Device Groups” in *Sun Cluster Concepts Guide for Solaris OS*.

Add this planning information to the “[Device Group Configurations Worksheet](#)” on page 279.

Consider the following points when you plan device groups.

- **Failover** – You can configure multihost disks and properly configured volume-manager devices as failover devices. Proper configuration of a volume-manager device includes multihost disks and correct setup of the volume manager itself. This configuration ensures that multiple nodes can host the exported device. You cannot configure tape drives, CD-ROMs or DVD-ROMs, or single-ported devices as failover devices.
- **Mirroring** – You must mirror the disks to protect the data from disk failure. See “[Mirroring Guidelines](#)” on page 42 for additional guidelines. See “[Configuring Solaris Volume Manager Software](#)” on page 127 or “[Installing and Configuring VxVM Software](#)” on page 159 and your volume-manager documentation for instructions about mirroring.
- **Storage-based replication** - Disks in a device group must be either all replicated or none replicated. A device group cannot use a mix of replicated and nonreplicated disks.

Cluster File Systems

For information about the purpose and function of cluster file systems, see “Cluster File Systems” in *Sun Cluster Overview for Solaris OS* and “Cluster File Systems” in *Sun Cluster Concepts Guide for Solaris OS*.

Note – You can alternatively configure highly available local file systems. This can provide better performance to support a data service with high I/O, or to permit use of certain file-system features that are not supported in a cluster file system.. For more information, see “Enabling Highly Available Local File Systems” in *Sun Cluster Data Services Planning and Administration Guide for Solaris OS*.

Consider the following points when you plan cluster file systems.

- **Quotas** - Quotas are not supported on cluster file systems. However, quotas are supported on highly available local file systems.
- **Non-global zones** - If a cluster file system is to be accessed from a non-global zone, it must first be mounted in the global zone. The cluster file system is then mounted in the non-global zone by using a loopback mount. Therefore, the loopback file system (LOFS) must be enabled in a cluster that contains non-global zones.
- **Loopback file system (LOFS)** - During cluster creation with the Solaris 9 version of Sun Cluster software, LOFS is disabled by default. During cluster creation with the Solaris 10 version of Sun Cluster software, LOFS is enabled by default.

You must manually disable LOFS on each cluster node if the cluster meets both of the following conditions:

- Sun Cluster HA for NFS is configured on a highly available local file system.
- The automountd daemon is running.

If the cluster meets both of these conditions, you must disable LOFS to avoid switchover problems or other failures. If the cluster meets only one of these conditions, you can safely enable LOFS.

If you require both LOFS and the `automountd` daemon to be enabled, exclude from the automounter map all files that are part of the highly available local file system that is exported by Sun Cluster HA for NFS.

- **Process accounting log files** - Do not locate process accounting log files on a cluster file system or on a highly available local file system. A switchover would be blocked by writes to the log file, which would cause the node to hang. Use only a local file system to contain process accounting log files.
- **Communication endpoints** - The cluster file system does not support any of the file-system features of Solaris software by which one would put a communication endpoint in the file-system namespace.
 - Although you can create a UNIX domain socket whose name is a path name into the cluster file system, the socket would not survive a node failover.
 - Any FIFOs or named pipes that you create on a cluster file system would not be globally accessible.

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

- **Device special files** - Neither block special files nor character special files are supported in a cluster file system. To specify a path name to a device node in a cluster file system, create a symbolic link to the device name in the `/dev` directory. Do not use the `mknod` command for this purpose.
- **atime** - Cluster file systems do not maintain `atime`.
- **ctime** - When a file on a cluster file system is accessed, the update of the file's `ctime` might be delayed.
- **Installing applications** - If you want the binaries of a highly available application to reside on a cluster file system, wait to install the application until after the cluster file system is configured. Also, if the application is installed by using the Sun Java System `installer` program and the application depends on any shared components, install those shared components on all nodes in the cluster that are not installed with the application.

Choosing Mount Options for Cluster File Systems

This section describes requirements and restrictions for the following types of cluster file systems:

- “UFS Cluster File Systems” on page 34
- “QFS Shared File Systems” on page 35
- “VxFS Cluster File Systems” on page 35

Note – You can alternatively configure these and other types of file systems as highly available local file systems. For more information, see “Enabling Highly Available Local File Systems” in *Sun Cluster Data Services Planning and Administration Guide for Solaris OS*.

Follow these guidelines to determine what mount options to use when you create your cluster file systems.

UFS Cluster File Systems

Mount Option	Usage	Description
<code>global</code>	Required	This option makes the file system globally visible to all nodes in the cluster.
<code>logging</code>	Required	This option enables logging.
<code>forcedirectio</code>	Conditional	This option is required only for cluster file systems that will host Oracle RAC RDBMS data files, log files, and control files. Note – Oracle RAC is supported for use only in SPARC based clusters.
<code>onerror=panic</code>	Required	You do not have to explicitly specify the <code>onerror=panic</code> mount option in the <code>/etc/vfstab</code> file. This mount option is already the default value if no other <code>onerror</code> mount option is specified. Note – Only the <code>onerror=panic</code> mount option is supported by Sun Cluster software. Do not use the <code>onerror=umount</code> or <code>onerror=lock</code> mount options. These mount options are not supported on cluster file systems for the following reasons: <ul style="list-style-type: none"> ■ Use of the <code>onerror=umount</code> or <code>onerror=lock</code> mount option might cause the cluster file system to lock or become inaccessible. This condition might occur if the cluster file system experiences file corruption. ■ The <code>onerror=umount</code> or <code>onerror=lock</code> mount option might cause the cluster file system to become unmountable. This condition might thereby cause applications that use the cluster file system to hang or prevent the applications from being killed. <p>A node might require rebooting to recover from these states.</p>
<code>syncdir</code>	Optional	If you specify <code>syncdir</code> , you are guaranteed POSIX-compliant file system behavior for the <code>write()</code> system call. If a <code>write()</code> succeeds, then this mount option ensures that sufficient space is on the disk. If you do not specify <code>syncdir</code> , the same behavior occurs that is seen with UFS file systems. When you do not specify <code>syncdir</code> , performance of writes that allocate disk blocks, such as when appending data to a file, can significantly improve. However, in some cases, without <code>syncdir</code> you would not discover an out-of-space condition (<code>ENOSPC</code>) until you close a file. You see <code>ENOSPC</code> on close only during a very short time after a failover. With <code>syncdir</code> , as with POSIX behavior, the out-of-space condition would be discovered before the close.

See the `mount_ufs(1M)` man page for more information about UFS mount options.

QFS Shared File Systems

Mount Parameter	Usage	Description
shared	Required	This option specifies that this is a shared file system, therefore globally visible to all nodes in the cluster.



Caution – Ensure that settings in the `/etc/vfstab` file do not conflict with settings in the `/etc/opt/SUNWsamfs/samfs.cmd` file. Settings in the `/etc/vfstab` file override settings in the `/etc/opt/SUNWsamfs/samfs.cmd` file.

Certain data services such as Sun Cluster Support for Oracle RAC have additional requirements and guidelines for QFS mount parameters. See your data service manual for any additional requirements.

Note – Logging is not enabled by an `/etc/vfstab` mount parameter, nor does Sun Cluster software require logging for QFS shared file systems.

See the `mount_samfs(1M)` man page for more information about QFS mount parameters.

VxFS Cluster File Systems

Mount Option	Usage	Description
global	Required	This option makes the file system globally visible to all nodes in the cluster.
log	Required	This option enables logging.

See the VxFS `mount_vxfs` man page and “Overview of the Administering Cluster File Systems” in *Sun Cluster System Administration Guide for Solaris OS* for more information about VxFS mount options.

Mount Information for Cluster File Systems

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

- **Mount-point location** – Create mount points for cluster file systems in the `/global` directory, unless you are prohibited by other software products. By using the `/global` directory, you can more easily distinguish cluster file systems, which are globally available, from local file systems.

- **SPARC: VxFS mount requirement** – If you use VERITAS File System (VxFS), globally mount and unmount a VxFS file system from the primary node. The primary node is the node that masters the disk on which the VxFS file system resides. This method ensures that the mount or unmount operation succeeds. A VxFS file-system mount or unmount operation that is performed from a secondary node might fail.
- The following VxFS features are not supported in a Sun Cluster 3.2 cluster file system. They are, however, supported in a local file system.
 - Quick I/O
 - Snapshots
 - Storage checkpoints
 - VxFS-specific mount options:
 - `convosync` (Convert `O_SYNC`)
 - `mincache`
 - `qlog, delaylog, tmplog`
 - VERITAS cluster file system (requires VxVM cluster feature & VERITAS Cluster Server)

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

All other VxFS features and options that are supported in a cluster file system are supported by Sun Cluster 3.2 software. See VxFS documentation for details about VxFS options that are supported in a cluster configuration.

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

Planning Volume Management

Add this planning information to the “[Device Group Configurations Worksheet](#)” on page 279 and the “[Volume-Manager Configurations Worksheet](#)” on page 281. For Solaris Volume Manager, also add this planning information to the “[Volumes Worksheet \(Solaris Volume Manager\)](#)” on page 283.

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

- “[Guidelines for Volume-Manager Software](#)” on page 37
- “[Guidelines for Solaris Volume Manager Software](#)” on page 38
- “[Guidelines for VERITAS Volume Manager Software](#)” on page 39

- “File-System Logging” on page 41
- “Mirroring Guidelines” on page 42

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

TABLE 1–4 Supported Use of Volume Managers With Sun Cluster Software

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

See your volume-manager documentation and [“Configuring Solaris Volume Manager Software” on page 127](#) or [“Installing and Configuring VxVM Software” on page 159](#) for instructions about how to install and configure the volume-manager software. For more information about the use of volume management in a cluster configuration, see “Multihost Devices” in *Sun Cluster Concepts Guide for Solaris OS* and “Device Groups” in *Sun Cluster Concepts Guide for Solaris OS*.

Guidelines for Volume-Manager Software

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

- **Software RAID** – Sun Cluster software does not support software RAID 5.
- **Mirrored multihost disks** – You must mirror all multihost disks across disk expansion units. See [“Guidelines for Mirroring Multihost Disks” on page 42](#) for guidelines on mirroring multihost disks. You do not need to use software mirroring if the storage device provides hardware RAID as well as redundant paths to devices.
- **Mirrored root** – Mirroring the root disk ensures high availability, but such mirroring is not required. See [“Mirroring Guidelines” on page 42](#) for guidelines about deciding whether to mirror the root disk.

- **Unique naming** – You might have local Solaris Volume Manager or VxVM volumes that are used as devices on which the `/global/.devices/node@nodeid` file systems are mounted. If so, the name of each local volume on which a `/global/.devices/node@nodeid` file system is to be mounted must be unique throughout the cluster.
- **Node lists** – To ensure high availability of a 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 or zones than its associated device group, make the scalable resource group's node list a superset of the device group's node list. See the resource group planning information in the *Sun Cluster Data Services Planning and Administration Guide for Solaris OS* for information about node lists.
- **Multihost disks** – You must connect, or port, all devices that are used to construct a device group to all of the nodes that are configured in the node list for that device group. Solaris Volume Manager software can automatically check for this connection at the time that devices are added to a disk set. However, configured VxVM disk groups do not have an association to any particular set of nodes.
- **Hot spare disks** – You can use hot spare disks to increase availability, but hot spare disks are not required.

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

Guidelines for Solaris Volume Manager Software

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

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

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

- `/kernel/drv/md.conf` **settings** – SPARC: On the Solaris 9 OS, Solaris Volume Manager volumes that are used by each disk set are created in advance, at reconfiguration boot time. This reconfiguration is based on the configuration parameters that exist in the `/kernel/drv/md.conf` file.

Note – With the Solaris 10 release, Solaris Volume Manager has been enhanced to configure volumes dynamically. You no longer need to edit the `nmd` and the `md_nsets` parameters in the `/kernel/drv/md.conf` file. New volumes are dynamically created, as needed.

You must modify the `nmd` and `md_nsets` fields as follows to support a Sun Cluster configuration on the Solaris 9 OS:



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

- `md_nsets` – The `md_nsets` field defines the total number of disk sets that can be created for a system to meet the needs of the entire cluster. Set the value of `md_nsets` to the expected number of disk sets in the cluster plus one additional disk set. Solaris Volume Manager software uses the additional disk set to manage the private disks on the local host.

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

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

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

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

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

See “System Files and Startup Files” in *Solaris Volume Manager Administration Guide* (Solaris 9 or Solaris 10) 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.

- **Accessibility to nodes** - You must configure all volume-manager disk groups as either Sun Cluster device groups or as local-only disk groups. If you do not configure the disk group in one of these ways, the devices in the disk group will not be accessible to any node in the cluster.
 - A device group enables a secondary node to host multihost disks if the primary node fails.
 - A local-only disk group functions outside the control of Sun Cluster software and can be accessed from only one node at a time.
- **Enclosure-Based Naming** – If you use Enclosure-Based Naming of devices, ensure that you use consistent device names on all cluster nodes that share the same storage. VxVM does not coordinate these names, so the administrator must ensure that VxVM assigns the same names to the same devices from different nodes. Failure to assign consistent names does not interfere with correct cluster behavior. However, inconsistent names greatly complicate cluster administration and greatly increase the possibility of configuration errors, potentially leading to loss of data.
- **Root disk group** – The creation of a root disk group is optional.

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

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

The root disk group must be local to the node.

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

See your VxVM installation documentation for additional information.

File-System Logging

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

- Solaris UFS logging – See the `mount_ufs(1M)` man page for more information.
- (Solaris 9 only) SPARC: Solaris Volume Manager transactional-volume logging – See “Transactional Volumes (Overview)” in *Solaris Volume Manager Administration Guide* for more information.

Note – Solaris Volume Manager transactional-volume logging is removed from the Solaris 10 OS. Solaris UFS logging provides the same capabilities but superior performance, as well as lower system administration requirements and overhead.

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

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

TABLE 1-5 Supported File System Logging Matrix

Volume Manager	Supported File System Logging
Solaris Volume Manager	<ul style="list-style-type: none"> ▪ Solaris UFS logging ▪ SPARC: Solaris Volume Manager transactional-volume logging (Solaris 9 only) ▪ VxFS logging
VERITAS Volume Manager	<ul style="list-style-type: none"> ▪ Solaris UFS logging ▪ VxFS logging

SPARC: On the Solaris 9 OS, consider the following points when you choose between Solaris UFS logging and Solaris Volume Manager transactional-volume logging for UFS cluster file systems:

- Solaris UFS logging always allocates the log by using free space on the UFS file system, and depending on the size of the file system.
 - On file systems less than 1 Gbyte, the log occupies 1 Mbyte.
 - On file systems 1 Gbyte or greater, the log occupies 1 Mbyte per Gbyte on the file system, to a maximum of 64 Mbytes.
- A transactional volume manages UFS logging. The logging device component of a transactional volume is a volume that you can mirror and stripe. You can create a maximum 1-Gbyte log size, although 64 Mbytes is sufficient for most file systems. The minimum log size is 1 Mbyte.

Mirroring Guidelines

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

- [“Guidelines for Mirroring Multihost Disks”](#) on page 42
- [“Guidelines for Mirroring the Root Disk”](#) on page 42

Guidelines for Mirroring Multihost Disks

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

Consider the following points when you mirror multihost disks:

- **Separate disk expansion units** – Each submirror of a given mirror or plex should reside in a different multihost expansion unit.
- **Disk space** – Mirroring doubles the amount of necessary disk space.
- **Three-way mirroring** – Solaris Volume Manager software and VERITAS Volume Manager (VxVM) software support three-way mirroring. However, Sun Cluster software requires only two-way mirroring.
- **Differing device sizes** – If you mirror to a device of a different size, your mirror capacity is limited to the size of the smallest submirror or plex.

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

Guidelines for Mirroring the Root Disk

Add this planning information to the [“Local File System Layout Worksheet”](#) on page 275.

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

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

See your volume-manager documentation and [“Configuring Solaris Volume Manager Software”](#) on page 127 or [“Installing and Configuring VxVM Software”](#) on page 159 for instructions about how to mirror the root disk.

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

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

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

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

Installing Software on the Cluster

This chapter provides procedures for how to install software on cluster nodes and the administration console.

Installing the Software

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

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

TABLE 2-1 Task Map: Installing the Software

Task	Instructions
1. Plan the layout of your cluster configuration and prepare to install software.	“How to Prepare for Cluster Software Installation” on page 46
2. <i>(Optional)</i> Install Cluster Control Panel (CCP) software on the administrative console.	“How to Install Cluster Control Panel Software on an Administrative Console” on page 47
3. Install the Solaris OS on all nodes.	“How to Install Solaris Software” on page 50
4. <i>(Optional)</i> Configure internal disk mirroring.	“How to Configure Internal Disk Mirroring” on page 54
5. <i>(Optional)</i> SPARC: Install and configure Sun multipathing software.	“How to Install Sun Multipathing Software” on page 55
6. <i>(Optional)</i> SPARC: Install VERITAS File System software.	“SPARC: How to Install VERITAS File System Software” on page 57
7. Install Sun Cluster software and any data services that you will use.	“How to Install Sun Cluster Framework and Data-Service Software Packages” on page 58
8. Set up directory paths.	“How to Set Up the Root Environment” on page 61

▼ How to Prepare for Cluster Software Installation

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

1 Ensure that the combination of hardware and software that you choose for your cluster is currently a supported Sun Cluster configuration.

Contact your Sun sales representative for the most current information about supported cluster configurations.

2 Read the following manuals for information that can help you plan your cluster configuration and prepare your installation strategy.

- *Sun Cluster 3.2 Release Notes for Solaris OS* - Restrictions, bug workarounds, and other late-breaking information.
- *Sun Cluster Overview for Solaris OS* and *Sun Cluster Concepts Guide for Solaris OS* - Overviews of the Sun Cluster product.
- *Sun Cluster Software Installation Guide for Solaris OS* (this manual) - Planning guidelines and procedures for installing and configuring Solaris, Sun Cluster, and volume-manager software.
- *Sun Cluster Data Services Planning and Administration Guide for Solaris OS* - Planning guidelines and procedures to install and configure data services.

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

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

- Solaris OS
- Solaris Volume Manager software
- Sun StorEdge QFS software
- VERITAS Volume Manager
- Third-party applications

4 Plan your cluster configuration.



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

For example, the Oracle Real Application Clusters Guard option of Oracle RAC 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 Data Services Planning and Administration Guide for Solaris OS* to determine how to install and configure your cluster.

- Fill out the cluster framework and data-services configuration worksheets that are referenced in the planning guidelines. Use your completed worksheets for reference during the installation and configuration tasks.

5 Obtain all necessary patches for your cluster configuration.

See “Patches and Required Firmware Levels” in *Sun Cluster 3.2 Release Notes for Solaris OS* for the location of patches and installation instructions.

Next Steps If you want to use Cluster Control Panel software to connect from an administrative console to your cluster nodes, go to “[How to Install Cluster Control Panel Software on an Administrative Console](#)” on page 47.

Otherwise, choose the Solaris installation procedure to use.

- To install Sun Cluster software by using the `scinstall(1M)` utility, go to “[How to Install Solaris Software](#)” on page 50 to first install Solaris software.
- To install Solaris and Sun Cluster software in the same operation (JumpStart method), go to “[How to Install Solaris and Sun Cluster Software \(JumpStart\)](#)” on page 79.

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

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

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

You can use any desktop machine that runs a version of the Solaris OS that is supported by Sun Cluster 3.2 software as an administrative console. If you are using Sun Cluster software on a SPARC based system, you can also use the administrative console as a Sun Management Center console or server as well. See Sun Management Center documentation for information about how to install Sun Management Center software.

Before You Begin Ensure that a supported version of the Solaris OS and any Solaris patches are installed on the administrative console. All platforms require at least the End User Solaris Software Group.

1 Become superuser on the administrative console.

2 Load the Sun Java Availability Suite DVD-ROM into the DVD-ROM drive.

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

3 Change to the `Solaris_arch/Product/sun_cluster/Solaris_ver/Packages/` directory, where `arch` is `sparc` or `x86` (Solaris 10 only), and where `ver` is 9 for Solaris 9 or 10 for Solaris 10 .

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

4 Install the `SUNWcccon` package.

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

5 (Optional) Install the `SUNWscman` package.

```
adminconsole# 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 Unload the Sun Java Availability Suite DVD-ROM from the DVD-ROM drive.

a. To ensure that the DVD-ROM is not being used, change to a directory that does *not* reside on the DVD-ROM.

b. Eject the DVD-ROM.

```
adminconsole# eject cdrom
```

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

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

```
adminconsole# 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, and the port number. Examples of a console-access device are a terminal concentrator (TC), a System Service Processor (SSP), and a Sun Fire system controller.

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

`node1, node2` Physical names of the cluster nodes.

`ca-dev-hostname` Hostname of the console-access device.

`port` Serial port number.

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

- For a Sun Fire 15000 system controller, use `telnet(1)` port number 23 for the serial port number of each entry.
- For all other console-access devices, use the `telnet` serial port number, not the physical port number. To determine the telnet serial port number, add 5000 to the physical port number. For example, if a physical port number is 6, the telnet serial port number is 5006.
- For Sun Enterprise 10000 servers, also see the `/opt/SUNWcluster/bin/serialports(4)` man page for details and special considerations.
- To connect the administrative console directly to the cluster nodes or through a management network, specify for each node its hostname and the port number that the node uses to connect to the administrative console or the management network.

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

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

10 Start the CCP utility.

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

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

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

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

Next Steps Determine whether the Solaris OS is already installed to meet Sun Cluster software requirements. See “[Planning the Solaris OS](#)” on page 14 for information about Sun Cluster installation requirements for the Solaris OS.

- If the Solaris OS meets Sun Cluster requirements, go to “[How to Install Sun Cluster Framework and Data-Service Software Packages](#)” on page 58.
- If the Solaris OS does not meet Sun Cluster requirements, install, reconfigure, or reinstall the Solaris OS as needed.
 - To install the Solaris OS alone, go to “[How to Install Solaris Software](#)” on page 50.

- To use the `scinstall` custom JumpStart method to install both the Solaris OS and Sun Cluster software, go to [“How to Install Solaris and Sun Cluster Software \(JumpStart\)” on page 79](#)

▼ How to Install Solaris Software

If you do not use the `scinstall` custom JumpStart installation method to install software, perform this procedure to install the Solaris OS on each node in the cluster. See [“How to Install Solaris and Sun Cluster Software \(JumpStart\)” on page 79](#) for more information about JumpStart installation of a cluster.

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

If your nodes are already installed with the Solaris OS but do not meet Sun Cluster installation requirements, you might need to reinstall the Solaris software. Follow the steps in this procedure to ensure subsequent successful installation of Sun Cluster software. See [“Planning the Solaris OS” on page 14](#) for information about required root-disk partitioning and other Sun Cluster installation requirements.

Before You Begin Perform the following tasks:

- Ensure that the hardware setup is complete and that connections are verified before you install Solaris software. See the *Sun Cluster Hardware Administration Collection* and your server and storage device documentation for details.
- Ensure that your cluster configuration planning is complete. See [“How to Prepare for Cluster Software Installation” on page 46](#) for requirements and guidelines.
- Complete the [“Local File System Layout Worksheet” on page 275](#).
- If you use a naming service, add address-to-name mappings for all public hostnames and logical addresses to any naming services that clients use for access to cluster services. See [“Public Network IP Addresses” on page 21](#) for planning guidelines. See your Solaris system-administrator documentation for information about using Solaris naming services.

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

- **If Cluster Control Panel (CCP) software is installed and configured on your administrative console, use the `cconsole(1M)` utility to display the individual console screens.**

As superuser, use the following command to start the `cconsole` utility:

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

The `cconsole` utility 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 the `cconsole` utility, connect to the consoles of each node individually.

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

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

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

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

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

See “Solaris Software Group Considerations” on page 16 for information about additional Solaris software requirements.

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

- Create a file system of at least 512 Mbytes for use by the global-device subsystem.

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

- Specify that slice 7 is at least 20 Mbytes in size.
- Create any other file-system partitions that you need, as described in “System Disk Partitions” on page 16.

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

3 If you will use role-based access control (RBAC) instead of superuser to access the cluster nodes, set up an RBAC role that provides authorization for all Sun Cluster commands.

This series of installation procedures requires the following Sun Cluster RBAC authorizations if the user is not superuser:

- `solaris.cluster.modify`
- `solaris.cluster.admin`
- `solaris.cluster.read`

See “Role-Based Access Control (Overview)” in *System Administration Guide: Security Services* for more information about using RBAC roles. See the Sun Cluster man pages for the RBAC authorization that each Sun Cluster subcommand requires.

4 If you are adding a node to an existing cluster, add mount points for cluster file systems to the new node.

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

```
phys-schost-1# 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.

```
phys-schost-new# 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.

5 If you are adding a node and VxVM is installed on any node in the cluster, perform the following tasks.

- a. Ensure that the same `vxio` number is used on the VxVM-installed nodes.

```
phys-schost# grep vxio /etc/name_to_major
vxio NNN
```

- b. Ensure that the `vxio` number is available for use on each of the nodes that do not have VxVM installed.

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

6 If you installed the End User Solaris Software Group and you want to use any of the following Sun Cluster features, install additional Solaris software packages to support these features.

- Remote Shared Memory Application Programming Interface (RSMAPI)
- RSMRDT drivers
- SPARC: SCI-PCI adapters

- **SPARC: For the Solaris 9 OS, use the following command:**

```
phys-schost# pkgadd -d . SUNWrsm SUNWrsmc SUNWrsmo SUNWrsmox
```

- **For the Solaris 10 OS, use the following command:**

```
phys-schost# pkgadd -G -d . SUNWrsm SUNWrsmo
```

You must add these packages only to the global zone. The `-G` option adds packages to the current zone only. This option also specifies that the packages are *not* propagated to any existing non-global zone or to any non-global zone that is created later.

- 7 Install any required Solaris OS patches and hardware-related firmware and patches, including those for storage-array support. Also download any needed firmware that is contained in the hardware patches.**

See “Patches and Required Firmware Levels” in *Sun Cluster 3.2 Release Notes for Solaris OS* for the location of patches and installation instructions.

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

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

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

- 9 Update the `/etc/inet/ipnodes` file on each node with all public IP addresses that are used in the cluster.**

Perform this step regardless of whether you are using a naming service. The `ipnodes` file can contain both IPv4 and IPv6 addresses. See “Public Network IP Addresses” on page 21 for a listing of Sun Cluster components whose IP addresses you must add.

Note – During establishment of a new cluster or new cluster node, the `scinstall` utility automatically adds the public IP address of each node that is being configured to the `/etc/inet/hosts` file. Adding these IP addresses to the `/etc/inet/ipnodes` file is optional.

- 10 If you will use `ce` adapters for the cluster interconnect, add the following entry to the `/etc/system` file.**

```
set ce:ce_taskq_disable=1
```

This entry becomes effective after the next system reboot.

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

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

```
set kernel_cage_enable=1
```

This entry becomes effective after the next system reboot. See your server documentation for more information about dynamic reconfiguration.

- 12 (Optional) Configure public-network adapters in IPMP groups.**

If you do not want to use the multiple-adapter IPMP groups that the `scinstall` utility configures during cluster creation, configure custom IPMP groups as you would in a standalone system. See Part VI, “IPMP,” in *System Administration Guide: IP Services* for details.

During cluster creation, the `scinstall` utility configures each set of public-network adapters that use the same subnet and are not already configured in an IPMP group into a single multiple-adapter IPMP group. The `scinstall` utility ignores any existing IPMP groups.

Next Steps If your server supports the mirroring of internal hard drives and you want to configure internal disk mirroring, go to [“How to Configure Internal Disk Mirroring”](#) on page 54.

Otherwise, to use Sun multipathing software, go to [“How to Install Sun Multipathing Software”](#) on page 55.

Otherwise, to install VxFS, go to [“SPARC: How to Install VERITAS File System Software”](#) on page 57.

Otherwise, install the Sun Cluster software packages. Go to [“How to Install Sun Cluster Framework and Data-Service Software Packages”](#) on page 58.

See Also See the *Sun Cluster System Administration Guide for Solaris OS* for procedures to perform dynamic reconfiguration tasks in a Sun Cluster configuration.

▼ How to Configure Internal Disk Mirroring

Perform this procedure on each node of the cluster to configure internal hardware RAID disk mirroring to mirror the system disk. This procedure is optional.

Note – Do not perform this procedure under either of the following circumstances:

- Your servers do not support the mirroring of internal hard drives.
 - You have already established the cluster. Instead, perform “Mirroring Internal Disks on Servers that Use Internal Hardware Disk Mirroring or Integrated Mirroring” in *Sun Cluster 3.1 - 3.2 Hardware Administration Manual for Solaris OS*.
-

Before You Begin Ensure that the Solaris operating system and any necessary patches are installed.

1 Become superuser.

2 Configure an internal mirror.

```
phys-schost# raidctl -c clt0d0 clt1d0
```

```
-c clt0d0 clt1d0
```

Creates the mirror of primary disk to the mirror disk. Enter the name of your primary disk as the first argument. Enter the name of the mirror disk as the second argument.

For specifics about how to configure your server’s internal disk mirroring, refer to the documents that shipped with your server and the `raidctl(1M)` man page.

Next Steps To use Sun multipathing software, go to [“How to Install Sun Multipathing Software”](#) on page 55.

Otherwise, to install VxFS, go to “SPARC: How to Install VERITAS File System Software” on page 57.

Otherwise, install the Sun Cluster software packages. Go to “How to Install Sun Cluster Framework and Data-Service Software Packages” on page 58.

▼ How to Install Sun Multipathing Software

Perform this procedure on each node of the cluster to install and configure Sun multipathing software for fiber channel (FC) storage. Multipathing software manages multiple I/O paths to the shared cluster storage. This procedure is optional.

- SPARC: For the Solaris 9 OS, you install and configure Sun StorEdge Traffic Manager software.
- For the Solaris 10 OS, you enable the Solaris multipathing feature, which is installed by default as part of the Solaris 10 software.

Before You Begin Perform the following tasks:

- Ensure that the Solaris OS is installed to support Sun Cluster software.
If Solaris software is already installed on the node, you must ensure that the Solaris installation meets the requirements for Sun Cluster software and any other software that you intend to install on the cluster. See “How to Install Solaris Software” on page 50 for more information about installing Solaris software to meet Sun Cluster software requirements.
- SPARC: For the Solaris 9 OS, have available your software packages, patches, and documentation for Sun StorEdge Traffic Manager software and Sun StorEdge SAN Foundation software. See <http://www.sun.com/products-n-solutions/hardware/docs/> for links to documentation.
- For the Solaris 10 OS, have available the at *Solaris Fibre Channel Storage Configuration and Multipathing Administration Guide*.

1 Become superuser.

2 SPARC: For the Solaris 9 OS, install on each node Sun StorEdge Traffic Manager software and any necessary patches.

- For the procedure about how to install Sun StorEdge Traffic Manager software, see the *Sun StorEdge Traffic Manager Installation and Configuration Guide* at <http://www.sun.com/products-n-solutions/hardware/docs/>.
- For a list of required patches for Sun StorEdge Traffic Manager software, see the *Sun StorEdge Traffic Manager Software Release Notes* at <http://www.sun.com/storage/san/>.

3 Enable multipathing functionality.

- **For the Solaris 9 OS, change the value of the `mpxio-disable` parameter to `no`.**

Modify this entry in the `/kernel/drv/scsi_vhci.conf` file on each node.

```
set mpxio-disable=no
```

- **For the Solaris 10 OS, issue the following command on each node:**



Caution – If Sun Cluster software is already installed, do not issue this command. Running the `stmsboot` command on an active cluster node might cause Solaris services to go into the maintenance state. Instead, follow instructions in the `stmsboot(1M)` man page for using the `stmsboot` command in a Sun Cluster environment.

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

`-e` Enables Solaris I/O multipathing.

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

4 SPARC: For the Solaris 9 OS, determine whether your version of Sun StorEdge SAN Foundation software includes built-in support for your storage array.

If the software does not include built-in support for your storage array, edit the `/kernel/drv/scsi_vhci.conf` file on each node to include the necessary entries. For more information, see the release notes for your storage device.

5 SPARC: For the Solaris 9 OS, shut down each node and perform a reconfiguration boot.

The reconfiguration boot creates the new Solaris device files and links.

```
phys-schost# shutdown -y -g0 -i0
```

```
ok boot -r
```

6 After the reconfiguration reboot is finished on all nodes, perform any additional tasks that are necessary to complete the configuration of your storage array.

See installation instructions for your storage array in the *Sun Cluster Hardware Administration Collection* for details.

Troubleshooting

If you installed Sun multipathing software after Sun Cluster software was installed on the cluster, DID mappings might require updating. Issue the following commands on each node of the cluster to regenerate the DID namespace.

```
phys-schost# cldevice clear
```



```
phys-schost# cldevice refresh
(Solaris 9 only) phys-schost# cfgadm -c configure
phys-schost# cldevice populate
```

See the `cfgadm(1M)` and `cldevice(1CL)` man pages for more information.

Next Steps To install VxFS, go to [“SPARC: How to Install VERITAS File System Software”](#) on page 57.

Otherwise, install the Sun Cluster software packages. Go to [“How to Install Sun Cluster Framework and Data-Service Software Packages”](#) on page 58.

▼ SPARC: How to Install VERITAS File System Software

To use VERITAS File System (VxFS) software in the cluster, perform this procedure on each node of the cluster.

1 Follow the procedures in your VxFS installation documentation to install VxFS software on each node of the cluster.

2 Install any Sun Cluster patches that are required to support VxFS.

See “Patches and Required Firmware Levels” in *Sun Cluster 3.2 Release Notes for Solaris OS* for the location of patches and installation instructions.

3 In the `/etc/system` file on each node, set the following values.

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

These changes become effective at the next system reboot.

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

Next Steps Install the Sun Cluster software packages. Go to [“How to Install Sun Cluster Framework and Data-Service Software Packages”](#) on page 58.

▼ How to Install Sun Cluster Framework and Data-Service Software Packages

Follow this procedure to use the Sun Java™ Enterprise System (Java ES) installer program to perform one or more of the following installation tasks:

- To install the Sun Cluster framework software packages on each node in the cluster.
- To install Sun Cluster framework software on the master node that you will flash archive for a JumpStart installation. See “[How to Install Solaris and Sun Cluster Software \(JumpStart\)](#)” on page 79 for more information about JumpStart installation of a cluster.
- To install data services.

Note – For the Solaris 10 OS, this procedure installs data services only to the global zone. To install data services to be visible only from within a certain non-global zone, see “[How to Create a Non-Global Zone on a Cluster Node](#)” on page 179.

Note – This procedure uses the interactive form of the installer program. To use the noninteractive form of the installer program, such as when developing installation scripts, see Chapter 5, “Installing in Silent Mode,” in *Sun Java Enterprise System 5 Installation Guide for UNIX*.

Before You Begin Perform the following tasks:

- Ensure that the Solaris OS is installed to support Sun Cluster software.
If Solaris software is already installed on the node, you must ensure that the Solaris installation meets the requirements for Sun Cluster software and any other software that you intend to install on the cluster. See “[How to Install Solaris Software](#)” on page 50 for more information about installing Solaris software to meet Sun Cluster software requirements.
- Have available the Sun Java Availability Suite DVD-ROM.

1 (Optional) To use the installer program with a GUI, ensure that the display environment of the cluster node to install is set to display the GUI.

```
% xhost +
% setenv DISPLAY nodename:0.0
```

If you do not make these settings, the installer program runs in text-based mode.

2 Become superuser on the cluster node to install.

3 Load the Sun Java Availability Suite DVD-ROM into the DVD-ROM drive.

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

4 Change to the installation wizard directory of the DVD-ROM.

- If you are installing the software packages on the SPARC platform, type the following command:

```
phys-schost# cd /cdrom/cdrom0/Solaris_sparc
```

- If you are installing the software packages on the x86 platform, type the following command:

```
phys-schost# cd /cdrom/cdrom0/Solaris_x86
```

5 Start the installation wizard program.

```
phys-schost# ./installer
```

See the *Sun Java Enterprise System 5 Installation Guide for UNIX* for additional information about using the different forms and features of the Java ES installer program.

6 Follow instructions on the screen to install Sun Cluster framework software and data services on the node.

- If you do not want to install Sun Cluster Manager, formerly SunPlex Manager, deselect it.

Note – You must install Sun Cluster Manager either on all nodes of the cluster or on none.

- If you want to install Sun Cluster Geographic Edition software, select it.

After the cluster is established, see *Sun Cluster Geographic Edition Installation Guide* for further installation procedures.

- Choose Configure Later when prompted whether to configure Sun Cluster framework software.

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

7 Install additional packages to use any of the following features.

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

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

a. Determine which packages you must install.

The following table lists the Sun Cluster 3.2 packages that each feature requires, in the order in which you must install each group of packages. The Java ES installer program does not automatically install these packages.

Note – Install packages in the order in which they are listed in the following table.

Feature	Additional Sun Cluster 3.2 Packages to Install
RSMAPI	SUNWscrif
SCI-PCI adapters	<ul style="list-style-type: none"> ■ <i>Solaris 9</i>: SUNWsci SUNWscid SUNWscidx ■ <i>Solaris 10</i>: SUNWscir SUNWsci SUNWscidr SUNWscid
RSMRDT drivers	SUNWscrdt

- b. Change to the `Solaris_arch/Product/sun_cluster/Solaris_ver/Packages/` directory, where `arch` is `sparc` or `x86` (Solaris 10 only), and where `ver` is 9 for Solaris 9 or 10 for Solaris 10.**

```
phys-schost# cd /cdrom/cdrom0/Solaris_arch/Product/sun_cluster/Solaris_ver/Packages/
```

- c. Install the additional packages.**

- **SPARC: For the Solaris 9 OS, use the following command:**

```
phys-schost# pkgadd -d . packages
```

- **For the Solaris 10 OS, use the following command:**

```
phys-schost# pkgadd -G -d . packages
```

- 8 Unload the Sun Java Availability Suite DVD-ROM from the DVD-ROM drive.**

- a. To ensure that the DVD-ROM is not being used, change to a directory that does *not* reside on the DVD-ROM.**

- b. Eject the DVD-ROM.**

```
phys-schost# eject cdrom
```

- 9 Apply any necessary patches to support Sun Cluster software.**

See “Patches and Required Firmware Levels” in *Sun Cluster 3.2 Release Notes for Solaris OS* for the location of patches and installation instructions.

Next Steps If you want to install Sun StorEdge QFS file system software, follow the procedures for initial installation in the *Sun StorEdge QFS Installation and Upgrade Guide*.

Otherwise, to set up the root user environment, go to “How to Set Up the Root Environment” on page 61.

▼ How to Set Up the Root Environment

Note – In a Sun Cluster configuration, user initialization files for the various shells must verify that they are run from an interactive shell. The files must verify this before they attempt to output to the terminal. Otherwise, unexpected behavior or interference with data services might occur. See “Customizing a User’s Work Environment” in *System Administration Guide: Basic Administration* (Solaris 9 or Solaris 10) for more information.

Perform this procedure on each node in the cluster.

- 1 Become superuser on a cluster node.**
- 2 Modify PATH and MANPATH entries in the `.cshrc` or `.profile` file.**
 - a. Add `/usr/sbin/` and `/usr/cluster/bin/` to the PATH.**
 - b. Add `/usr/cluster/man/` to the MANPATH.**

See your Solaris OS documentation, volume manager documentation, and other application documentation for additional file paths to set.
- 3 (Optional) For ease of administration, set the same root password on each node, if you have not already done so.**

Next Steps Configure Sun Cluster software on the cluster nodes. Go to “[Establishing a New Cluster or New Cluster Node](#)” on page 63.

Establishing the Cluster

This chapter provides procedures for how to establish a cluster or a new cluster node. The following procedures are in this chapter:

- “How to Configure Sun Cluster Software on All Nodes (`scinstall`)” on page 64
- “How to Configure Sun Cluster Software on All Nodes (XML)” on page 72
- “How to Install Solaris and Sun Cluster Software (JumpStart)” on page 79
- “How to Prepare the Cluster for Additional Cluster Nodes” on page 95
- “How to Change the Private Network Configuration When Adding Nodes or Private Networks” on page 98
- “How to Configure Sun Cluster Software on Additional Cluster Nodes (`scinstall`)” on page 103
- “How to Configure Sun Cluster Software on Additional Cluster Nodes (XML)” on page 110
- “How to Update Quorum Devices After Adding a Node to a Cluster” on page 114
- “How to Configure Quorum Devices” on page 117
- “How to Verify the Quorum Configuration and Installation Mode” on page 121
- “How to Change Private Hostnames” on page 122
- “How to Configure Network Time Protocol (NTP)” on page 123

Establishing a New Cluster or New Cluster Node

This section provides information and procedures to establish a new cluster or to add a node to an existing cluster. Before you start to perform these tasks, ensure that you installed software packages for the Solaris OS, Sun Cluster framework, and other products as described in “Installing the Software” on page 45.

The following task map lists the tasks to perform. Complete the procedures in the order that is indicated.

TABLE 3-1 Task Map: Establish the Cluster

Method	Instructions
1. Use one of the following methods to establish a new cluster or add a node to an existing cluster:	
<ul style="list-style-type: none"> ■ <i>(New clusters only)</i> Use the <code>scinstall</code> utility to establish the cluster. 	“How to Configure Sun Cluster Software on All Nodes (<code>scinstall</code>)” on page 64
<ul style="list-style-type: none"> ■ <i>(New clusters only)</i> Use an XML configuration file to establish the cluster. 	“How to Configure Sun Cluster Software on All Nodes (XML)” on page 72
<ul style="list-style-type: none"> ■ <i>(New clusters or added nodes)</i> Set up a JumpStart install server. Then create a flash archive of the installed system. Finally, use the <code>scinstall</code> JumpStart option to install the flash archive on each node and establish the cluster. 	“How to Install Solaris and Sun Cluster Software (JumpStart)” on page 79
<ul style="list-style-type: none"> ■ <i>(Added nodes only)</i> Use the <code>clsetup</code> command to add the new node to the cluster authorized-nodes list. If necessary, also configure the cluster interconnect and reconfigure the private network address range. Configure Sun Cluster software on a new node by using the <code>scinstall</code> utility or by using an XML configuration file. 	<p>“How to Prepare the Cluster for Additional Cluster Nodes” on page 95</p> <p>“How to Change the Private Network Configuration When Adding Nodes or Private Networks” on page 98</p> <p>“How to Configure Sun Cluster Software on Additional Cluster Nodes (<code>scinstall</code>)” on page 103</p> <p>“How to Configure Sun Cluster Software on Additional Cluster Nodes (XML)” on page 110</p>
2. If you added a node to a cluster, update the quorum configuration information.	“How to Update Quorum Devices After Adding a Node to a Cluster” on page 114
3. Assign quorum votes and remove the cluster from installation mode, if this operation was not already performed.	“How to Configure Quorum Devices” on page 117
4. Validate the quorum configuration.	“How to Verify the Quorum Configuration and Installation Mode” on page 121
5. <i>(Optional)</i> Change a node’s private hostname.	“How to Change Private Hostnames” on page 122
6. Create or modify the NTP configuration file, if not already configured.	“How to Configure Network Time Protocol (NTP)” on page 123

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

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

Note – This procedure uses the interactive form of the `scinstall` command. To use the noninteractive forms of the `scinstall` command, such as when developing installation scripts, see the `scinstall(1M)` man page.

Ensure that Sun Cluster software packages are installed on the node, either manually or by using the silent-mode form of the Java ES installer program, before you run the `scinstall` command. For information about running the Java ES installer program from an installation script, see Chapter 5, “Installing in Silent Mode,” in *Sun Java Enterprise System 5 Installation Guide for UNIX*.

Before You Begin Perform the following tasks:

- Ensure that the Solaris OS is installed to support Sun Cluster software.
If Solaris software is already installed on the node, you must ensure that the Solaris installation meets the requirements for Sun Cluster software and any other software that you intend to install on the cluster. See [“How to Install Solaris Software” on page 50](#) for more information about installing Solaris software to meet Sun Cluster software requirements.
- Ensure that Sun Cluster software packages and patches are installed on the node. See [“How to Install Sun Cluster Framework and Data-Service Software Packages” on page 58](#).
- Determine which mode of the `scinstall` utility you will use, Typical or Custom. For the Typical installation of Sun Cluster software, `scinstall` automatically specifies the following configuration defaults.

Component	Default Value
Private-network address	172.16.0.0
Private-network netmask	255.255.248.0
Cluster-transport switches	switch1 and switch2
Global-devices file-system name	/globaldevices
Installation security (DES)	Limited

- Complete one of the following cluster configuration worksheets, depending on whether you run the `scinstall` utility in Typical mode or Custom mode.
 - **Typical Mode Worksheet** - If you will use Typical mode and accept all defaults, complete the following worksheet.

Component	Description/Example	Answer
Cluster Name	What is the name of the cluster that you want to establish?	

Component	Description/Example	Answer	
Cluster Nodes	List the names of the other cluster nodes planned for the initial cluster configuration. <i>(For a single-node cluster, press Control-D alone.)</i>		
Cluster Transport Adapters and Cables	What are the names of the two cluster-transport adapters that attach the node to the private interconnect?	<i>First</i>	<i>Second</i>
<i>(VLAN adapters only)</i>	Will this be a dedicated cluster transport adapter? <i>(Answer No if using tagged VLAN adapters.)</i>	Yes No	Yes No
	If no, what is the VLAN ID for this adapter?		
Quorum Configuration <i>(two-node cluster only)</i>	Do you want to disable automatic quorum device selection? <i>(Answer Yes if any shared storage is not qualified to be a quorum device or if you want to configure a quorum server or a Network Appliance NAS device as a quorum device.)</i>	Yes No	
Check	Do you want to interrupt cluster creation for sccheck errors?	Yes No	

- **Custom Mode Worksheet** - If you will use Custom mode and customize the configuration data, complete the following worksheet.

Note – If you are installing a single-node cluster, the `scinstall` utility automatically assigns the default private network address and netmask, even though the cluster does not use a private network.

Component	Description/Example	Answer
Cluster Name	What is the name of the cluster that you want to establish?	
Cluster Nodes	List the names of the other cluster nodes planned for the initial cluster configuration. <i>(For a single-node cluster, press Control-D alone.)</i>	
Authenticating Requests to Add Nodes <i>(multiple-node cluster only)</i>	Do you need to use DES authentication?	No Yes
Network Address for the Cluster Transport <i>(multiple-node cluster only)</i>	Do you want to accept the default network address (172.16.0.0)?	Yes No
	If no, which private network address do you want to use?	____.____.____.____
	Do you want to accept the default netmask (255.255.248.0)?	Yes No
	If no, what are the maximum number of nodes and private networks that you expect to configure in the cluster?	____ nodes ____ networks
	Which netmask do you want to use? <i>Choose from the values calculated by scinstall or supply your own.</i>	____.____.____.____

Component	Description/Example	Answer	
Point-to-Point Cables (multiple-node cluster only)	If this is a two-node cluster, does this cluster use switches?	Yes No	
Cluster Switches (multiple-node cluster only)	If used, what are the names of the two switches? Defaults: <code>switch1</code> and <code>switch2</code>	First	Second
Cluster Transport Adapters and Cables (multiple-node cluster only)	Node name (<i>the node from which you run <code>scinstall</code></i>):		
	Transport adapters:	First	Second
(VLAN adapters only)	Will this be a dedicated cluster transport adapter? (<i>Answer No if using tagged VLAN adapters.</i>)	Yes No	Yes No
	If no, what is the VLAN ID for this adapter?		
	Where does each transport adapter connect to (<i>a switch or another adapter</i>)? Switch defaults: <code>switch1</code> and <code>switch2</code>		
	For transport switches, do you want to use the default port name?	Yes No	Yes No
	If no, what is the name of the port that you want to use?		
	Do you want to use autodiscovery to list the available adapters for the other nodes? If no, supply the following information for each additional node:	Yes No	
Specify for each additional node (multiple-node cluster only)	Node name:		
	Transport adapters:	First	Second
(VLAN adapters only)	Will this be a dedicated cluster transport adapter? (<i>Answer No if using tagged VLAN adapters.</i>)	Yes No	Yes No
	If no, what is the VLAN ID for this adapter?		
	Where does each transport adapter connect to (<i>a switch or another adapter</i>)? Defaults: <code>switch1</code> and <code>switch2</code>		
	For transport switches, do you want to use the default port name?	Yes No	Yes No
	If no, what is the name of the port that you want to use?		
Quorum Configuration (two-node cluster only)	Do you want to disable automatic quorum device selection? (<i>Answer Yes if any shared storage is not qualified to be a quorum device or if you want to configure a quorum server or a Network Appliance NAS device as a quorum device.</i>)	Yes No	Yes No

Component	Description/Example	Answer
Global Devices File System (specify for each node)	Do you want to use the default name of the global-devices file system (/globaldevices)?	Yes No
	If no, do you want to use an already-existing file system?	Yes No
	What is the name of the file system that you want to use?	
Check (multiple-node cluster only)	Do you want to interrupt cluster creation for sccheck errors?	Yes No
(single-node cluster only)	Do you want to run the sccheck utility to validate the cluster?	Yes No
Automatic Reboot (single-node cluster only)	Do you want scinstall to automatically reboot the node after installation?	Yes No

Follow these guidelines to use the interactive `scinstall` utility in this procedure:

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

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

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

2 On the cluster node from which you intend to configure the cluster, become superuser.

3 Start the `scinstall` utility.

```
phys-schost# /usr/cluster/bin/scinstall
```

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

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

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

- * 1) Create a new cluster or add a cluster node
- 2) Configure a cluster to be JumpStarted from this install server
- 3) Manage a dual-partition upgrade
- 4) Upgrade this cluster node
- * 5) Print release information for this cluster node

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

Option: **1**

The New Cluster and Cluster Node Menu is displayed.

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

The Typical or Custom Mode menu is displayed.

6 Type the number that corresponds to the option for either Typical or Custom and press the Return key.

The Create a New Cluster screen is displayed. Read the requirements, then press Control-D to continue.

7 Follow the menu prompts to supply your answers from the configuration planning worksheet.

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

8 For the Solaris 10 OS, verify on each node that multi-user services for the Service Management Facility (SMF) are online.

If services are not yet online for a node, wait until the state becomes online before you proceed to the next step.

```
phys-schost# svcs multi-user-server
STATE          STIME      FMRI
online         17:52:55  svc:/milestone/multi-user-server:default
```

9 On one node, become superuser.

10 Verify that all nodes have joined the cluster.

```
phys-schost# clnode status
```

Output resembles the following.

```
=== Cluster Nodes ===
```

```
--- Node Status ---
```

Node Name	Status
-----	-----
phys-schost-1	Online
phys-schost-2	Online
phys-schost-3	Online

For more information, see the `clnode(1CL)` man page.

11 (Optional) Enable automatic node reboot if all monitored disk paths fail.**a. Enable the automatic reboot feature.**

```
phys-schost# clnode set -p reboot_on_path_failure=enabled
```

-p Specifies the property to set

reboot_on_path_failure=enable Specifies that the node will reboot if all monitored disk paths fail, provided that at least one of the disks is accessible from a different node in the cluster.

b. Verify that automatic reboot on disk-path failure is enabled.

```
phys-schost# clnode show
```

```
=== Cluster Nodes ===
```

```
Node Name:                               node
...
reboot_on_path_failure:                   enabled
...
```

12 If you intend to use Sun Cluster HA for NFS on a highly available local file system, ensure that the loopback file system (LOFS) is disabled.

To disable LOFS, add the following entry to the `/etc/system` file on each node of the cluster.

```
exclude:lofs
```

The change to the `/etc/system` file becomes effective after the next system reboot.

Note – You cannot have LOFS enabled if you use Sun Cluster HA for NFS on a highly available local file system *and* have `automountd` running. LOFS can cause switchover problems for Sun Cluster HA for NFS. If you choose to add Sun Cluster HA for NFS on a highly available local file system, you must make one of the following configuration changes.

However, if you configure non-global zones in your cluster, you must enable LOFS on all cluster nodes. If Sun Cluster HA for NFS on a highly available local file system must coexist with LOFS, use one of the other solutions instead of disabling LOFS.

- Disable LOFS.
- Disable the `automountd` daemon.
- Exclude from the automounter map all files that are part of the highly available local file system that is exported by Sun Cluster HA for NFS. This choice enables you to keep both LOFS and the `automountd` daemon enabled.

See “The Loopback File System” in *System Administration Guide: Devices and File Systems* (Solaris 9 or Solaris 10) for more information about loopback file systems.

Example 3-1 Configuring Sun Cluster Software on All Nodes

The following example shows the `scinstall` progress messages that are logged as `scinstall` completes configuration tasks on the two-node cluster, `schost`. The cluster is installed from `phys-schost-1` by using the `scinstall` Typical mode. The other cluster node is `phys-schost-2`. The adapter names are `qfe2` and `qfe3`. The automatic selection of a quorum device is enabled.

Installation and Configuration

```
Log file - /var/cluster/logs/install/scinstall.log.24747
```

```
Testing for "/globaldevices" on "phys-schost-1" ... done
Testing for "/globaldevices" on "phys-schost-2" ... done
Checking installation status ... done
```

```
The Sun Cluster software is already installed on "phys-schost-1".
The Sun Cluster software is already installed on "phys-schost-2".
Starting discovery of the cluster transport configuration.
```

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

```
phys-schost-1:qfe2  switch1  phys-schost-2:qfe2
phys-schost-1:qfe3  switch2  phys-schost-2:qfe3
```

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

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

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

```
Removing the downloaded files ... done
```

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

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

```
Log file - /var/cluster/logs/install/scinstall.log.24747
```

```
Rebooting ...
```

Troubleshooting **Unsuccessful configuration** - If one or more nodes cannot join the cluster, or if the wrong configuration information was specified, first attempt to rerun this procedure. If that does not correct the problem, perform the procedure [“How to Unconfigure Sun Cluster Software to Correct](#)

[Installation Problems](#)” on page 263 on each misconfigured node to remove it from the cluster configuration. You do not need to uninstall the Sun Cluster software packages. Then rerun this procedure.

- Next Steps**
- If you installed a single-node cluster, cluster establishment is complete. Go to [“Creating Cluster File Systems”](#) on page 175 to install volume management software and configure the cluster.
 - If you installed a multiple-node cluster and chose automatic quorum configuration, postinstallation setup is complete. Go to [“How to Verify the Quorum Configuration and Installation Mode”](#) on page 121.
 - If you installed a multiple-node cluster and declined automatic quorum configuration, perform postinstallation setup. Go to [“How to Configure Quorum Devices”](#) on page 117.

▼ **How to Configure Sun Cluster Software on All Nodes (XML)**

Perform this procedure to configure a new cluster by using an XML cluster configuration file. The new cluster can be a duplication of an existing cluster that runs Sun Cluster 3.2 software.

This procedure configures the following cluster components:

- Cluster name
- Cluster node membership
- Cluster interconnect
- Global devices

Before You Begin Perform the following tasks:

- Ensure that the Solaris OS is installed to support Sun Cluster software.
If Solaris software is already installed on the node, you must ensure that the Solaris installation meets the requirements for Sun Cluster software and any other software that you intend to install on the cluster. See [“How to Install Solaris Software”](#) on page 50 for more information about installing Solaris software to meet Sun Cluster software requirements.
Ensure that the Solaris OS is installed to support Sun Cluster software.
If Solaris software is already installed on the node, you must ensure that the Solaris installation meets the requirements for Sun Cluster software and any other software that you intend to install on the cluster. See [“How to Install Solaris Software”](#) on page 50 for more information about installing Solaris software to meet Sun Cluster software requirements.
- Ensure that Sun Cluster 3.2 software and patches are installed on each node that you will configure. See [“How to Install Sun Cluster Framework and Data-Service Software Packages”](#) on page 58.

- 1 **Ensure that Sun Cluster 3.2 software is not yet configured on each potential cluster node.**
 - a. **Become superuser on a potential node that you want to configure in the new cluster.**
 - b. **Determine whether Sun Cluster 3.2 software is already configured on the potential node.**

```
phys-schost# /usr/sbin/clinfo -n
```

 - **If the command returns the following message, proceed to Step c.**

```
clinfo: node is not configured as part of acluster: Operation not applicable
```

This message indicates that Sun Cluster software is not yet configured on the potential node.
 - **If the command returns the node ID number, do not perform this procedure.**

The return of a node ID indicates that Sun Cluster software is already configured on the node.

If the cluster is running an older version of Sun Cluster software and you want to install Sun Cluster 3.2 software, instead perform upgrade procedures in [Chapter 8](#).
 - c. **Repeat Step a and Step b on each remaining potential node that you want to configure in the new cluster.**

If Sun Cluster 3.2 software is not yet configured on any of the potential cluster nodes, proceed to Step 2.
- 2 **If you are duplicating an existing cluster than runs Sun Cluster 3.2 software, use a node in that cluster to create a cluster configuration XML file.**
 - a. **Become superuser on an active member of the cluster that you want to duplicate.**
 - b. **Export the existing cluster's configuration information to a file.**

```
phys-schost# cluster export -o clconfigfile
```

`-o` Specifies the output destination.

`clconfigfile` The name of the cluster configuration XML file. The specified file name can be an existing file or a new file that the command will create.

For more information, see the `cluster(1CL)` man page.
 - c. **Copy the configuration file to the potential node from which you will configure the new cluster.**

You can store the file in any directory that is accessible to the other hosts that you will configure as cluster nodes.
- 3 **Become superuser on the potential node from which you will configure the new cluster.**

4 Modify the cluster configuration XML file as needed.

a. Open your cluster configuration XML file for editing.

- **If you are duplicating an existing cluster, open the file that you created with the `cluster export` command.**
- **If you are not duplicating an existing cluster, create a new file.**

Base the file on the element hierarchy that is shown in the `clconfiguration(5CL)` man page. You can store the file in any directory that is accessible to the other hosts that you will configure as cluster nodes.

b. Modify the values of the XML elements to reflect the cluster configuration that you want to create.

- To establish a cluster, the following components must have valid values in the cluster configuration XML file:
 - Cluster name
 - Cluster nodes
 - Cluster transport
- The cluster is created with the assumption that the partition `/globaldevices` exists on each node that you configure as a cluster node. The `global-devices` namespace is created on this partition. If you need to use a different file-system name on which to create the global devices, add the following property to the `<propertyList>` element for each node that does not have a partition that is named `/globaldevices`.

```
...
<nodeList>
  <node name="node" id="N">
    <propertyList>
      ...
      <property name="globaldevfs" value="/filesystem-name"/>
      ...
    </propertyList>
  </node>
  ...

```

- If you are modifying configuration information that was exported from an existing cluster, some values that you must change to reflect the new cluster, such as node names, are used in the definitions of more than one cluster object.

See the `clconfiguration(5CL)` man page for details about the structure and content of the cluster configuration XML file.

5 Validate the cluster configuration XML file.

```
phys-schost# /usr/share/src/xmllint --valid --noout clconfigfile
```

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

6 From the potential node that contains the cluster configuration XML file, create the cluster.

```
phys-schost# cluster create -i clconfigfile
```

`-i clconfigfile` Specifies the name of the cluster configuration XML file to use as the input source.

7 For the Solaris 10 OS, verify on each node that multi-user services for the Service Management Facility (SMF) are online.

If services are not yet online for a node, wait until the state becomes online before you proceed to the next step.

```
phys-schost# svcs multi-user-server
STATE      STIME      FMRI
online     17:52:55   svc:/milestone/multi-user-server:default
```

8 On one node, become superuser.

9 Verify that all nodes have joined the cluster.

```
phys-schost# clnode status
```

Output resembles the following.

```
=== Cluster Nodes ===

--- Node Status ---

Node Name                               Status
-----
phys-schost-1                           Online
phys-schost-2                           Online
phys-schost-3                           Online
```

For more information, see the `clnode(1CL)` man page.

10 Install any necessary patches to support Sun Cluster software, if you have not already done so.

See “Patches and Required Firmware Levels” in *Sun Cluster 3.2 Release Notes for Solaris OS* for the location of patches and installation instructions.

11 If you intend to use Sun Cluster HA for NFS on a highly available local file system, ensure that the loopback file system (LOFS) is disabled.

To disable LOFS, add the following entry to the `/etc/system` file on each node of the cluster.

```
exclude:lofs
```

The change to the `/etc/system` file becomes effective after the next system reboot.

Note – You cannot have LOFS enabled if you use Sun Cluster HA for NFS on a highly available local file system *and* have automountd running. LOFS can cause switchover problems for Sun Cluster HA for NFS. If you choose to add Sun Cluster HA for NFS on a highly available local file system, you must make one of the following configuration changes.

However, if you configure non-global zones in your cluster, you must enable LOFS on all cluster nodes. If Sun Cluster HA for NFS on a highly available local file system must coexist with LOFS, use one of the other solutions instead of disabling LOFS.

- Disable LOFS.
- Disable the automountd daemon.
- Exclude from the automounter map all files that are part of the highly available local file system that is exported by Sun Cluster HA for NFS. This choice enables you to keep both LOFS and the automountd daemon enabled.

See “The Loopback File System” in *System Administration Guide: Devices and File Systems* (Solaris 9 or Solaris 10) for more information about loopback file systems.

12 To duplicate quorum information from an existing cluster, configure the quorum device by using the cluster configuration XML file.

You must configure a quorum device if you created a two-node cluster. If you choose not to use the cluster configuration XML file to create a required quorum device, go instead to “[How to Configure Quorum Devices](#)” on page 117.

a. If you are using a quorum server for the quorum device, ensure that the quorum server is set up and running.

Follow instructions in *Sun Cluster Quorum Server User’s Guide*.

b. If you are using a Network Appliance NAS device for the quorum device, ensure that the NAS device is set up and operational.

i. Observe the requirements for using a NAS device as a quorum device.

See “Requirements, Recommendations, and Restrictions for Network Appliance NAS Devices” in *Sun Cluster 3.1 - 3.2 With Network-Attached Storage Devices Manual for Solaris OS*.

ii. Follow instructions in your device’s documentation to set up the NAS device.

c. Ensure that the quorum configuration information in the cluster configuration XML file reflects valid values for the cluster that you created.

d. If you made changes to the cluster configuration XML file, validate the file.

```
phys-schost# xmllint --valid --noout clconfigfile
```

e. Configure the quorum device.

```
phys-schost# clquorum add -i clconfigfile devicename
```

devicename Specifies the name of the device to configure as a quorum device.

13 Remove the cluster from installation mode.

```
phys-schost# clquorum reset
```

14 (Optional) Enable automatic node reboot if all monitored disk paths fail.**a. Enable the automatic reboot feature.**

```
phys-schost# clnode set -p reboot_on_path_failure=enabled
```

-p Specifies the property to set

reboot_on_path_failure=enable Specifies that the node will reboot if all monitored disk paths fail, provided that at least one of the disks is accessible from a different node in the cluster.

b. Verify that automatic reboot on disk-path failure is enabled.

```
phys-schost# clnode show
```

```
=== Cluster Nodes ===
```

```
Node Name:                               node
...
reboot_on_path_failure:                   enabled
...
```

Example 3-2 Configuring Sun Cluster Software on All Nodes By Using an XML File

The following example duplicates the cluster configuration and quorum configuration of an existing two-node cluster to a new two-node cluster. The new cluster is installed with the Solaris 10 OS and is not configured with non-global zones. The cluster configuration is exported from the existing cluster node, `phys-oldhost-1`, to the cluster configuration XML file `clusterconf.xml`. The node names of the new cluster are `phys-newhost-1` and `phys-newhost-2`. The device that is configured as a quorum device in the new cluster is `d3`.

The prompt name `phys-newhost-N` in this example indicates that the command is performed on both cluster nodes.

```
phys-newhost-N# /usr/sbin/clinfo -n
clinfo: node is not configured as part of acluster: Operation not applicable
```

```
phys-oldhost-1# cluster export -o clusterconf.xml
```

Copy clusterconf.xml to phys-newhost-1 and modify the file with valid values

```
phys-newhost-1# xmllint --valid --noout clusterconf.xml
```

No errors are reported

```
phys-newhost-1# cluster create -i clusterconf.xml
```

```
phys-newhost-1# svcs multi-user-server
```

```
STATE          STIME          FMRI
```

```
online         17:52:55      svc:/milestone/multi-user-server:default
```

```
phys-newhost-1# clnode status
```

Output shows that both nodes are online

```
phys-newhost-1# clquorum add -i clusterconf.xml d3
```

```
phys-newhost-1# clquorum reset
```

Troubleshooting **Unsuccessful configuration** - If one or more nodes cannot join the cluster, or if the wrong configuration information was specified, first attempt to rerun this procedure. If that does not correct the problem, perform the procedure [“How to Unconfigure Sun Cluster Software to Correct Installation Problems”](#) on page 263 on each misconfigured node to remove it from the cluster configuration. You do not need to uninstall the Sun Cluster software packages. Then rerun this procedure.

Next Steps Go to [“How to Verify the Quorum Configuration and Installation Mode”](#) on page 121.

See Also After the cluster is fully established, you can duplicate the configuration of the other cluster components from the existing cluster. If you did not already do so, modify the values of the XML elements that you want to duplicate to reflect the cluster configuration you are adding the component to. For example, if you are duplicating resource groups, ensure that the `<resourcegroupNodeList>` entry contains the valid node names for the new cluster, and not the node names from the cluster that you duplicated unless the node names are the same.

To duplicate a cluster component, run the `export` subcommand of the object-oriented command for the cluster component that you want to duplicate. For more information about the command syntax and options, see the man page for the cluster object that you want to duplicate. The following table lists the cluster components that you can create from a cluster configuration XML file after the cluster is established and the man page for the command that you use to duplicate the component.

Note – This table provides the long forms of the Sun Cluster commands. Most commands also have short forms. Except for the forms of the command names, the commands are identical. For a list of the commands and their short forms, see Appendix A, “Sun Cluster Object-Oriented Commands,” in *Sun Cluster System Administration Guide for Solaris OS*.

Cluster Component	Man Page	Special Instructions
Device groups: Solaris Volume Manager and VERITAS Volume Manager	cldevicegroup(1CL)	For Solaris Volume Manager, first create the disk sets that you specify in the cluster configuration XML file. For VxVM, first install and configure VxVM software and create the disk groups that you specify in the cluster configuration XML file.
Resources	clresource(1CL)	You can use the <code>-a</code> option of the <code>clresource</code> , <code>clressharedaddress</code> , or <code>clreslogicalhostname</code> command to also duplicate the resource type and resource group that are associated with the resource that you duplicate. Otherwise, you must first add the resource type and resource group to the cluster before you add the resource.
Shared address resources	clressharedaddress(1CL)	
Logical hostname resources	clreslogicalhostname(1CL)	
Resource types	clresourcetype(1CL)	
Resource groups	clresourcegroup(1CL)	
NAS devices	clnasdevice(1CL)	You must first set up the NAS device as described in the device's documentation.
SNMP hosts	clsnmphost(1CL)	The <code>clsnmphost create -i</code> command requires that you specify a user password file with the <code>-f</code> option.
SNMP users	clsnmpuser(1CL)	
Thresholds for monitoring system resources on cluster objects	cltelemetryattribute(1CL)	

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

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

Before You Begin Perform the following tasks:

- Ensure that the hardware setup is complete and connections are verified before you install Solaris software. See the *Sun Cluster Hardware Administration Collection* and your server and storage device documentation for details on how to set up the hardware.
- Determine the Ethernet address of each cluster node.
- If you use a naming service, ensure that the following information is added to any naming services that clients use to access cluster services. See “[Public Network IP Addresses](#)” on [page 21](#) for planning guidelines. See your Solaris system-administrator documentation for information about using Solaris naming services.

- Address-to-name mappings for all public hostnames and logical addresses
- The IP address and hostname of the JumpStart install server
- Ensure that your cluster configuration planning is complete. See [“How to Prepare for Cluster Software Installation” on page 46](#) for requirements and guidelines.
- On the server from which you will create the flash archive, ensure that all Solaris OS software, patches, and firmware that is necessary to support Sun Cluster software is installed.
 If Solaris software is already installed on the server, you must ensure that the Solaris installation meets the requirements for Sun Cluster software and any other software that you intend to install on the cluster. See [“How to Install Solaris Software” on page 50](#) for more information about installing Solaris software to meet Sun Cluster software requirements.
- Ensure that Sun Cluster software packages and patches are installed on the server from which you will create the flash archive. See [“How to Install Sun Cluster Framework and Data-Service Software Packages” on page 58](#).
- Determine which mode of the `scinstall` utility you will use, Typical or Custom. For the Typical installation of Sun Cluster software, `scinstall` automatically specifies the following configuration defaults.

Component	Default Value
Private-network address	172.16.0.0
Private-network netmask	255.255.248.0
Cluster-transport switches	switch1 and switch2
Global-devices file-system name	/globaldevices
Installation security (DES)	Limited

- Complete one of the following cluster configuration worksheets, depending on whether you run the `scinstall` utility in Typical mode or Custom mode. See [“Planning the Sun Cluster Environment” on page 20](#) for planning guidelines.
 - **Typical Mode Worksheet** - If you will use Typical mode and accept all defaults, complete the following worksheet.

Component	Description/Example	Answer
JumpStart Directory	What is the name of the JumpStart directory to use?	
Cluster Name	What is the name of the cluster that you want to establish?	
Cluster Nodes	List the names of the cluster nodes that are planned for the initial cluster configuration. <i>(For a single-node cluster, press Control-D alone.)</i>	

Component	Description/Example	Answer	
Cluster Transport Adapters and Cables	First node name:		
	Transport adapters:	<i>First</i>	<i>Second</i>
<i>VLAN adapters only</i>	Will this be a dedicated cluster transport adapter? (<i>Answer No if using tagged VLAN adapters.</i>)	Yes No	Yes No
	If no, what is the VLAN ID for this adapter?		
<i>Specify for each additional node</i>	Node name:		
	Transport adapters:	<i>First</i>	<i>Second</i>
Quorum Configuration <i>(two-node cluster only)</i>	Do you want to disable automatic quorum device selection? (<i>Answer Yes if any shared storage is not qualified to be a quorum device or if you want to configure a quorum server or a Network Appliance NAS device as a quorum device.</i>)	Yes No	Yes No

- **Custom Mode Worksheet** - If you will use Custom mode and customize the configuration data, complete the following worksheet.

Note – If you are installing a single-node cluster, the `scinstall` utility automatically uses the default private network address and netmask, even though the cluster does not use a private network.

Component	Description/Example	Answer
JumpStart Directory	What is the name of the JumpStart directory to use?	
Cluster Name	What is the name of the cluster that you want to establish?	
Cluster Nodes	List the names of the cluster nodes that are planned for the initial cluster configuration. (<i>For a single-node cluster, press Control-D alone.</i>)	
Authenticating Requests to Add Nodes <i>(multiple-node cluster only)</i>	Do you need to use DES authentication?	No Yes

Component	Description/Example	Answer	
Network Address for the Cluster Transport <i>(multiple-node cluster only)</i>	Do you want to accept the default network address (172.16.0.0)?	Yes No	
	If no, which private network address do you want to use?	____.____.____.____	
	Do you want to accept the default netmask (255.255.248.0)?	Yes No	
	If no, what are the maximum number of nodes and private networks that you expect to configure in the cluster?	____ nodes	____ networks
	Which netmask do you want to use? <i>Choose from the values that are calculated by scinstall or supply your own.</i>	____.____.____.____	
Point-to-Point Cables <i>(two-node cluster only)</i>	Does this cluster use switches?	Yes No	
Cluster Switches <i>(multiple-node cluster only)</i>	If used, what are the names of the two switches? Defaults: swi t ch1 and swi t ch2	<i>First</i>	<i>Second</i>
Cluster Transport Adapters and Cables <i>(multiple-node cluster only)</i>	First node name:		
	Transport adapters:	<i>First</i>	<i>Second</i>
<i>(VLAN adapters only)</i>	Will this be a dedicated cluster transport adapter? <i>(Answer No if using tagged VLAN adapters.)</i>	Yes No	Yes No
	If no, what is the VLAN ID for this adapter?		
	Where does each transport adapter connect to <i>(a switch or another adapter)</i> ? Switch defaults: swi t ch1 and swi t ch2		
	For transport switches, do you want to use the default port name?	Yes No	Yes No
	If no, what is the name of the port that you want to use?		
<i>Specify for each additional node</i> <i>(multiple-node cluster only)</i>	Node name:		
	Transport adapters:	<i>First</i>	<i>Second</i>
	Where does each transport adapter connect to <i>(a switch or another adapter)</i> ? Switch defaults: swi t ch1 and swi t ch2		
	For transport switches, do you want to use the default port name?	Yes No	Yes No
	If no, what is the name of the port that you want to use?		

Component	Description/Example	Answer	
Global Devices File System <i>Specify for each node</i>	Do you want to use the default name of the global-devices file system (/globaldevices)?	Yes No	
	If no, do you want to use an already-existing file system?	Yes No	
	If no, do you want to create a new file system on an unused partition?	Yes No	
	What is the name of the file system?		
Quorum Configuration <i>(two-node cluster only)</i>	Do you want to disable automatic quorum device selection? <i>(Answer Yes if any shared storage is not qualified to be a quorum device or if you want to configure a quorum server or a Network Appliance NAS device as a quorum device.)</i>	Yes No	Yes No

Follow these guidelines to use the interactive `scinstall` utility in this procedure:

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

1 Set up your JumpStart install server.

Ensure that the JumpStart install server meets the following requirements.

- The install server is on the same subnet as the cluster nodes, or on the Solaris boot server for the subnet that the cluster nodes use.
- The install server is not itself a cluster node.
- The install server installs a release of the Solaris OS that is supported by the Sun Cluster software.
- A custom JumpStart directory exists for JumpStart installation of Sun Cluster software. This `jumpstart-dir` directory must meet the following requirements:
 - Contain a copy of the check utility.
 - Be NFS exported for reading by the JumpStart install server.
- Each new cluster node is configured as a custom JumpStart installation client that uses the custom JumpStart directory that you set up for Sun Cluster installation.

Follow the appropriate instructions for your software platform and OS version to set up the JumpStart install server. See “Creating a Profile Server for Networked Systems” in *Solaris 9 9/04 Installation Guide* or “Creating a Profile Server for Networked Systems” in *Solaris 10 Installation Guide: Custom JumpStart and Advanced Installations*.

See also the `setup_install_server(1M)` and `add_install_client(1M)` man pages.

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

- a. **Switch to another cluster node that is active and start the `clsetup` utility.**
- b. **Use the `clsetup` utility to add the new node's name to the list of authorized cluster nodes.**

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

3 On a cluster node or another machine of the same server platform, install the Solaris OS and any necessary patches, if you have not already done so.

If Solaris software is already installed on the server, you must ensure that the Solaris installation meets the requirements for Sun Cluster software and any other software that you intend to install on the cluster. See “How to Install Solaris Software” on page 50 for more information about installing Solaris software to meet Sun Cluster software requirements.

Follow procedures in “How to Install Solaris Software” on page 50.

4 On the installed system, install Sun Cluster software and any necessary patches, if you have not done so already.

Follow procedures in “How to Install Sun Cluster Framework and Data-Service Software Packages” on page 58.

See “Patches and Required Firmware Levels” in *Sun Cluster 3.2 Release Notes for Solaris OS* for the location of patches and installation instructions.

5 Enable the common agent container daemon to start automatically during system boots.

```
machine# cacaoadm enable
```

6 On the installed system, update the `/etc/inet/ipnodes` file with all public IP addresses that are used in the cluster.

Perform this step regardless of whether you are using a naming service. See “Public Network IP Addresses” on page 21 for a listing of Sun Cluster components whose IP addresses you must add.

7 Create the flash archive of the installed system.

```
machine# flarcreate -n name archive
```

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

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

Follow procedures in one of the following manuals:

- Chapter 21, “Creating Solaris Flash Archives (Tasks),” in *Solaris 9 9/04 Installation Guide*
- Chapter 3, “Creating Solaris Flash Archives (Tasks),” in *Solaris 10 Installation Guide: Solaris Flash Archives (Creation and Installation)*

8 Ensure that the flash archive is NFS exported for reading by the JumpStart install server.

See “Managing Network File Systems (Overview),” in *System Administration Guide: Network Services* (Solaris 9 or Solaris 10) for more information about automatic file sharing.

See also the `share(1M)` and `dfstab(4)` man pages.

9 On the JumpStart install server, become superuser.

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

The path `/export/suncluster/sc31/` is used here as an example of the JumpStart installation directory that you created. In the media path, replace `arch` with `sparc` or `x86` (Solaris 10 only) and replace `ver` with `9` for Solaris 9 or `10` for Solaris 10.

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

The `scinstall` Main Menu is displayed.

11 Type the number that corresponds to the option for Configure a cluster to be JumpStarted from this install server and press the Return key.

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

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

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

- ```
* 1) Create a new cluster or add a cluster node
* 2) Configure a cluster to be JumpStarted from this install server
 3) Manage a dual-partition upgrade
 4) Upgrade this cluster node
* 5) Print release information for this cluster node

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

```
Option: 2
```

**12 Follow the menu prompts to supply your answers from the configuration planning worksheet.**

The `scinstall` command stores your configuration information and copies the `autoscinstall.class` default class file in the `/jumpstart-dir/autoscinstall.d/3.2/` directory. This file is similar to the following example.

```
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
```

**13 If necessary, make adjustments to the `autoscinstall.class` file to configure JumpStart to install the flash archive.****a. Modify entries as necessary to match configuration choices that you made when you installed the Solaris OS on the flash archive machine or when you ran the `scinstall` utility.**

For example, if you assigned slice 4 for the global-devices file system and specified to `scinstall` that the file-system name is `/gdevs`, you would change the `/globaldevices` entry of the `autoscinstall.class` file to the following:

```
filesystems rootdisk.s4 512 /gdevs
```

**b. Change the following entries in the `autoscinstall.class` file.**

| Existing Entry to Replace |                              | New Entry to Add              |                                      |
|---------------------------|------------------------------|-------------------------------|--------------------------------------|
| <code>install_type</code> | <code>initial_install</code> | <code>install_type</code>     | <code>flash_install</code>           |
| <code>system_type</code>  | <code>standalone</code>      | <code>archive_location</code> | <code>retrieval_type location</code> |

See “archive\_location Keyword” in *Solaris 9 9/04 Installation Guide* or *Solaris 10 Installation Guide: Custom JumpStart and Advanced Installations* for information about valid values for `retrieval_type` and `location` when used with the `archive_location` keyword.

**c. Remove all entries that would install a specific package, such as the following entries.**

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

**d. If your configuration has additional Solaris software requirements, change the `autoscinstall.class` file accordingly.**

The `autoscinstall.class` file installs the End User Solaris Software Group (SUNWCuser).

**e. If you install the End User Solaris Software Group (SUNWcuser), add to the `autoscinstall.class` file any additional Solaris software packages that you might need.**

The following table lists Solaris packages that are required to support some Sun Cluster functionality. These packages are not included in the End User Solaris Software Group. See “Solaris Software Group Considerations” on page 16 for more information.

| Feature                                                                 | Mandatory Solaris Software Packages                                                                   |
|-------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------|
| RSMAPI, RSMRDT drivers, or SCI-PCI adapters (SPARC based clusters only) | SPARC: <i>Solaris 9</i> : SUNWrsn SUNWrsmx SUNWrsmo SUNWrsmax<br><i>Solaris 10</i> : SUNWrsn SUNWrsmo |
| Sun Cluster Manager (formerly SunPlex™ Manager)                         | SUNWapchr SUNWapchu                                                                                   |

You can change the default `class` file 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 OS installation profile meets minimum Sun Cluster file-system allocation requirements, Sun Cluster software places no restrictions on other changes to the installation profile. See “System Disk Partitions” on page 16 for partitioning guidelines and requirements to support Sun Cluster software.

For more information about JumpStart profiles, see Chapter 26, “Preparing Custom JumpStart Installations (Tasks),” in *Solaris 9 9/04 Installation Guide* or Chapter 6, “Preparing Custom JumpStart Installations (Tasks),” in *Solaris 10 Installation Guide: Custom JumpStart and Advanced Installations*.

**14 To install required packages for any of the following features or to perform other postinstallation tasks, set up your own finish script.**

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

---

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

---

Your own finish script runs after the standard finish script that is installed by the `scinstall` command. See Preparing Custom JumpStart Installations in Chapter 26, “Preparing Custom JumpStart Installations (Tasks),” in *Solaris 9 9/04 Installation Guide* or Chapter 6, “Preparing Custom

JumpStart Installations (Tasks),” in *Solaris 10 Installation Guide: Custom JumpStart and Advanced Installations* for information about creating a JumpStart finish script.

- a. **Ensure that any dependency Solaris packages will be installed by the default `class` file.**

See [Step 13](#).

- b. **Name your finish script `finish`.**

- c. **Modify the `finish` script to install the software packages listed in the following table that support the features that you intend to use.**

| Feature          | Additional Sun Cluster 3.2 Packages to Install                                                                                                                     |
|------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| RSMAPI           | SUNWscrif                                                                                                                                                          |
| SCI-PCI adapters | <ul style="list-style-type: none"> <li>■ <i>Solaris 9</i>: SUNWsci SUNWscid SUNWscidx</li> <li>■ <i>Solaris 10</i>: SUNWscir SUNWsci SUNWscidr SUNWscid</li> </ul> |
| RSMRDT drivers   | SUNWscrdt                                                                                                                                                          |

- Install the packages in the order that is used in the table.
- Install the packages from the `Solaris_arch/Product/sun_cluster/Solaris_ver/Packages/` directory, where *arch* is `sparc` or `x86` (Solaris 10 only), and where *ver* is 9 for Solaris 9 or 10 for Solaris 10.

- d. **Make any additional modifications for other postinstallation tasks that you want the `finish` script to perform.**

- e. **Copy your `finish` script to each `jumpstart-dir/autosinstall.d/nodes/node` directory.**

Create one *node* directory for each node in the cluster. Or, use this naming convention to create symbolic links to a shared `finish` script.

- 15 Exit from the JumpStart install server.**

- 16 If you are using a cluster administrative console, display a console screen for each node in the cluster.**

- **If Cluster Control Panel (CCP) software is installed and configured on your administrative console, use the `cconsole(1M)` utility to display the individual console screens.**

As superuser, use the following command to start the `cconsole` utility:

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

The `cconsole` utility 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 the `cconsole` utility, connect to the consoles of each node individually.**



**17 Shut down each node.**

```
phys-schost# shutdown -g0 -y -i0
```

**18 Boot each node to start the JumpStart installation.**

- **On SPARC based systems, do the following:**

```
ok boot net - install
```

---

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

---

- **On x86 based systems, do the following:**

- a. **Press any key to begin the booting sequence.**

Press any key to reboot.  
*keystroke*

- b. **As soon as the BIOS information screen appears, immediately press Esc+2 or press the F2 key.**

After the initialization sequence completes, the BIOS Setup Utility screen appears.

- c. **In the BIOS Setup Utility menu bar, navigate to the Boot menu item.**

The list of boot devices is displayed.

- d. **Navigate to the listed IBA that is connected to the same network as the JumpStart PXE install server and move it to the top of the boot order.**

The lowest number to the right of the IBA boot choices corresponds to the lower Ethernet port number. The higher number to the right of the IBA boot choices corresponds to the higher Ethernet port number.

- e. **Save your change and exit the BIOS.**

The boot sequence begins again. After further processing, the GRUB menu is displayed.

- f. **Immediately select the Solaris JumpStart entry and press Enter.**

---

**Note** – If the Solaris JumpStart entry is the only entry listed, you can alternatively wait for the selection screen to time out. If you do not respond in 30 seconds, the system automatically continues the boot sequence.

---

```
GNU GRUB version 0.95 (631K lower / 2095488K upper memory)
```

```
+-----+
| Solaris_10 Jumpstart |
| |
| |
+-----+
```

Use the ^ and v keys to select which entry is highlighted.  
Press enter to boot the selected OS, 'e' to edit the  
commands before booting, or 'c' for a command-line.

After further processing, the installation type menu is displayed.

**g. From the installation type menu, immediately type the menu number for Custom JumpStart.**

---

**Note** – If you do not type the number for Custom JumpStart before the 30-second timeout period ends, the system automatically begins the Solaris interactive installation.

---

Select the type of installation you want to perform:

- 1 Solaris Interactive
- 2 Custom JumpStart
- 3 Solaris Interactive Text (Desktop session)
- 4 Solaris Interactive Text (Console session)
- 5 Apply driver updates
- 6 Single user shell

Enter the number of your choice.

2

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

**h. When the BIOS screen again appears, immediately press Esc+2 or press the F2 key.**

---

**Note** – If you do not interrupt the BIOS at this point, it automatically returns to the installation type menu. There, if no choice is typed within 30 seconds, the system automatically begins an interaction installation.

---

After further processing, the BIOS Setup Utility is displayed.

**i. In the menu bar, navigate to the Boot menu.**

The list of boot devices is displayed.

**j. Navigate to the Hard Drive entry and move it back to the top of the boot order.**

**k. Save your change and exit the BIOS.**

The boot sequence begins again. No further interaction with the GRUB menu is needed to complete booting into cluster mode.

**19 For the Solaris 10 OS, verify on each node that multi-user services for the Service Management Facility (SMF) are online.**

If services are not yet online for a node, wait until the state becomes online before you proceed to the next step.

```
phys-schost# svcs multi-user-server
STATE STIME FMRI
online 17:52:55 svc:/milestone/multi-user-server:default
```

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

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

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

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

```
phys-schost-new# mkdir -p mountpoint
```

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

---

**Note** – The mount points become active after you reboot the cluster in [Step 24](#).

---

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

```
phys-schost# grep vxio /etc/name_to_major
vxio NNN
```

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

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

```
set kernel_cage_enable=1
```

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

**22 If you intend to use Sun Cluster HA for NFS on a highly available local file system, ensure that the loopback file system (LOFS) is disabled.**

To disable LOFS, add the following entry to the `/etc/system` file on each node of the cluster.

```
exclude:lofs
```

The change to the `/etc/system` file becomes effective after the next system reboot.

---

**Note** – You cannot have LOFS enabled if you use Sun Cluster HA for NFS on a highly available local file system *and* have `automountd` running. LOFS can cause switchover problems for Sun Cluster HA for NFS. If you choose to add Sun Cluster HA for NFS on a highly available local file system, you must make one of the following configuration changes.

However, if you configure non-global zones in your cluster, you must enable LOFS on all cluster nodes. If Sun Cluster HA for NFS on a highly available local file system must coexist with LOFS, use one of the other solutions instead of disabling LOFS.

- Disable LOFS.
- Disable the `automountd` daemon.
- Exclude from the automounter map all files that are part of the highly available local file system that is exported by Sun Cluster HA for NFS. This choice enables you to keep both LOFS and the `automountd` daemon enabled.

---

See “The Loopback File System” in *System Administration Guide: Devices and File Systems* (Solaris 9 or Solaris 10) for more information about loopback file systems.

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

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

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

**24 If you performed a task that requires a cluster reboot, follow these steps to reboot the cluster.**

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

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

**a. On one node, become superuser.**

**b. Shut down the cluster.**

```
phys-schost-1# cluster shutdown -y -g0 clustername
```

---

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

Cluster nodes remain in installation mode until the first time that you run the `clsetup` command. You run this command during the procedure “[How to Configure Quorum Devices](#)” on page 117.

---

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

- **On SPARC based systems, do the following:**

```
ok boot
```

- **On x86 based systems, do the following:**

When the GRUB menu is displayed, select the appropriate Solaris entry and press Enter. The GRUB menu appears similar to the following:

```
GNU GRUB version 0.95 (631K lower / 2095488K upper memory)
+-----+
| Solaris 10 /sol_10_x86 |
| Solaris failsafe |
| |
+-----+
```

Use the `^` and `v` keys to select which entry is highlighted.  
Press enter to boot the selected OS, `'e'` to edit the commands before booting, or `'c'` for a command-line.

For more information about GRUB based booting, see Chapter 11, “GRUB Based Booting (Tasks),” in *System Administration Guide: Basic Administration*.

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

**25 (Optional) If you did not perform [Step 24](#) to reboot the nodes, start the Sun Java Web Console web server manually on each node.**

```
phys-schost# smcwebserver start
```

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

**26 On one node, become superuser.**

**27 Verify that all nodes have joined the cluster.**

```
phys-schost# clnode status
```

Output resembles the following.

```
=== Cluster Nodes ===

--- Node Status ---

Node Name Status

phys-schost-1 Online
phys-schost-2 Online
phys-schost-3 Online
```

For more information, see the `clnode(1CL)` man page.

## 28 (Optional) On each node, enable automatic node reboot if all monitored disk paths fail.

### a. Enable the automatic reboot feature.

```
phys-schost# clnode set -p reboot_on_path_failure=enabled
```

`-p` Specifies the property to set

`reboot_on_path_failure=enable` Specifies that the node will reboot if all monitored disk paths fail, provided that at least one of the disks is accessible from a different node in the cluster.

### b. Verify that automatic reboot on disk-path failure is enabled.

```
phys-schost# clnode show
=== Cluster Nodes ===

Node Name: node
...
reboot_on_path_failure: enabled
...
```

**Next Steps** If you added a node to a two-node cluster, go to [“How to Update Quorum Devices After Adding a Node to a Cluster”](#) on page 114.

Otherwise, go to the next appropriate procedure:

- If you installed a multiple-node cluster and chose automatic quorum configuration, postinstallation setup is complete. Go to [“How to Verify the Quorum Configuration and Installation Mode”](#) on page 121.
- If you installed a multiple-node cluster and declined automatic quorum configuration, perform postinstallation setup. Go to [“How to Configure Quorum Devices”](#) on page 117.
- If you added a new node to an existing cluster that uses a quorum device, go to [“How to Update Quorum Devices After Adding a Node to a Cluster”](#) on page 114.

- If you added a new node to an existing cluster that does not use a quorum device, verify the state of the cluster. Go to [“How to Verify the Quorum Configuration and Installation Mode”](#) on page 121.
- If you installed a single-node cluster, cluster establishment is complete. Go to [“Creating Cluster File Systems”](#) on page 175 to install volume management software and configure the cluster.

**Troubleshooting** **Disabled scinstall option** - If the JumpStart option of the `scinstall` command does not have an asterisk in front, the option is disabled. This condition indicates that JumpStart setup is not complete or that the setup has an error. To correct this condition, first quit the `scinstall` utility. Repeat [Step 1](#) through [Step 14](#) to correct JumpStart setup, then restart the `scinstall` utility.

**Error messages about nonexistent nodes** - Unless you have installed your own `/etc/inet/ntp.conf` file, the `scinstall` command installs a default `ntp.conf` file for you. The default file is shipped with references to the maximum number of nodes. Therefore, the `xntpd(1M)` daemon might issue error messages regarding some of these references at boot time. You can safely ignore these messages. See [“How to Configure Network Time Protocol \(NTP\)”](#) on page 123 for information about how to suppress these messages under otherwise normal cluster conditions.

## ▼ How to Prepare the Cluster for Additional Cluster Nodes

Perform this procedure on existing cluster nodes to prepare the cluster for the addition of new cluster nodes.

**Before You Begin** Perform the following tasks:

- Ensure that all necessary hardware is installed.
  - Ensure that the host adapter is installed on the new node. See the *Sun Cluster 3.1 - 3.2 Hardware Administration Manual for Solaris OS*.
  - Verify that any existing cluster interconnects can support the new node. See the *Sun Cluster 3.1 - 3.2 Hardware Administration Manual for Solaris OS*.
  - Ensure that any additional storage is installed. See the appropriate manual from the Sun Cluster 3.1 - 3.2 Hardware Administration Collection.

- 1 **If you use the Cluster Control Panel (CCP), update the configuration files on the administrative console.**
  - a. **Add to the cluster’s entry in the `/etc/clusters` file the name of the node that you are adding.**
  - b. **Add to the `/etc/serialports` files an entry with the new node name, the host name of the node’s console-access device, and the port number.**

**2 Add the name of the new node to the cluster's authorized-nodes list.**

a. On any node, become superuser.

b. Start the `clsetup` utility.

```
phys-schost# clsetup
```

The Main Menu is displayed.

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

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

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

The `clsetup` utility displays the message `Command completed successfully` if the task is completed without error.

f. Quit the `clsetup` utility.

**3 If you are adding a node to a single-node cluster, ensure that two cluster interconnects already exist by displaying the interconnect configuration.**

```
phys-schost# clinterconnect show
```

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

- If the output shows configuration information for two cables or for two adapters, proceed to [Step 4](#).
- If the output shows no configuration information for either cables or adapters, or shows configuration information for only one cable or adapter, configure new cluster interconnects.
  - a. On one node, start the `clsetup` utility.

```
phys-schost# clsetup
```
  - b. Choose the menu item, **Cluster interconnect**.
  - c. Choose the menu item, **Add a transport cable**.

Follow the instructions to specify the name of the node to add to the cluster, the name of a transport adapter, and whether to use a transport switch.
  - d. If necessary, repeat [Step c](#) to configure a second cluster interconnect.
  - e. When finished, quit the `clsetup` utility.



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

```
phys-schost# clinterconnect show
```

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

**4 Ensure that the private-network configuration can support the nodes and private networks that you are adding.**

**a. Display the maximum number of nodes and private networks that the current private-network configuration supports.**

```
phys-schost# cluster show-netprops
```

Output looks similar to the following, which shows the default values:

```
=== Private Network ===
```

```
private_netaddr: 172.16.0.0
private_netmask: 255.255.248.0
max_nodes: 64
max_privatenets: 10
```

**b. Determine whether the current private-network configuration can support the increased number of nodes, including non-global zones, and private networks.**

▪ **If the current IP address range is sufficient, you are ready to install the new node.**

Go to [“How to Configure Sun Cluster Software on Additional Cluster Nodes \(scinstall\)”](#) on page 103.

▪ **If the current IP address range is not sufficient, reconfigure the private IP address range.**

Go to [“How to Change the Private Network Configuration When Adding Nodes or Private Networks”](#) on page 98. You must shut down the cluster to change the private IP address range. This involves switching each resource group offline, disabling all resources in the cluster, then rebooting into noncluster mode before you reconfigure the IP address range.

**Next Steps** Configure Sun Cluster software on the new cluster nodes. Go to [“How to Configure Sun Cluster Software on Additional Cluster Nodes \(scinstall\)”](#) on page 103 or [“How to Configure Sun Cluster Software on Additional Cluster Nodes \(XML\)”](#) on page 110.

## ▼ How to Change the Private Network Configuration When Adding Nodes or Private Networks

Perform this task to change the cluster private IP address range to accommodate an increase in the number of nodes or non-global zones or in the number of private networks, or a combination. You can also use this procedure to decrease the private IP address range.

---

**Note** – This procedure requires you to shut down the entire cluster.

---

- 1 Become superuser on a node of the cluster.**
- 2 From one node, start the `clsetup` utility.**
  - # `clsetup`**
  - The `clsetup` Main Menu is displayed.
- 3 Switch each resource group offline.**

If the node contains non-global zones, any resource groups in the zones are also switched offline.

  - a. Type the number that corresponds to the option for Resource groups and press the Return key.**

The Resource Group Menu is displayed.
  - b. Type the number that corresponds to the option for Online/Offline or Switchover a resource group and press the Return key.**
  - c. Follow the prompts to take offline all resource groups and to put them in the unmanaged state.**
  - d. When all resource groups are offline, type `q` to return to the Resource Group Menu.**
- 4 Disable all resources in the cluster.**
  - a. Type the number that corresponds to the option for Enable/Disable a resource and press the Return key.**
  - b. Choose a resource to disable and follow the prompts.**
  - c. Repeat the previous step for each resource to disable.**
  - d. When all resources are disabled, type `q` to return to the Resource Group Menu.**
- 5 Quit the `clsetup` utility.**

**6 Verify that all resources on all nodes are Offline and that all resource groups are in the Unmanaged state.**

```
cluster status -t resource,resourcegroup
```

```
-t Limits output to the specified cluster object
```

```
resource Specifies resources
```

```
resourcegroup Specifies resource groups
```

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

```
cluster shutdown -g0 -y
```

```
-g Specifies the wait time in seconds
```

```
-y Prevents the prompt that asks you to confirm a shutdown from being issued
```

**8 Boot each node into noncluster mode.**

- On SPARC based systems, perform the following command:

```
ok boot -x
```

- On x86 based systems, perform the following commands:

- a. In the GRUB menu, use the arrow keys to select the appropriate Solaris entry and type **e** to edit its commands.

The GRUB menu appears similar to the following:

```
GNU GRUB version 0.95 (631K lower / 2095488K upper memory)
```

```
+-----+
| Solaris 10 /sol_10_x86 |
| Solaris failsafe |
| |
+-----+

```

Use the **^** and **v** keys to select which entry is highlighted.

Press enter to boot the selected OS, 'e' to edit the commands before booting, or 'c' for a command-line.

For more information about GRUB based booting, see Chapter 11, “GRUB Based Booting (Tasks),” in *System Administration Guide: Basic Administration*.

- b. In the boot parameters screen, use the arrow keys to select the kernel entry and type **e** to edit the entry.

The GRUB boot parameters screen appears similar to the following:

```
GNU GRUB version 0.95 (615K lower / 2095552K upper memory)
```

```
+-----+
| root (hd0,0,a) |
| kernel /platform/i86pc/multiboot |

```

```
| module /platform/i86pc/boot_archive |
+-----+
Use the ^ and v keys to select which entry is highlighted.
Press 'b' to boot, 'e' to edit the selected command in the
boot sequence, 'c' for a command-line, 'o' to open a new line
after ('O' for before) the selected line, 'd' to remove the
selected line, or escape to go back to the main menu.
```

**c. Add -x to the command to specify that the system boot into noncluster mode.**

[ Minimal BASH-like line editing is supported. For the first word, TAB lists possible command completions. Anywhere else TAB lists the possible completions of a device/filename. ESC at any time exits. ]

```
grub edit> kernel /platform/i86pc/multiboot -x
```

**d. Press Enter to accept the change and return to the boot parameters screen.**

The screen displays the edited command.

```
GNU GRUB version 0.95 (615K lower / 2095552K upper memory)
+-----+
| root (hd0,0,a) |
| kernel /platform/i86pc/multiboot -x |
| module /platform/i86pc/boot_archive |
+-----+
Use the ^ and v keys to select which entry is highlighted.
Press 'b' to boot, 'e' to edit the selected command in the
boot sequence, 'c' for a command-line, 'o' to open a new line
after ('O' for before) the selected line, 'd' to remove the
selected line, or escape to go back to the main menu.-
```

**e. Type b to boot the node into noncluster mode.**

---

**Note** – This change to the kernel boot parameter command does not persist over the system boot. The next time you reboot the node, it will boot into cluster mode. To boot into noncluster mode instead, perform these steps to again to add the -x option to the kernel boot parameter command.

---

**9 From one node, start the clsetup utility.**

When run in noncluster mode, the clsetup utility displays the Main Menu for noncluster-mode operations.

**10 Type the number that corresponds to the option for Change IP Address Range and press the Return key.**

The clsetup utility displays the current private-network configuration, then asks if you would like to change this configuration.

**11 To change either the private-network IP address or the IP address range, type `yes` and press the Return key.**

The `clsetup` utility displays the default private-network IP address, `172.16.0.0`, and asks if it is okay to accept this default.

**12 Change or accept the private-network IP address.**

- **To accept the default private-network IP address and proceed to changing the IP address range, type `yes` and press the Return key.**

The `clsetup` utility will ask if it is okay to accept the default netmask. Skip to the next step to enter your response.

- **To change the default private-network IP address, perform the following substeps.**

- a. **Type `no` in response to the `clsetup` utility question about whether it is okay to accept the default address, then press the Return key.**

The `clsetup` utility will prompt for the new private-network IP address.

- b. **Type the new IP address and press the Return key.**

The `clsetup` utility displays the default netmask and then asks if it is okay to accept the default netmask.

**13 Change or accept the default private-network IP address range.**

The default netmask is `255.255.248.0`. This default IP address range supports up to 64 nodes and up to 10 private networks in the cluster.

- **To accept the default IP address range, type `yes` and press the Return key.**

Then skip to the next step.

- **To change the IP address range, perform the following substeps.**

- a. **Type `no` in response to the `clsetup` utility's question about whether it is okay to accept the default address range, then press the Return key.**

When you decline the default netmask, the `clsetup` utility prompts you for the number of nodes and private networks that you expect to configure in the cluster.

- b. **Enter the number of nodes and private networks that you expect to configure in the cluster.**

From these numbers, the `clsetup` utility calculates two proposed netmasks:

- The first netmask is the minimum netmask to support the number of nodes and private networks that you specified.
- The second netmask supports twice the number of nodes and private networks that you specified, to accommodate possible future growth.

- c. Specify either of the calculated netmasks, or specify a different netmask that supports the expected number of nodes and private networks.

**14** Type `yes` in response to the `clsetup` utility's question about proceeding with the update.

**15** When finished, exit the `clsetup` utility.

**16** Reboot each node back into the cluster.

- a. Shut down each node.

```
shutdown -g0 -y
```

- b. Boot each node into cluster mode.

- On SPARC based systems, do the following:

```
ok boot
```

- On x86 based systems, do the following:

When the GRUB menu is displayed, select the appropriate Solaris entry and press Enter. The GRUB menu appears similar to the following:

```
GNU GRUB version 0.95 (631K lower / 2095488K upper memory)
```

```
+-----+
| Solaris 10 /sol_10_x86 |
| Solaris failsafe |
| |
+-----+
```

Use the `^` and `v` keys to select which entry is highlighted.

Press enter to boot the selected OS, 'e' to edit the commands before booting, or 'c' for a command-line.

For more information about GRUB based booting, see Chapter 11, "GRUB Based Booting (Tasks)," in *System Administration Guide: Basic Administration*.

**17** From one node, start the `clsetup` utility.

```
clsetup
```

The `clsetup` Main Menu is displayed.

**18** Re-enable all disabled resources.

- a. Type the number that corresponds to the option for Resource groups and press the Return key.

The Resource Group Menu is displayed.

- b. Type the number that corresponds to the option for Enable/Disable a resource and press the Return key.

- c. Choose a resource to enable and follow the prompts.
- d. Repeat for each disabled resource.
- e. When all resources are re-enabled, type **q** to return to the Resource Group Menu.

**19 Bring each resource group back online.**

If the node contains non-global zones, also bring online any resource groups that are in those zones.

- a. Type the number that corresponds to the option for Online/Offline or Switchover a resource group and press the Return key.
- b. Follow the prompts to put each resource group into the managed state and then bring the resource group online.

**20 When all resource groups are back online, exit the `clsetup` utility.**

Type **q** to back out of each submenu, or press Ctrl-C.

**Next Steps** To add a node to an existing cluster, go to one of the following procedures:

- “How to Configure Sun Cluster Software on Additional Cluster Nodes (`scinstall`)” on page 103
- “How to Install Solaris and Sun Cluster Software (JumpStart)” on page 79
- “How to Configure Sun Cluster Software on Additional Cluster Nodes (XML)” on page 110

To create a non-global zone on a cluster node, go to “Configuring a Non-Global Zone on a Cluster Node” on page 179.

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

Perform this procedure to add a new node to an existing cluster. To use JumpStart to add a new node, instead follow procedures in “How to Install Solaris and Sun Cluster Software (JumpStart)” on page 79.

---

**Note** – This procedure uses the interactive form of the `scinstall` command. To use the noninteractive forms of the `scinstall` command, such as when developing installation scripts, see the `scinstall(1M)` man page.

Ensure that Sun Cluster software packages are installed on the node, either manually or by using the silent-mode form of the Java ES installer program, before you run the `scinstall` command. For information about running the Java ES installer program from an installation script, see Chapter 5, “Installing in Silent Mode,” in *Sun Java Enterprise System 5 Installation Guide for UNIX*.

---

**Before You Begin** Perform the following tasks:

- Ensure that the Solaris OS is installed to support Sun Cluster software.  
If Solaris software is already installed on the node, you must ensure that the Solaris installation meets the requirements for Sun Cluster software and any other software that you intend to install on the cluster. See [“How to Install Solaris Software” on page 50](#) for more information about installing Solaris software to meet Sun Cluster software requirements.
- Ensure that Sun Cluster software packages and patches are installed on the node. See [“How to Install Sun Cluster Framework and Data-Service Software Packages” on page 58](#).
- Ensure that the cluster is prepared for the addition of the new node. See [“How to Prepare the Cluster for Additional Cluster Nodes” on page 95](#).
- Determine which mode of the `scinstall` utility you will use, Typical or Custom. For the Typical installation of Sun Cluster software, `scinstall` automatically specifies the following configuration defaults.

| Component                       | Default Value       |
|---------------------------------|---------------------|
| Cluster-transport switches      | switch1 and switch2 |
| Global-devices file-system name | /globaldevices      |

- Complete one of the following configuration planning worksheets. See [“Planning the Solaris OS” on page 14](#) and [“Planning the Sun Cluster Environment” on page 20](#) for planning guidelines.
  - **Typical Mode Worksheet** - If you will use Typical mode and accept all defaults, complete the following worksheet.

| Component                          | Description/Example                                                                                                         | Answer       |               |
|------------------------------------|-----------------------------------------------------------------------------------------------------------------------------|--------------|---------------|
| Sponsoring Node                    | What is the name of the sponsoring node?<br><i>Choose any node that is active in the cluster.</i>                           |              |               |
| Cluster Name                       | What is the name of the cluster that you want the node to join?                                                             |              |               |
| Check                              | Do you want to run the <code>sccheck</code> validation utility?                                                             | Yes   No     |               |
| Autodiscovery of Cluster Transport | Do you want to use autodiscovery to configure the cluster transport?<br>If no, supply the following additional information: | Yes   No     |               |
| Point-to-Point Cables              | Does the node that you are adding to the cluster make this a two-node cluster?                                              | Yes   No     |               |
|                                    | Does the cluster use switches?                                                                                              | Yes   No     |               |
| Cluster Switches                   | If used, what are the names of the two switches?<br>Defaults: switch1 and switch2                                           | <i>First</i> | <i>Second</i> |



| Component                             | Description/Example                                                                                                                                    | Answer       |               |
|---------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------|--------------|---------------|
| Cluster Transport Adapters and Cables | What are the names of the two transport adapters?                                                                                                      | <i>First</i> | <i>Second</i> |
|                                       | Where does each transport adapter connect to ( <i>a switch or another adapter</i> )?<br>Switch defaults: <code>switch1</code> and <code>switch2</code> |              |               |
|                                       | For transport switches, do you want to use the default port name?                                                                                      | Yes   No     | Yes   No      |
|                                       | If no, what is the name of the port that you want to use?                                                                                              |              |               |
| Automatic Reboot                      | Do you want <code>scinstall</code> to automatically reboot the node after installation?                                                                | Yes   No     |               |

- **Custom Mode Worksheet** - If you will use Custom mode and customize the configuration data, complete the following worksheet.

| Component                             | Description/Example                                                                                                                                    | Answer       |               |
|---------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------|--------------|---------------|
| Sponsoring Node                       | What is the name of the sponsoring node?<br><i>Choose any node that is active in the cluster.</i>                                                      |              |               |
| Cluster Name                          | What is the name of the cluster that you want the node to join?                                                                                        |              |               |
| Check                                 | Do you want to run the <code>sccheck</code> validation utility?                                                                                        | Yes   No     |               |
| Autodiscovery of Cluster Transport    | Do you want to use autodiscovery to configure the cluster transport?<br>If no, supply the following additional information:                            | Yes   No     |               |
| Point-to-Point Cables                 | Does the node that you are adding to the cluster make this a two-node cluster?                                                                         | Yes   No     |               |
|                                       | Does the cluster use switches?                                                                                                                         | Yes   No     |               |
| Cluster Switches                      | If used, what are the names of the two switches?<br>Defaults: <code>switch1</code> and <code>switch2</code>                                            | <i>First</i> | <i>Second</i> |
| Cluster Transport Adapters and Cables | What are the names of the two transport adapters?                                                                                                      | <i>First</i> | <i>Second</i> |
|                                       | Where does each transport adapter connect to ( <i>a switch or another adapter</i> )?<br>Switch defaults: <code>switch1</code> and <code>switch2</code> |              |               |
|                                       | For transport switches, do you want to use the default port name?                                                                                      | Yes   No     | Yes   No      |
|                                       | If no, what is the name of the port that you want to use?                                                                                              |              |               |
| Global Devices File System            | What is the name of the global-devices file system?<br>Default: <code>/globaldevices</code>                                                            |              |               |
| Automatic Reboot                      | Do you want <code>scinstall</code> to automatically reboot the node after installation?                                                                | Yes   No     |               |

Follow these guidelines to use the interactive `scinstall` utility in this procedure:

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

**1 On the cluster node to configure, become superuser.**

**2 Start the `scinstall` utility.**

```
phys-schost-new# /usr/cluster/bin/scinstall
```

The `scinstall` Main Menu is displayed.

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

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

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

- \* 1) Create a new cluster or add a cluster node
- 2) Configure a cluster to be JumpStarted from this install server
- 3) Manage a dual-partition upgrade
- 4) Upgrade this cluster node
- \* 5) Print release information for this cluster node
  
- \* ?) Help with menu options
- \* q) Quit

```
Option: 1
```

The New Cluster and Cluster Node Menu is displayed.

**4 Type the number that corresponds to the option for Add this machine as a node in an existing cluster and press the Return key.**

**5 Follow the menu prompts to supply your answers from the configuration planning worksheet.**

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

**6 Unload the Sun Java Availability Suite DVD-ROM from the DVD-ROM drive.**

- a. To ensure that the DVD-ROM is not being used, change to a directory that does *not* reside on the DVD-ROM.

**b. Eject the DVD-ROM.**

```
phys-schost# eject cdrom
```

**7 Repeat this procedure on any other node to add to the cluster until all additional nodes are fully configured.****8 For the Solaris 10 OS, verify on each node that multi-user services for the Service Management Facility (SMF) are online.**

If services are not yet online for a node, wait until the state becomes online before you proceed to the next step.

```
phys-schost# svcs multi-user-server
STATE STIME FMRI
online 17:52:55 svc:/milestone/multi-user-server:default
```

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

```
phys-schost# claccess deny-all
```

Alternately, you can use the `clsetup` utility. See “How to Add a Node to the Authorized Node List” in *Sun Cluster System Administration Guide for Solaris OS* for procedures.

**10 On one node, become superuser.****11 Verify that all nodes have joined the cluster.**

```
phys-schost# clnode status
```

Output resembles the following.

```
=== Cluster Nodes ===

--- Node Status ---

Node Name Status

phys-schost-1 Online
phys-schost-2 Online
phys-schost-3 Online
```

For more information, see the `clnode(1CL)` man page.

**12 Verify that all necessary patches are installed.**

```
phys-schost# showrev -p
```

**13 (Optional) Enable automatic node reboot if all monitored disk paths fail.****a. Enable the automatic reboot feature.**

```
phys-schost# clnode set -p reboot_on_path_failure=enabled
```

|                               |                                                                                                                                                                   |
|-------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| -p                            | Specifies the property to set                                                                                                                                     |
| reboot_on_path_failure=enable | Specifies that the node will reboot if all monitored disk paths fail, provided that at least one of the disks is accessible from a different node in the cluster. |

**b. Verify that automatic reboot on disk-path failure is enabled.**

```
phys-schost# clnode show
=== Cluster Nodes ===

Node Name: node
...
reboot_on_path_failure: enabled
...
```

**14 If you intend to use Sun Cluster HA for NFS on a highly available local file system, ensure that the loopback file system (LOFS) is disabled.**

To disable LOFS, add the following entry to the `/etc/system` file on each node of the cluster.

```
exclude:lofs
```

The change to the `/etc/system` file becomes effective after the next system reboot.

---

**Note** – You cannot have LOFS enabled if you use Sun Cluster HA for NFS on a highly available local file system *and* have `automountd` running. LOFS can cause switchover problems for Sun Cluster HA for NFS. If you choose to add Sun Cluster HA for NFS on a highly available local file system, you must make one of the following configuration changes.

However, if you configure non-global zones in your cluster, you must enable LOFS on all cluster nodes. If Sun Cluster HA for NFS on a highly available local file system must coexist with LOFS, use one of the other solutions instead of disabling LOFS.

- Disable LOFS.
- Disable the `automountd` daemon.
- Exclude from the automounter map all files that are part of the highly available local file system that is exported by Sun Cluster HA for NFS. This choice enables you to keep both LOFS and the `automountd` daemon enabled.

---

See “The Loopback File System” in *System Administration Guide: Devices and File Systems* (Solaris 9 or Solaris 10) for more information about loopback file systems.

### Example 3–3 Configuring Sun Cluster Software on an Additional Node

The following example shows the node `phys-schost-3` added to the cluster `schost`. The sponsoring node is `phys-schost-1`.

\*\*\* Adding a Node to an Existing Cluster \*\*\*

Fri Feb 4 10:17:53 PST 2005

```
scinstall -ik -C schost -N phys-schost-1 -A trtype=dlpi,name=qfe2 -A trtype=dlpi,name=qfe3
-m endpoint=:qfe2,endpoint=switch1 -m endpoint=:qfe3,endpoint=switch2
```

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

Adding node "phys-schost-3" to the cluster configuration ... done

Adding adapter "qfe2" to the cluster configuration ... done

Adding adapter "qfe3" 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

Copying the postconfig file from "phys-schost-1" if it exists ... done

Copying the Common Agent Container keys from "phys-schost-1" ... done

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

Setting the major number for the "did" driver ...

Obtaining the major number for the "did" driver from "phys-schost-1" ... done

"did" driver major number set to 300

Checking for global devices global file system ... done

Updating vfstab ... done

Verifying that NTP is configured ... done

Initializing NTP configuration ... done

Updating nsswitch.conf ...

done

Adding clusternode entries to /etc/inet/hosts ... done

Configuring IP Multipathing groups in "/etc/hostname.<adapter>" files

Updating "/etc/hostname.hme0".

Verifying that power management is NOT configured ... done

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

The "local-mac-address?" parameter setting has been changed to "true".

Ensure network routing is disabled ... done

Updating file ("ntp.conf.cluster") on node phys-schost-1 ... done

Updating file ("hosts") on node phys-schost-1 ... done

Rebooting ...

**Troubleshooting** **Unsuccessful configuration** - If one or more nodes cannot join the cluster, or if the wrong configuration information was specified, first attempt to rerun this procedure. If that does not correct the problem, perform the procedure [“How to Unconfigure Sun Cluster Software to Correct Installation Problems” on page 263](#) on each misconfigured node to remove it from the cluster configuration. You do not need to uninstall the Sun Cluster software packages. Then rerun this procedure.

**Next Steps** If you added a node to an existing cluster that uses a quorum device, go to [“How to Update Quorum Devices After Adding a Node to a Cluster” on page 114](#).

Otherwise, go to [“How to Verify the Quorum Configuration and Installation Mode” on page 121](#).

## ▼ **How to Configure Sun Cluster Software on Additional Cluster Nodes (XML)**

Perform this procedure to configure a new cluster node by using an XML cluster configuration file. The new node can be a duplication of an existing cluster node that runs Sun Cluster 3.2 software.

This procedure configures the following cluster components on the new node:

- Cluster node membership
- Cluster interconnect
- Global devices

**Before You Begin** Perform the following tasks:

- Ensure that the Solaris OS is installed to support Sun Cluster software.  
If Solaris software is already installed on the node, you must ensure that the Solaris installation meets the requirements for Sun Cluster software and any other software that you intend to install on the cluster. See [“How to Install Solaris Software” on page 50](#) for more information about installing Solaris software to meet Sun Cluster software requirements.
- Ensure that Sun Cluster software packages and any necessary patches are installed on the node. See [“How to Install Sun Cluster Framework and Data-Service Software Packages” on page 58](#).

- Ensure that the cluster is prepared for the addition of the new node. See “How to Prepare the Cluster for Additional Cluster Nodes” on page 95.

**1 Ensure that Sun Cluster 3.2 software is not yet configured on the potential node that you want to add to a cluster.**

**a. Become superuser on the potential node.**

**b. Determine whether Sun Cluster 3.2 software is configured on the potential node.**

```
phys-schost-new# /usr/sbin/clinfo -n
```

- **If the command fails, go to Step 2.**

Sun Cluster 3.2 software is not yet configured on the node. You can add the potential node to the cluster.

- **If the command returns a node ID number, proceed to Step c.**

Sun Cluster software is already configured on the node. Before you can add the node to a different cluster, you must remove the existing cluster configuration information.

**c. Boot the potential node into noncluster mode.**

- **On SPARC based systems, perform the following command:**

```
ok boot -x
```

- **On x86 based systems, perform the following commands:**

- i. **In the GRUB menu, use the arrow keys to select the appropriate Solaris entry and type e to edit its commands.**

The GRUB menu appears similar to the following:

```
GNU GRUB version 0.95 (631K lower / 2095488K upper memory)
```

```
+-----+
| Solaris 10 /sol_10_x86 |
| Solaris failsafe |
| |
+-----+
```

Use the ^ and v keys to select which entry is highlighted.

Press enter to boot the selected OS, 'e' to edit the commands before booting, or 'c' for a command-line.

For more information about GRUB based booting, see Chapter 11, “GRUB Based Booting (Tasks),” in *System Administration Guide: Basic Administration*.

- ii. **In the boot parameters screen, use the arrow keys to select the kernel entry and type e to edit the entry.**

The GRUB boot parameters screen appears similar to the following:

```
GNU GRUB version 0.95 (615K lower / 2095552K upper memory)
+-----+
| root (hd0,0,a) |
| kernel /platform/i86pc/multiboot |
| module /platform/i86pc/boot_archive |
+-----+
```

Use the ^ and v keys to select which entry is highlighted. Press 'b' to boot, 'e' to edit the selected command in the boot sequence, 'c' for a command-line, 'o' to open a new line after ('O' for before) the selected line, 'd' to remove the selected line, or escape to go back to the main menu.

- iii. **Add -x to the command to specify that the system boot into noncluster mode.**

[ Minimal BASH-like line editing is supported. For the first word, TAB lists possible command completions. Anywhere else TAB lists the possible completions of a device/filename. ESC at any time exits. ]

```
grub edit> kernel /platform/i86pc/multiboot -x
```

- iv. **Press Enter to accept the change and return to the boot parameters screen.**

The screen displays the edited command.

```
GNU GRUB version 0.95 (615K lower / 2095552K upper memory)
+-----+
| root (hd0,0,a) |
| kernel /platform/i86pc/multiboot -x |
| module /platform/i86pc/boot_archive |
+-----+
```

Use the ^ and v keys to select which entry is highlighted. Press 'b' to boot, 'e' to edit the selected command in the boot sequence, 'c' for a command-line, 'o' to open a new line after ('O' for before) the selected line, 'd' to remove the selected line, or escape to go back to the main menu.-

- v. **Type b to boot the node into noncluster mode.**

---

**Note** – This change to the kernel boot parameter command does not persist over the system boot. The next time you reboot the node, it will boot into cluster mode. To boot into noncluster mode instead, perform these steps to again to add the -x option to the kernel boot parameter command.

---



**d. Unconfigure Sun Cluster software from the potential node.**

```
phys-schost-new# /usr/cluster/bin/clnode remove
```

**2 If you are duplicating a node that runs Sun Cluster 3.2 software, create a cluster configuration XML file.****a. Become superuser on the cluster node that you want to duplicate.****b. Export the existing node's configuration information to a file.**

```
phys-schost# clnode export -o clconfigfile
```

`-o` Specifies the output destination.

`clconfigfile` The name of the cluster configuration XML file. The specified file name can be an existing file or a new file that the command will create.

For more information, see the `clnode(1CL)` man page.

**c. Copy the cluster configuration XML file to the potential node that you will configure as a new cluster node.****3 Become superuser on the potential node.****4 Modify the cluster configuration XML file as needed.****a. Open your cluster configuration XML file for editing.**

- **If you are duplicating an existing cluster node, open the file that you created with the `clnode export` command.**

- **If you are not duplicating an existing cluster node, create a new file.**

Base the file on the element hierarchy that is shown in the `clconfiguration(5CL)` man page. You can store the file in any directory.

**b. Modify the values of the XML elements to reflect the node configuration that you want to create.**

See the `clconfiguration(5CL)` man page for details about the structure and content of the cluster configuration XML file.

**5 Validate the cluster configuration XML file.**

```
phys-schost-new# xmllint --valid --noout clconfigfile
```

**6 Configure the new cluster node.**

```
phys-schost-new# clnode add -n sponsornode -i clconfigfile
```

- n *sponsornode* Specifies the name of an existing cluster member to act as the sponsor for the new node.
- i *clconfigfile* Specifies the name of the cluster configuration XML file to use as the input source.

## 7 (Optional) Enable automatic node reboot if all monitored disk paths fail.

### a. Enable the automatic reboot feature.

```
phys-schost# clnode set -p reboot_on_path_failure=enabled
```

-p Specifies the property to set

reboot\_on\_path\_failure=enable Specifies that the node will reboot if all monitored disk paths fail, provided that at least one of the disks is accessible from a different node in the cluster.

### b. Verify that automatic reboot on disk-path failure is enabled.

```
phys-schost# clnode show
=== Cluster Nodes ===
```

```
Node Name: node
...
 reboot_on_path_failure: enabled
...
```

**Troubleshooting** **Unsuccessful configuration** - If one or more nodes cannot join the cluster, or if the wrong configuration information was specified, first attempt to rerun this procedure. If that does not correct the problem, perform the procedure [“How to Unconfigure Sun Cluster Software to Correct Installation Problems” on page 263](#) on each misconfigured node to remove it from the cluster configuration. You do not need to uninstall the Sun Cluster software packages. Then rerun this procedure.

**Next Steps** If you added a node to a cluster that uses a quorum device, go to [“How to Update Quorum Devices After Adding a Node to a Cluster” on page 114](#).

Otherwise, go to [“How to Verify the Quorum Configuration and Installation Mode” on page 121](#).

## ▼ How to Update Quorum Devices After Adding a Node to a Cluster

If you added a node to a cluster, you must update the configuration information of the quorum devices, regardless of whether you use SCSI devices, NAS devices, a quorum server, or a combination. To do this, you remove all quorum devices and update the global-devices namespace.

You can optionally reconfigure any quorum devices that you still want to use. This registers the new node with each quorum device, which can then recalculate its vote count based on the new number of nodes in the cluster.

Any newly configured SCSI quorum devices will be set to SCSI-3 reservations.

**Before You Begin** Ensure that you have completed installation of Sun Cluster software on the added node.

**1 On any node of the cluster, become superuser.**

**2 View the current quorum configuration.**

Command output lists each quorum device and each node. The following example output shows the current SCSI quorum device, d3.

```
phys-schost# clquorum list
d3
...
```

**3 Note the name of each quorum device that is listed.**

**4 Remove the original quorum device.**

Perform this step for each quorum device that is configured.

```
phys-schost# clquorum remove devicename
devicename Specifies the name of the quorum device.
```

**5 Verify that all original quorum devices are removed.**

If the removal of the quorum devices was successful, no quorum devices are listed.

```
phys-schost# clquorum status
```

**6 Update the global-devices namespace.**

```
phys-schost# cldevice populate
```

---

**Note** – This step is necessary to prevent possible node panic.

---

**7 On each node, verify that the cldevice populate command has completed processing before you attempt to add a quorum device.**

The cldevice populate command executes remotely on all nodes, even through the command is issued from just one node. To determine whether the cldevice populate command has completed processing, run the following command on each node of the cluster.

```
phys-schost# ps -ef | grep scgdevs
```

**8 (Optional) Add a quorum device.**

You can configure the same device that was originally configured as the quorum device or choose a new shared device to configure.

**a. (Optional) If you want to choose a new shared device to configure as a quorum device, display all devices that the system checks.**

Otherwise, skip to [Step c](#).

```
phys-schost# cldevice list -v
```

Output resembles the following:

| DID Device | Full Device Path                |
|------------|---------------------------------|
| -----      | -----                           |
| d1         | phys-schost-1:/dev/rdisk/c0t0d0 |
| d2         | phys-schost-1:/dev/rdisk/c0t6d0 |
| d3         | phys-schost-2:/dev/rdisk/c1t1d0 |
| d3         | phys-schost-1:/dev/rdisk/c1t1d0 |
| ...        |                                 |

**b. From the output, choose a shared device to configure as a quorum device.****c. Configure the shared device as a quorum device.**

```
phys-schost# clquorum add -t type devicename
```

`-t type` Specifies the type of quorum device. If this option is not specified, the default type `scsi` is used.

**d. Repeat for each quorum device that you want to configure.****e. Verify the new quorum configuration.**

```
phys-schost# clquorum list
```

Output should list each quorum device and each node.

**Example 3-4 Updating SCSI Quorum Devices After Adding a Node to a Two-Node Cluster**

The following example identifies the original SCSI quorum device `d2`, removes that quorum device, lists the available shared devices, updates the global-device namespace, configures `d3` as a new SCSI quorum device, and verifies the new device.

```
phys-schost# clquorum list
d2
phys-schost-1
phys-schost-2

phys-schost# clquorum remove d2
```

```

phys-schost# clquorum status
...
--- Quorum Votes by Device ---

Device Name Present Possible Status

phys-schost# cldevice list -v
DID Device Full Device Path

...
d3 phys-schost-2:/dev/rdisk/c1t1d0
d3 phys-schost-1:/dev/rdisk/c1t1d0
...
phys-schost# cldevice populate
phys-schost# ps -ef | grep scgdevs
phys-schost# clquorum add d3
phys-schost# clquorum list
d3
phys-schost-1
phys-schost-2

```

**Next Steps** Go to [“How to Verify the Quorum Configuration and Installation Mode”](#) on page 121.

## ▼ How to Configure Quorum Devices

**Note** – You do not need to configure quorum devices in the following circumstances:

- You chose automatic quorum configuration during Sun Cluster software configuration.
- You installed a single-node cluster.
- You added a node to an existing cluster and already have sufficient quorum votes assigned.

Instead, proceed to [“How to Verify the Quorum Configuration and Installation Mode”](#) on page 121.

Perform this procedure one time only, after the cluster is fully formed. Use this procedure to assign quorum votes and then to remove the cluster from installation mode.

**Before You Begin** Perform the following preparations to configure a quorum server or a NAS device as a quorum device.

- To configure a quorum server as a quorum device, do the following:
  - Install the Sun Cluster Quorum Server software on the quorum server host machine and start the quorum server. For information about installing and starting the quorum server, see *Sun Cluster Quorum Server User’s Guide*.

- Ensure that the spanning tree algorithm is disabled on any router or switch between the quorum server and the cluster nodes. An enabled spanning tree algorithm might cause the router or switch to not permit communications for extended periods of time. This prevention of communication to the quorum server is eventually interpreted by the cluster as loss of the quorum device.
- Have available the following information:
  - A name to assign to the configured quorum device
  - The IP address of the quorum server host machine
  - The port number of the quorum server
- To configure a Network Appliance network-attached storage (NAS) device as a quorum device, do the following:
  - Install the NAS device hardware and software. See Chapter 1, “Installing and Maintaining Network Appliance Network-Attached Storage Devices in a Sun Cluster Environment,” in *Sun Cluster 3.1 - 3.2 With Network-Attached Storage Devices Manual for Solaris OS* and your device documentation for requirements and installation procedures for NAS hardware and software.
  - Have available the following information:
    - The name of the NAS device
    - The LUN ID of the NAS device

See the following Network Appliance NAS documentation for information about creating and setting up a Network Appliance NAS device and LUN. You can access the following documents at <http://now.netapp.com>.

| Task                                                      | Network Appliance Documentation                                                                                                                                            |
|-----------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Setting up a NAS device                                   | <i>System Administration File Access Management Guide</i>                                                                                                                  |
| Setting up a LUN                                          | <i>Host Cluster Tool for Unix Installation Guide</i>                                                                                                                       |
| Installing ONTAP software                                 | <i>Software Setup Guide, Upgrade Guide</i>                                                                                                                                 |
| Exporting volumes for the cluster                         | <i>Data ONTAP Storage Management Guide</i>                                                                                                                                 |
| Installing NAS support software packages on cluster nodes | Log in to <a href="http://now.netapp.com">http://now.netapp.com</a> . From the Software Download page, download the <i>Host Cluster Tool for Unix Installation Guide</i> . |

## 1 On one node, become superuser.

**2 To use a shared SCSI disk as a quorum device, verify device connectivity to the cluster nodes and choose the device to configure.**

**a. From one node of the cluster, 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.

```
phys-schost-1# cldevice list -v
```

Output resembles the following:

| DID Device | Full Device Path             |
|------------|------------------------------|
| -----      | -----                        |
| d1         | pcircinus1:/dev/rdisk/c0t0d0 |
| d2         | pcircinus1:/dev/rdisk/c0t6d0 |
| d3         | pcircinus2:/dev/rdisk/c1t1d0 |
| d3         | pcircinus1:/dev/rdisk/c1t1d0 |
| ...        |                              |

**b. Ensure that the output shows all connections between cluster nodes and storage devices.**

**c. Determine the global device-ID name of each shared disk that you are configuring as a quorum device.**

---

**Note** – Any shared disk that you choose must be qualified for use as a quorum device. See “[Quorum Devices](#)” on page 29 for further information about choosing quorum devices.

---

Use the `scsidadm` output from [Step a](#) to identify the device-ID name of each shared disk that you are configuring as a quorum device. For example, the output in [Step a](#) shows that global device `d2` is shared by `phys-schost-1` and `phys-schost-2`.

**3 Start the `clsetup` utility.**

```
phys-schost# clsetup
```

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 8](#).

---

**4 Answer the prompt Do you want to add any quorum disks?.**

- **If your cluster is a two-node cluster, you must configure at least one shared quorum device. Type `Yes` to configure one or more quorum devices.**
- **If your cluster has three or more nodes, quorum device configuration is optional.**
  - **Type `No` if you do not want to configure additional quorum devices. Then skip to [Step 7](#).**

- **Type Yes to configure additional quorum devices. Then proceed to [Step 5](#).**
- 5 Specify what type of device you want to configure as a quorum device.**
- **Choose `scsi` to configure a shared SCSI disk.**
  - **Choose `quorum_server` to configure a quorum server as a quorum device.**
  - **Choose `netapp_nas` to configure a Network Appliance NAS device.**
- 6 Specify the name of the device to configure as a quorum device.**
- For a quorum server, also specify the following information:
    - The IP address of the quorum server host
    - The port number that is used by the quorum server to communicate with the cluster nodes
  - For a Network Appliance NAS device, also specify the following information:
    - The name of the NAS device
    - The LUN ID of the NAS device
- 7 At the prompt Is it okay to reset "installmode"?, type Yes.**
- After the `clsetup` utility sets the quorum configurations and vote counts for the cluster, the message Cluster initialization is complete is displayed. The utility returns you to the Main Menu.
- 8 Quit the `clsetup` utility.**

**Next Steps** Verify the quorum configuration and that installation mode is disabled. Go to [“How to Verify the Quorum Configuration and Installation Mode”](#) on page 121.

**Troubleshooting** **Interrupted `clsetup` processing** - If the quorum setup process is interrupted or fails to be completed successfully, rerun `clsetup`.

**Changes to quorum vote count** - If you later increase or decrease the number of node attachments to a quorum device, the quorum vote count is not automatically recalculated. You can reestablish the correct quorum vote by removing each quorum device and then add it back into the configuration, one quorum device at a time. For a two-node cluster, temporarily add a new quorum device before you remove and add back the original quorum device. Then remove the temporary quorum device. See the procedure “How to Modify a Quorum Device Node List” in Chapter 6, “Administering Quorum,” in *Sun Cluster System Administration Guide for Solaris OS*.



## ▼ How to Verify the Quorum Configuration and Installation Mode

Perform this procedure to verify that quorum configuration was completed successfully and that cluster installation mode is disabled.

You do not need to be superuser to run these commands.

### 1 From any node, verify the device and node quorum configurations.

```
phys-schost% clquorum list
```

Output lists each quorum device and each node.

### 2 From any node, verify that cluster installation mode is disabled.

```
phys-schost% cluster show -t global | grep installmode
installmode: disabled
```

Cluster installation is complete.

**Next Steps** Determine from the following list the next task to perform that applies to your cluster configuration. If you need to perform more than one task from this list, go to the first of those tasks in this list.

- If you want to change any private hostnames, go to [“How to Change Private Hostnames” on page 122](#).
- If you did not install your own `/etc/inet/ntp.conf` file before you installed Sun Cluster software, install or create the NTP configuration file. Go to [“How to Configure Network Time Protocol \(NTP\)” on page 123](#).
- To install a volume manager, go to [Chapter 4](#) and [Chapter 5](#) to install volume management software.

---

**Note** – If you added a new node to a cluster that uses VxVM, you must do one of the following tasks:

- Install VxVM on that node.
- Modify that node’s `/etc/name_to_major` file to support coexistence with VxVM.

Follow procedures in [“How to Install VERITAS Volume Manager Software” on page 160](#) to perform one of these required tasks.

- 
- To create cluster file systems, go to [“How to Create Cluster File Systems” on page 175](#).
  - To create non-global zones on a node, go to [“How to Create a Non-Global Zone on a Cluster Node” on page 179](#).
  - SPARC: To configure Sun Management Center to monitor the cluster, go to [“SPARC: Installing the Sun Cluster Module for Sun Management Center” on page 183](#).

- Install third-party applications, register resource types, set up resource groups, and configure data services. See the documentation that is supplied with the application software and the *Sun Cluster Data Services Planning and Administration Guide for Solaris OS*.

**See Also** Make a backup of your cluster configuration.

An archived backup of your cluster configuration facilitates easier recovery of the your cluster configuration,

For more information, see “How to Back Up the Cluster Configuration” in *Sun Cluster System Administration Guide for Solaris OS*.

## ▼ How to Change Private Hostnames

Perform this task if you do not want to use the default private hostnames, `clusternodeid-priv`, that are assigned during Sun Cluster software installation.

---

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

---

Perform this procedure on one active node of the cluster.

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

**2 Start the `clsetup` utility.**

```
phys-schost# clsetup
```

The `clsetup` Main Menu is displayed.

**3 Type the number that corresponds to the option for Private hostnames and press the Return key.**

The Private Hostname Menu is displayed.

**4 Type the number that corresponds to the option for Change a private hostname and press the Return key.**

**5 Follow the prompts to change the private hostname.**

Repeat for each private hostname to change.

**6 Verify the new private hostnames.**

```
phys-schost# clnode show -t node | grep privatehostname
privatehostname: clusternode1-priv
privatehostname: clusternode2-priv
privatehostname: clusternode3-priv
```

**Next Steps** Determine from the following list the next task to perform that applies to your cluster configuration. If you need to perform more than one task from this list, go to the first of those tasks in this list.

- If you did not install your own `/etc/inet/ntp.conf` file before you installed Sun Cluster software, install or create the NTP configuration file. Go to “[How to Configure Network Time Protocol \(NTP\)](#)” on page 123.
- To install a volume manager, go to [Chapter 4](#) and [Chapter 5](#) to install volume management software.

---

**Note** – If you added a new node to a cluster that uses VxVM, you must do one of the following tasks:

- Install VxVM on that node.
- Modify that node’s `/etc/name_to_major` file to support coexistence with VxVM.

Follow procedures in “[How to Install VERITAS Volume Manager Software](#)” on page 160 to perform one of these required tasks.

- 
- To create cluster file systems, go to “[How to Create Cluster File Systems](#)” on page 175.
  - To create non-global zones on a node, go to “[How to Create a Non-Global Zone on a Cluster Node](#)” on page 179.
  - SPARC: To configure Sun Management Center to monitor the cluster, go to “[SPARC: Installing the Sun Cluster Module for Sun Management Center](#)” on page 183.
  - Install third-party applications, register resource types, set up resource groups, and configure data services. See the documentation that is supplied with the application software and the *Sun Cluster Data Services Planning and Administration Guide for Solaris OS*.

## ▼ How to Configure Network Time Protocol (NTP)

---

**Note** – If you installed your own `/etc/inet/ntp.conf` file **before** you installed Sun Cluster software, you do not need to perform this procedure. Determine your next step:

---

Perform this task to create or modify the NTP configuration file after you perform any of the following tasks:

- Install Sun Cluster software
- Add a node to an existing cluster
- Change the private hostname of a node in the cluster

If you added a node to a single-node cluster, you must ensure that the NTP configuration file that you use is copied to the original cluster node as well as to the new node.

The primary requirement when you configure NTP, or any time synchronization facility within the cluster, is that all cluster nodes must be synchronized to the same time. Consider accuracy of time on

individual nodes to be of secondary importance to the synchronization of time among nodes. You are free to configure NTP as best meets your individual needs if this basic requirement for synchronization is met.

See the *Sun Cluster Concepts Guide for Solaris OS* for further information about cluster time. See the `/etc/inet/ntp.cluster` template file for additional guidelines on how to configure NTP for a Sun Cluster configuration.

- 1 Become superuser on a cluster node.**
- 2 If you have your own `/etc/inet/ntp.conf` file, copy your file to each node of the cluster.**
- 3 If you do not have your own `/etc/inet/ntp.conf` file to install, use the `/etc/inet/ntp.conf.cluster` file as your NTP configuration file.**

---

**Note** – Do not rename the `ntp.conf.cluster` file as `ntp.conf`.

---

If the `/etc/inet/ntp.conf.cluster` file does not exist on the node, you might have an `/etc/inet/ntp.conf` file from an earlier installation of Sun Cluster software. Sun Cluster software creates the `/etc/inet/ntp.conf.cluster` file as the NTP configuration file if an `/etc/inet/ntp.conf` file is not already present on the node. If so, perform the following edits instead on that `ntp.conf` file.

- a. **Use your preferred text editor to open the NTP configuration file on one node of the cluster for editing.**
  - b. **Ensure that an entry exists for the private hostname of each cluster node.**  
If you changed any node's private hostname, ensure that the NTP configuration file contains the new private hostname.
  - c. **If necessary, make other modifications to meet your NTP requirements.**
  - d. **Copy the NTP configuration file to all nodes in the cluster.**  
The contents of the NTP configuration file must be identical on all cluster nodes.
- 4 Stop the NTP daemon on each node.**

Wait for the command to complete successfully on each node before you proceed to [Step 5](#).

- SPARC: For the Solaris 9 OS, use the following command:

```
phys-schost# /etc/init.d/xntpd stop
```

- For the Solaris 10 OS, use the following command:

```
phys-schost# svcadm disable ntp
```

## 5 Restart the NTP daemon on each node.

- **If you use the `ntp.conf.cluster` file, run the following command:**

```
phys-schost# /etc/init.d/xntpd.cluster start
```

The `xntpd.cluster` startup script first looks for the `/etc/inet/ntp.conf` file.

- If the `ntp.conf` file exists, the script exits immediately without starting the NTP daemon.
- If the `ntp.conf` file does not exist but the `ntp.conf.cluster` file does exist, the script starts the NTP daemon. In this case, the script uses the `ntp.conf.cluster` file as the NTP configuration file.

- **If you use the `ntp.conf` file, run one of the following commands:**

- SPARC: For the Solaris 9 OS, use the following command:

```
phys-schost# /etc/init.d/xntpd start
```

- For the Solaris 10 OS, use the following command:

```
phys-schost# svcadm enable ntp
```

**Next Steps** Determine from the following list the next task to perform that applies to your cluster configuration. If you need to perform more than one task from this list, go to the first of those tasks in this list.

- To install a volume manager, go to [Chapter 4](#) and [Chapter 5](#) to install volume management software.

---

**Note** – If you added a new node to a cluster that uses VxVM, you must do one of the following tasks:

- Install VxVM on that node.
- Modify that node's `/etc/name_to_major` file to support coexistence with VxVM.

Follow procedures in “[How to Install VERITAS Volume Manager Software](#)” on page 160 to perform one of these required tasks.

- 
- To create cluster file systems, go to “[How to Create Cluster File Systems](#)” on page 175.
  - To create non-global zones on a node, go to “[How to Create a Non-Global Zone on a Cluster Node](#)” on page 179.
  - SPARC: To configure Sun Management Center to monitor the cluster, go to “[SPARC: Installing the Sun Cluster Module for Sun Management Center](#)” on page 183.
  - Install third-party applications, register resource types, set up resource groups, and configure data services. See the documentation that is supplied with the application software and the *Sun Cluster Data Services Planning and Administration Guide for Solaris OS*.



# Configuring Solaris Volume Manager Software

---

Configure your local and multihost disks for Solaris Volume Manager software by using the procedures in this chapter, along with the planning information in [“Planning Volume Management” on page 36](#). See your Solaris Volume Manager documentation for additional details.

---

**Note** – The Enhanced Storage module of Solaris Management Console is not compatible with Sun Cluster software. Use the command-line interface or Sun Cluster utilities to configure Solaris Volume Manager software.

---

The following sections are in this chapter:

- [“Configuring Solaris Volume Manager Software” on page 127](#)
- [“Creating Disk Sets in a Cluster” on page 146](#)
- [“Configuring Dual-String Mediators” on page 155](#)

## Configuring Solaris Volume Manager Software

The following table lists the tasks that you perform to configure Solaris Volume Manager software for Sun Cluster configurations.

TABLE 4-1 Task Map: Configuring Solaris Volume Manager Software

| Task                                                                                                                                                                        | Instructions                                                                             |
|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------------------------------------------------------------------------------|
| 1. Plan the layout of your Solaris Volume Manager configuration.                                                                                                            | <a href="#">“Planning Volume Management” on page 36</a>                                  |
| 2. ( <i>Solaris 9 only</i> ) Calculate the number of volume names and disk sets that you need for your configuration, and modify the <code>/kernel/drv/md.conf</code> file. | <a href="#">“SPARC: How to Set the Number of Volume Names and Disk Sets” on page 128</a> |

TABLE 4-1 Task Map: Configuring Solaris Volume Manager Software (Continued)

| Task                                                  | Instructions                                                        |
|-------------------------------------------------------|---------------------------------------------------------------------|
| 3. Create state database replicas on the local disks. | <a href="#">“How to Create State Database Replicas” on page 129</a> |
| 4. (Optional) Mirror file systems on the root disk.   | <a href="#">“Mirroring the Root Disk” on page 130</a>               |

## ▼ SPARC: How to Set the Number of Volume Names and Disk Sets

**Note** – This procedure is only required for the Solaris 9 OS. If the cluster runs on the Solaris 10 OS, proceed to [“How to Create State Database Replicas” on page 129](#).

With the Solaris 10 release, Solaris Volume Manager has been enhanced to configure volumes dynamically. You no longer need to edit the `nmd` and the `md_nsets` parameters in the `/kernel/drv/md.conf` file. New volumes are dynamically created, as needed.

This procedure describes how to determine the number of Solaris Volume Manager volume names and disk sets that you need 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 volume names per disk set is 128, but many configurations need more than the default. Increase this number before you implement a configuration, to save administration time later.

At the same time, keep the value of the `nmdfield` and the `md_nsets` field as low as possible. Memory structures exist for all possible devices as determined by `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 volumes that you plan to use.

**Before You Begin** Have available the completed [“Device Group Configurations Worksheet” on page 279](#).

### 1 Calculate the total number of disk sets that you expect to need in the cluster, then add one more disk set for private disk management.

The cluster can have a maximum of 32 disk sets, 31 disk sets for general use plus one disk set for private disk management. The default number of disk sets is 4. You supply this value for the `md_nsets` field in [Step 3](#).



**2 Calculate the largest volume name that you expect to need for any disk set in the cluster.**

Each disk set can have a maximum of 8192 volume names. You supply this value for the `nmd` field in [Step 3](#).

**a. Determine the quantity of volume names that you expect to need for each disk set.**

If you use local volumes, ensure that each local volume name on which a global-devices file system, `/global/.devices/node@nodeid`, is mounted is unique throughout the cluster and does not use the same name as any device-ID name in the cluster.

---

**Tip** – Choose a range of numbers to use exclusively for device-ID names and a range for each node to use exclusively for its local volume names. For example, device-ID names might use the range from `d1` to `d100`. Local volumes on node 1 might use names in the range from `d100` to `d199`. And local volumes on node 2 might use `d200` to `d299`.

---

**b. Calculate the highest of the volume names that you expect to use in any disk set.**

The quantity of volume names to set is based on the volume name *value* rather than on the *actual quantity*. For example, if your volume names range from `d950` to `d1000`, Solaris Volume Manager software requires that you set the value at 1000 names, not 50.

**3 On each node, become superuser and edit the `/kernel/drv/md.conf` file.**


---

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

---

**a. Set the `md_nsets` field to the value that you determined in [Step 1](#).****b. Set the `nmd` field to the value that you determined in [Step 2](#).****4 On each node, perform a reconfiguration reboot.**

```
phys-schost# touch /reconfigure
phys-schost# shutdown -g0 -y -i6
```

Changes to the `/kernel/drv/md.conf` file become operative after you perform a reconfiguration reboot.

**Next Steps** Create local state database replicas. Go to [“How to Create State Database Replicas”](#) on page 129.

## ▼ How to Create State Database Replicas

Perform this procedure on each node in the cluster.

**1 Become superuser.**

**2 Create state database replicas on one or more local devices for each cluster node.**

Use the physical name (cNtXdY sZ), not the device-ID name (dN), to specify the slices to use.

```
phys-schost# metadb -af slice-1 slice-2 slice-3
```

---

**Tip** – To provide protection of state data, which is necessary to run Solaris Volume Manager software, create at least three replicas for each node. Also, you can place replicas on more than one device to provide protection if one of the devices fails.

---

See the `metadb(1M)` man page and your Solaris Volume Manager documentation for details.

**3 Verify the replicas.**

```
phys-schost# metadb
```

The `metadb` command displays the list of replicas.

**Example 4–1 Creating State Database Replicas**

The following example shows three state database replicas. Each replica is created on a different device.

```
phys-schost# metadb -af c0t0d0s7 c0t1d0s7 c1t0d0s7
phys-schost# metadb
flags first blk block count
a u 16 8192 /dev/dsk/c0t0d0s7
a u 16 8192 /dev/dsk/c0t1d0s7
a u 16 8192 /dev/dsk/c1t0d0s7
```

**Next Steps** To mirror file systems on the root disk, go to [“Mirroring the Root Disk” on page 130](#).

Otherwise, go to [“Creating Disk Sets in a Cluster” on page 146](#) to create Solaris Volume Manager disk sets.

## 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 131](#)
- [“How to Mirror the Global Devices Namespace” on page 134](#)
- [“How to Mirror File Systems Other Than Root \(/\) That Cannot Be Unmounted” on page 138](#)
- [“How to Mirror File Systems That Can Be Unmounted” on page 142](#)




---

**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.

---

**Note** – This procedure provides the long forms of the Sun Cluster commands. Most commands also have short forms. Except for the forms of the command names, the commands are identical. For a list of the commands and their short forms, see Appendix A, “Sun Cluster Object-Oriented Commands,” in *Sun Cluster System Administration Guide for Solaris OS*.

---

### 1 Become superuser.

### 2 Place the root slice in a single-slice (one-way) concatenation.

Specify the physical disk name of the root-disk slice (`cNtXdYsZ`).

```
phys-schost# metainit -f submirror1 1 1 root-disk-slice
```

### 3 Create a second concatenation.

```
phys-schost# metainit submirror2 1 1 submirror-disk-slice
```

### 4 Create a one-way mirror with one submirror.

```
phys-schost# metainit mirror -m submirror1
```

---

**Note** – If the device is a local device to be used to mount a global-devices file system, `/global/.devices/node@nodeid`, the volume name for the mirror *must* be unique throughout the cluster.

---

### 5 Set up the system files for the root (/) directory.

```
phys-schost# metaroot mirror
```

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

### 6 Flush all file systems.

```
phys-schost# lockfs -fa
```

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

**7 Move any resource groups or device groups from the node.**

```
phys-schost# clnode evacuate from-node
```

*from-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.

```
phys-schost# shutdown -g0 -y -i6
```

**9 Attach the second submirror to the mirror.**

```
phys-schost# metattach mirror submirror2
```

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

**10 If the disk that is used to mirror the root disk is physically connected to more than one node (multihosted), modify the device group's properties to support its use as a mirror.**

Ensure that the device group meets the following requirements:

- The raw-disk device group must have only one node configured in its node list.
- The `localonly` property of the raw-disk device group must be enabled. The `localonly` property prevents unintentional fencing of a node from its boot device if the boot device is connected to multiple nodes.

**a. If necessary, use the `cldevice` command to determine the name of the raw-disk device group.**

```
phys-schost# cldevice show node:/dev/rdisk/cNtXdY
```

---

**Tip** – If you issue the command from a node that is physically connected to the disk, you can specify the disk name as `cNtXdY` instead of by its full device path name.

---

In the following example, the raw-disk device-group name `dsk/d2` is part of the DID device name.

```
=== DID Device Instances ===
```

```
DID Device Name: /dev/did/rdisk/d2
Full Device Path: phys-schost-1:/dev/rdisk/c1t1d0
Full Device Path: phys-schost-3:/dev/rdisk/c1t1d0
...
```

See the `cldevice(1CL)` man page for more information.

**b. View the node list of the raw-disk device group.**

```
phys-schost# cldevicegroup show dsk/dN
```

Output looks similar to the following for the device group `dsk/d2`:

```
Device Group Name: dsk/d2
...
Node List: phys-schost-1, phys-schost-3
...
localonly: false
```

- c. **If the node list contains more than one node name, remove all nodes from the node list except the node whose root disk you mirrored.**

Only the node whose root disk you mirrored should remain in the node list for the raw-disk device group.

```
phys-schost# cldevicegroup remove-node -n node devicegroup
```

`-n node` Specifies the node to remove from the device-group node list.

- d. **Enable the `localonly` property of the raw-disk device group, if it is not already enabled.**

When the `localonly` property is enabled, the raw-disk device group is used exclusively by the node in its node list. This usage prevents unintentional fencing of the node from its boot device if the boot device is connected to multiple nodes.

```
phys-schost# cldevicegroup set -p localonly=true devicegroup
```

`-p` Sets the value of a device-group property.

`localonly=true` Enables the `localonly` property of the device group.

For more information about the `localonly` property, see the `cldevicegroup(1CL)` man page.

## 11 Record the alternate boot path for possible future use.

If the primary boot device fails, you can then boot from this alternate boot device. See “Special Considerations for Mirroring root (/)” in *Solaris Volume Manager Administration Guide* or “Creating a RAID-1 Volume” in *Solaris Volume Manager Administration Guide* for more information about alternate boot devices.

```
phys-schost# ls -l /dev/rdisk/root-disk-slice
```

## 12 Repeat Step 1 through Step 11 on each remaining node of the cluster.

Ensure that each volume name for a mirror on which a global-devices file system, `/global/.devices/node@nodeid`, is to be mounted is unique throughout the cluster.

### Example 4–2 Mirroring the Root (/) File System

The following example shows the creation of mirror `d0` on the node `phys-schost-1`, which consists of submirror `d10` on partition `c0t0d0s0` and submirror `d20` on partition `c2t2d0s0`. Device `c2t2d0` is a multihost disk, so the `localonly` property is enabled. The example also displays the alternate boot path for recording.

```

phys-schost# metainit -f d10 1 1 c0t0d0s0
d11: Concat/Stripe is setup
phys-schost# metainit d20 1 1 c2t2d0s0
d12: Concat/Stripe is setup
phys-schost# metainit d0 -m d10
d10: Mirror is setup
phys-schost# metaroot d0
phys-schost# lockfs -fa
phys-schost# clnode evacuate phys-schost-1
phys-schost# shutdown -g0 -y -i6
phys-schost# metattach d0 d20
d0: Submirror d20 is attached
phys-schost# cldevicegroup show dsk/d2
Device Group Name: dsk/d2
...
Node List: phys-schost-1, phys-schost-3
...
localonly: false
phys-schost# cldevicegroup remove-node -n phys-schost-3 dsk/d2
phys-schost# cldevicegroup set -p localonly=true dsk/d2
phys-schost# 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

```

**Next Steps** To mirror the global devices namespace, `/global/.devices/node@nodeid`, go to [“How to Mirror the Global Devices Namespace” on page 134](#).

To mirror file systems that cannot be unmounted, go to [“How to Mirror File Systems Other Than Root \(/\) That Cannot Be Unmounted” on page 138](#).

To mirror user-defined file systems, go to [“How to Mirror File Systems That Can Be Unmounted” on page 142](#).

Otherwise, go to [“Creating Disk Sets in a Cluster” on page 146](#) to create a disk set.

**Troubleshooting** Some of the steps in this mirroring procedure might cause an error message similar to `metainit: dg-schost-1: d1s0: not a metadvice`. Such an error message is harmless and can be ignored.

## ▼ How to Mirror the Global Devices Namespace

Use this procedure to mirror the global devices namespace, `/global/.devices/node@nodeid/`.

---

**Note** – This procedure provides the long forms of the Sun Cluster commands. Most commands also have short forms. Except for the forms of the command names, the commands are identical. For a list of the commands and their short forms, see Appendix A, “Sun Cluster Object-Oriented Commands,” in *Sun Cluster System Administration Guide for Solaris OS*.

---

**1 Become superuser.**

**2 Place the global devices namespace slice in a single-slice (one-way) concatenation.**

Use the physical disk name of the disk slice (cNtXdY sZ).

```
phys-schost# metainit -f submirror1 1 1 diskslice
```

**3 Create a second concatenation.**

```
phys-schost# metainit submirror2 1 1 submirror-diskslice
```

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

```
phys-schost# metainit mirror -m submirror1
```

---

**Note** – The volume name for a mirror on which a global-devices file system, `/global/.devices/node@nodeid`, is to be mounted *must* be unique throughout the cluster.

---

**5 Attach the second submirror to the mirror.**

This attachment starts a synchronization of the submirrors.

```
phys-schost# 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.

```
phys-schost# vi /etc/vfstab
#device device mount FS fsck mount mount
#to mount to fsck point type pass at boot options
#
/dev/md/dsk/mirror /dev/md/rdisk/mirror /global/.devices/node@nodeid ufs 2 no global
```

**7 Repeat Step 1 through Step 6 on each remaining node of the cluster.**

**8 Wait for the synchronization of the mirrors, started in Step 5, to be completed.**

Use the `metastat(1M)` command to view mirror status and to verify that mirror synchronization is complete.

```
phys-schost# metastat mirror
```

- 9 If the disk that is used to mirror the global devices namespace is physically connected to more than one node (multihosted), ensure that the device-group node list contains only one node and that the `localonly` property is enabled.**

Ensure that the device group meets the following requirements:

- The raw-disk device group must have only one node configured in its node list.
- The `localonly` property of the raw-disk device group must be enabled. The `localonly` property prevents unintentional fencing of a node from its boot device if the boot device is connected to multiple nodes.

- a. If necessary, use the `cldevice` command to determine the name of the raw-disk device group.**

```
phys-schost# cldevice show node:/dev/rdisk/cNtXdY
```

---

**Tip** – If you issue the command from a node that is physically connected to the disk, you can specify the disk name as `cNtXdY` instead of by its full device path name.

---

In the following example, the raw-disk device-group name `dsk/d2` is part of the DID device name.

```
=== DID Device Instances ===
```

```
DID Device Name: /dev/did/rdisk/d2
 Full Device Path: phys-schost-1:/dev/rdisk/clt1d0
 Full Device Path: phys-schost-3:/dev/rdisk/clt1d0
 ...
```

See the `cldevice(1CL)` man page for more information.

- b. View the node list of the raw-disk device group.**

```
phys-schost# cldevicegroup show dsk/dN
```

Output looks similar to the following for the device group `dsk/d2`:

```
Device Group Name: dsk/d2
 ...
 Node List: phys-schost-1, phys-schost-3
 ...
 localonly: false
```

- c. If the node list contains more than one node name, remove all nodes from the node list except the node whose root disk you mirrored.**

Only the node whose root disk you mirrored should remain in the node list for the raw-disk device group.

```
phys-schost# cldevicegroup remove-node -n node devicegroup
```

`-n node` Specifies the node to remove from the device-group node list.



**d. Enable the `localonly` property of the raw-disk device group, if it is not already enabled.**

When the `localonly` property is enabled, the raw-disk device group is used exclusively by the node in its node list. This usage prevents unintentional fencing of the node from its boot device if the boot device is connected to multiple nodes.

```
phys-schost# cldevicegroup set -p localonly=true devicegroup
```

`-p` Sets the value of a device-group property.

`localonly=true` Enables the `localonly` property of the device group.

For more information about the `localonly` property, see the `cldevicegroup(1CL)` man page.

**Example 4-3 Mirroring the Global Devices Namespace**

The following example shows creation of mirror `d101`, which consists of submirror `d111` on partition `c0t0d0s3` and submirror `d121` on partition `c2t2d0s3`. The `/etc/vfstab` file entry for `/global/.devices/node@1` is updated to use the mirror name `d101`. Device `c2t2d0` is a multihost disk, so the `localonly` property is enabled.

```
phys-schost# metainit -f d111 1 1 c0t0d0s3
d111: Concat/Stripe is setup
phys-schost# metainit d121 1 1 c2t2d0s3
d121: Concat/Stripe is setup
phys-schost# metainit d101 -m d111
d101: Mirror is setup
phys-schost# metattach d101 d121
d101: Submirror d121 is attached
phys-schost# 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
phys-schost# metastat d101
d101: Mirror
 Submirror 0: d111
 State: Okay
 Submirror 1: d121
 State: Resyncing
 Resync in progress: 15 % done
...
phys-schost# cldevice show phys-schost-3:/dev/rdisk/c2t2d0
=== DID Device Instances ===

DID Device Name: /dev/did/rdisk/d2
Full Device Path: phys-schost-1:/dev/rdisk/c2t2d0
Full Device Path: phys-schost-3:/dev/rdisk/c2t2d0
...
```

```

phys-schost# cldevicegroup show | grep dsk/d2
Device Group Name: dsk/d2
...
Node List: phys-schost-1, phys-schost-3
...
localonly: false
phys-schost# cldevicegroup remove-node -n phys-schost-3 dsk/d2
phys-schost# cldevicegroup set -p localonly=true dsk/d2

```

**Next Steps** To mirror file systems other than root (/) that cannot be unmounted, go to [“How to Mirror File Systems Other Than Root \(/\) That Cannot Be Unmounted”](#) on page 138.

To mirror user-defined file systems, go to [“How to Mirror File Systems That Can Be Unmounted”](#) on page 142

Otherwise, go to [“Creating Disk Sets in a Cluster”](#) on page 146 to create a disk set.

**Troubleshooting** Some of the steps in this mirroring procedure might cause an error message similar to `metainit: dg-schost-1: d1s0: not a metadvice`. Such an error message is harmless and can be ignored.

## ▼ 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.

---

**Note** – This procedure provides the long forms of the Sun Cluster commands. Most commands also have short forms. Except for the forms of the command names, the commands are identical. For a list of the commands and their short forms, see Appendix A, “Sun Cluster Object-Oriented Commands,” in *Sun Cluster System Administration Guide for Solaris OS*.

---

- 1 **Become superuser.**
- 2 **Place the slice on which an unmountable file system resides in a single-slice (one-way) concatenation.**

Specify the physical disk name of the disk slice (cNtX dYsZ).

```
phys-schost# metainit -f submirror1 1 1 diskslice
```

- 3 **Create a second concatenation.**

```
phys-schost# metainit submirror2 1 1 submirror-diskslice
```

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

```
phys-schost# metainit mirror -m submirror1
```

---

**Note** – The volume name for this mirror does *not* need to be unique throughout the cluster.

---

**5 Repeat Step 1 through Step 4 for each remaining unmountable file system that you want to mirror.****6 On each node, edit the /etc/vfstab file entry for each unmountable file system you mirrored.**

Replace the names in the device to mount and device to fsck columns with the mirror name.

```
phys-schost# vi /etc/vfstab
#device device mount FS fsck mount mount
#to mount to fsck point type pass at boot options
#
/dev/md/dsk/mirror /dev/md/rdisk/mirror /filesystem ufs 2 no global
```

**7 Move any resource groups or device groups from the node.**

```
phys-schost# clnode evacuate from-node
```

*from-node* Specifies the name of the node from which to move resource or device groups.

**8 Reboot the node.**

```
phys-schost# shutdown -g0 -y -i6
```

**9 Attach the second submirror to each mirror.**

This attachment starts a synchronization of the submirrors.

```
phys-schost# metattach mirror submirror2
```

**10 Wait for the synchronization of the mirrors, started in Step 9, to complete.**

Use the `metastat(1M)` command to view mirror status and to verify that mirror synchronization is complete.

```
phys-schost# metastat mirror
```

**11 If the disk that is used to mirror the unmountable file system is physically connected to more than one node (multihosted), ensure that the device-group node list contains only one node and that the localonly property is enabled.**

Ensure that the device group meets the following requirements:

- The raw-disk device group must have only one node configured in its node list.

- The `localonly` property of the raw-disk device group must be enabled. The `localonly` property prevents unintentional fencing of a node from its boot device if the boot device is connected to multiple nodes.

**a. If necessary, use the `cldevice` command to determine the name of the raw-disk device group.**

```
phys-schost# cldevice show node:/dev/rdisk/cNtXdY
```

---

**Tip** – If you issue the command from a node that is physically connected to the disk, you can specify the disk name as `cNtXdY` instead of by its full device path name.

---

In the following example, the raw-disk device-group name `dsk/d2` is part of the DID device name.

```
=== DID Device Instances ===
```

```
DID Device Name: /dev/did/rdsk/d2
 Full Device Path: phys-schost-1:/dev/rdisk/c1t1d0
 Full Device Path: phys-schost-3:/dev/rdisk/c1t1d0
 ...
```

See the `cldevice(1CL)` man page for more information.

**b. View the node list of the raw-disk device group.**

```
phys-schost# cldevicegroup show dsk/dN
```

Output looks similar to the following for the device group `dsk/d2`:

```
Device Group Name: dsk/d2
 ...
 Node List: phys-schost-1, phys-schost-3
 ...
 localonly: false
```

**c. If the node list contains more than one node name, remove all nodes from the node list except the node whose root disk you mirrored.**

Only the node whose root disk you mirrored should remain in the node list for the raw-disk device group.

```
phys-schost# cldevicegroup remove-node -n node devicegroup
```

`-n node` Specifies the node to remove from the device-group node list.

**d. Enable the `localonly` property of the raw-disk device group, if it is not already enabled.**

When the `localonly` property is enabled, the raw-disk device group is used exclusively by the node in its node list. This usage prevents unintentional fencing of the node from its boot device if the boot device is connected to multiple nodes.

```
phys-schost# cldevicegroup set -p localonly=true devicegroup
```

-p                      Sets the value of a device-group property.

localonly=true        Enables the `localonly` property of the device group.

For more information about the `localonly` property, see the `cldevicegroup(1CL)` man page.

**Example 4-4 Mirroring File Systems That Cannot Be Unmounted**

The following example shows the creation of mirror `d1` on the node `phys-schost-1` to mirror `/usr`, which resides on `c0t0d0s1`. Mirror `d1` consists of submirror `d11` on partition `c0t0d0s1` and submirror `d21` on partition `c2t2d0s1`. The `/etc/vfstab` file entry for `/usr` is updated to use the mirror name `d1`. Device `c2t2d0` is a multihost disk, so the `localonly` property is enabled.

```
phys-schost# metainit -f d11 1 1 c0t0d0s1
d11: Concat/Stripe is setup
phys-schost# metainit d21 1 1 c2t2d0s1
d21: Concat/Stripe is setup
phys-schost# metainit d1 -m d11
d1: Mirror is setup
phys-schost# 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
...
phys-schost# clnode evacuate phys-schost-1
phys-schost# shutdown -g0 -y -i6
phys-schost# metattach d1 d21
d1: Submirror d21 is attached
phys-schost# metastat d1
d1: Mirror
 Submirror 0: d11
 State: Okay
 Submirror 1: d21
 State: Resyncing
 Resync in progress: 15 % done
...
phys-schost# cldevice show phys-schost-3:/dev/rdsk/c2t2d0
...
DID Device Name: /dev/did/rdsk/d2
phys-schost# cldevicegroup show dsk/d2
```

```

Device Group Name: dsk/d2
...
Node List: phys-schost-1, phys-schost-3
...
localonly: false
phys-schost# cldevicegroup remove-node -n phys-schost-3 dsk/d2
phys-schost# cldevicegroup set -p localonly=true dsk/d2

```

**Next Steps** To mirror user-defined file systems, go to [“How to Mirror File Systems That Can Be Unmounted” on page 142.](#)

Otherwise, go to [“Creating Disk Sets in a Cluster” on page 146](#) to create a disk set.

**Troubleshooting** Some of the steps in this mirroring procedure might cause an error message similar to `metainit: dg-schost-1: d1s0: not a metadvice`. Such an error message is harmless and can be ignored.

## ▼ How to Mirror File Systems That Can Be Unmounted

Use this procedure to mirror user-defined file systems that can be unmounted. In this procedure, the nodes do not need to be rebooted.

---

**Note** – This procedure provides the long forms of the Sun Cluster commands. Most commands also have short forms. Except for the forms of the command names, the commands are identical. For a list of the commands and their short forms, see Appendix A, “Sun Cluster Object-Oriented Commands,” in *Sun Cluster System Administration Guide for Solaris OS*.

---

### 1 Become superuser.

### 2 Unmount the file system to mirror.

Ensure that no processes are running on the file system.

```
phys-schost# umount /mount-point
```

See the `umount(1M)` man page and Chapter 19, “Mounting and Unmounting File Systems (Tasks),” in *System Administration Guide: Devices and File Systems* for more information.

### 3 Place in a single-slice (one-way) concatenation the slice that contains a user-defined file system that can be unmounted.

Specify the physical disk name of the disk slice (`cNtX dYsZ`).

```
phys-schost# metainit -f submirror1 1 1 diskslice
```

### 4 Create a second concatenation.

```
phys-schost# metainit submirror2 1 1 submirror-diskslice
```

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

```
phys-schost# metainit mirror -m submirror1
```

---

**Note** – The volume name for this mirror does *not* need to be unique throughout the cluster.

---

**6 Repeat Step 1 through Step 5 for each mountable file system to be mirrored.****7 On each node, edit the /etc/vfstab file entry for each file system you mirrored.**

Replace the names in the device to mount and device to fsck columns with the mirror name.

```
phys-schost# vi /etc/vfstab
#device device mount FS fsck mount mount
#to mount to fsck point type pass at boot options
#
/dev/md/dsk/mirror /dev/md/rdisk/mirror /filesystem ufs 2 no global
```

**8 Attach the second submirror to the mirror.**

This attachment starts a synchronization of the submirrors.

```
phys-schost# metattach mirror submirror2
```

**9 Wait for the synchronization of the mirrors, started in Step 8, to be completed.**

Use the `metastat(1M)` command to view mirror status.

```
phys-schost# metastat mirror
```

**10 If the disk that is used to mirror the user-defined file system is physically connected to more than one node (multihosted), ensure that the device-group node list contains only one node and that the localonly property is enabled.**

Ensure that the device group meets the following requirements:

- The raw-disk device group must have only one node configured in its node list.
- The `localonly` property of the raw-disk device group must be enabled. The `localonly` property prevents unintentional fencing of a node from its boot device if the boot device is connected to multiple nodes.

**a. If necessary, use the `cldevice` command to determine the name of the raw-disk device group.**

```
phys-schost# cldevice show node:/dev/rdisk/cNtXdY
```

---

**Tip** – If you issue the command from a node that is physically connected to the disk, you can specify the disk name as `cNtXdY` instead of by its full device path name.

---

In the following example, the raw-disk device-group name `dsk/d2` is part of the DID device name.

```
=== DID Device Instances ===
```

```
DID Device Name: /dev/did/rdsk/d2
 Full Device Path: phys-schost-1:/dev/rdsk/clt1d0
 Full Device Path: phys-schost-3:/dev/rdsk/clt1d0
 ...
```

See the `cldevice(1CL)` man page for more information.

**b. View the node list of the raw-disk device group.**

```
phys-schost# cldevicegroup show dsk/dN
```

Output looks similar to the following for the device group `dsk/d2`:

```
Device Group Name: dsk/d2
 ...
 Node List: phys-schost-1, phys-schost-3
 ...
 localonly: false
```

**c. If the node list contains more than one node name, remove all nodes from the node list except the node whose root disk you mirrored.**

Only the node whose root disk you mirrored should remain in the node list for the raw-disk device group.

```
phys-schost# cldevicegroup remove-node -n node devicegroup
```

`-n node` Specifies the node to remove from the device-group node list.

**d. Enable the `localonly` property of the raw-disk device group, if it is not already enabled.**

When the `localonly` property is enabled, the raw-disk device group is used exclusively by the node in its node list. This usage prevents unintentional fencing of the node from its boot device if the boot device is connected to multiple nodes.

```
phys-schost# cldevicegroup set -p localonly=true devicegroup
```

`-p` Sets the value of a device-group property.

`localonly=true` Enables the `localonly` property of the device group.

For more information about the `localonly` property, see the `cldevicegroup(1CL)` man page.

**11 Mount the mirrored file system.**

```
phys-schost# mount /mount-point
```

See the `mount(1M)` man page and Chapter 19, “Mounting and Unmounting File Systems (Tasks),” in *System Administration Guide: Devices and File Systems* for more information.



**Example 4–5** Mirroring File Systems That Can Be Unmounted

The following example shows creation of mirror d4 to mirror /export, which resides on c0t0d0s4. Mirror d4 consists of submirror d14 on partition c0t0d0s4 and submirror d24 on partition c2t2d0s4. The /etc/vfstab file entry for /export is updated to use the mirror name d4. Device c2t2d0 is a multihost disk, so the localOnly property is enabled.

```

phys-schost# umount /export
phys-schost# metainit -f d14 1 1 c0t0d0s4
d14: Concat/Stripe is setup
phys-schost# metainit d24 1 1 c2t2d0s4
d24: Concat/Stripe is setup
phys-schost# metainit d4 -m d14
d4: Mirror is setup
phys-schost# vi /etc/vfstab
#device device mount FS fsck mount mount
#to mount to fsck point type pass at boot options
#
/dev/md/dsk/d4 /dev/md/rdisk/d4 /export ufs 2 no global
phys-schost# metattach d4 d24
d4: Submirror d24 is attached
phys-schost# metastat d4
d4: Mirror
 Submirror 0: d14
 State: Okay
 Submirror 1: d24
 State: Resyncing
 Resync in progress: 15 % done
...
phys-schost# cldevice show phys-schost-3:/dev/rdsk/c2t2d0
...
DID Device Name: /dev/did/rdsk/d2
phys-schost# cldevicegroup show dsk/d2
Device Group Name: dsk/d2
...
 Node List: phys-schost-1, phys-schost-2
...
 localOnly: false
phys-schost# cldevicegroup remove-node -n phys-schost-3 dsk/d2
phys-schost# cldevicegroup set -p localOnly=true dsk/d2
phys-schost# mount /export

```

**Next Steps** If you need to create disk sets, go to one of the following:

- To create a Solaris Volume Manager for Sun Cluster disk set for use by Oracle RAC, go to “How to Create a Multi-Owner Disk Set in Solaris Volume Manager for Sun Cluster for the Oracle RAC Database” in *Sun Cluster Data Service for Oracle RAC Guide for Solaris OS*.

- To create a disk set for any other application, go to [“Creating Disk Sets in a Cluster”](#) on page 146.

If you have sufficient disk sets for your needs, go to one of the following:

- If your cluster contains disk sets that are configured with exactly two disk enclosures and two nodes, you must add dual-string mediators. Go to [“Configuring Dual-String Mediators”](#) on page 155.
- If your cluster configuration does not require dual-string mediators, go to [“How to Create Cluster File Systems”](#) on page 175.

**Troubleshooting**

Some of the steps in this mirroring procedure might cause an error message that is similar to `metainit: dg-schost-1: d1s0: not a metadvice`. Such an error message is harmless and can be ignored.

## Creating Disk Sets in a Cluster

This section describes how to create disk sets for a cluster configuration. When you create a Solaris Volume Manager disk set in a Sun Cluster environment, the disk set is automatically registered with the Sun Cluster software as a device group of type `svm`. To create or delete an `svm` device group, you must use Solaris Volume Manager commands and utilities to create or delete the underlying disk set of the device group.

---

**Note** – To create a Solaris Volume Manager for Sun Cluster disk set for use by Oracle RAC, do not use these procedures. Instead, perform the procedures in [“How to Create a Multi-Owner Disk Set in Solaris Volume Manager for Sun Cluster for the Oracle RAC Database”](#) in *Sun Cluster Data Service for Oracle RAC Guide for Solaris OS*.

---

The following table lists the tasks that you perform to create disk sets.

**TABLE 4–2** Task Map: Installing and Configuring Solaris Volume Manager Software

| Task                                                                                             | Instructions                                                            |
|--------------------------------------------------------------------------------------------------|-------------------------------------------------------------------------|
| 1. Create disk sets by using the <code>metaset</code> command.                                   | <a href="#">“How to Create a Disk Set”</a> on page 147                  |
| 2. Add drives to the disk sets.                                                                  | <a href="#">“How to Add Drives to a Disk Set”</a> on page 150           |
| 3. (Optional) Repartition drives in a disk set to allocate space to different slices.            | <a href="#">“How to Repartition Drives in a Disk Set”</a> on page 152   |
| 4. List DID pseudo-driver mappings and define volumes in the <code>/etc/lvm/md.tab</code> files. | <a href="#">“How to Create an <code>md.tab</code> File”</a> on page 152 |
| 5. Initialize the <code>md.tab</code> files.                                                     | <a href="#">“How to Activate Volumes”</a> on page 154                   |

## ▼ How to Create a Disk Set

Perform this procedure to create disk sets.

---

**Note** – This procedure provides the long forms of the Sun Cluster commands. Most commands also have short forms. Except for the forms of the command names, the commands are identical. For a list of the commands and their short forms, see Appendix A, “Sun Cluster Object-Oriented Commands,” in *Sun Cluster System Administration Guide for Solaris OS*.

---

- 1 **SPARC: (Solaris 9) Determine whether, after you create the new disk sets, the cluster will have more than three disk sets.**
  - If the cluster will have no more than three disk sets, skip to [Step 9](#).
  - If the cluster will have four or more disk sets, proceed to [Step 2](#) to prepare the cluster. You must perform this task whether you are installing disk sets for the first time or whether you are adding more disk sets to a fully configured cluster.
  - If the cluster runs on the Solaris 10 OS, Solaris Volume Manager automatically makes the necessary configuration changes. Skip to [Step 9](#).
- 2 **On any node of the cluster, check the value of the `md_nsets` variable in the `/kernel/drv/md.conf` file.**
- 3 **If the total number of disk sets in the cluster will be greater than the existing value of `md_nsets` minus one, increase the value of `md_nsets` to the desired value.**

The maximum permissible number of disk sets is one less than the configured value of `md_nsets`. The maximum possible value of `md_nsets` is 32, therefore the maximum permissible number of disk sets that you can create is 31.
- 4 **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 Solaris Volume Manager errors and possible loss of data.

---

- 5 **If you made changes to the `md.conf` file on any node, perform the following steps to make those changes active.**
  - a. **On one node, become superuser.**
  - b. **From one node, shut down the cluster.**

```
phys-schost# cluster shutdown -g0 -y
```

**c. Reboot each node of the cluster.**

- **On SPARC based systems, do the following:**

```
ok boot
```

- **On x86 based systems, do the following:**

When the GRUB menu is displayed, select the appropriate Solaris entry and press Enter. The GRUB menu appears similar to the following:

```
GNU GRUB version 0.95 (631K lower / 2095488K upper memory)
+-----+
| Solaris 10 /sol_10_x86 |
| Solaris failsafe |
| |
+-----+
```

Use the ^ and v keys to select which entry is highlighted. Press enter to boot the selected OS, 'e' to edit the commands before booting, or 'c' for a command-line.

For more information about GRUB based booting, see Chapter 11, “GRUB Based Booting (Tasks),” in *System Administration Guide: Basic Administration*.

**6 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.

**7 From one node of the cluster, update the global-devices namespace.**

```
phys-schost# cldevice populate
```

See the `cldevice(1CL)` man page for more information.

**8 On each node, verify that the command has completed processing before you attempt to create any disk sets.**

The command executes remotely on all nodes, even though the command is run from just one node. To determine whether the command has completed processing, run the following command on each node of the cluster.

```
phys-schost# ps -ef | grep scgdevs
```

**9 Ensure that the disk set that you intend to create meets one of the following requirements.**

- If the disk set is configured with exactly two disk strings, the disk set must connect to exactly two nodes and use exactly two mediator hosts. These mediator hosts must be the same two hosts used for the disk set. See “[Configuring Dual-String Mediators](#)” on page 155 for details on how to configure dual-string mediators.
- If the disk set is configured with more than two disk strings, ensure that for any two disk strings S1 and S2, the sum of the number of drives on those strings exceeds the number of drives on the third string S3. Stated as a formula, the requirement is that  $\text{count}(S1) + \text{count}(S2) > \text{count}(S3)$ .

**10 Ensure that the local state database replicas exist.**

For instructions, see “How to Create State Database Replicas” on page 129.

**11 Become superuser on the cluster node that will master the disk set.****12 Create the disk set.**

The following command creates the disk set and registers the disk set as a Sun Cluster device group.

```
phys-schost# metaset -s setname -a -h node1 node2
```

|            |                                                                 |
|------------|-----------------------------------------------------------------|
| -s setname | Specifies the disk set name.                                    |
| -a         | Adds (creates) the disk set.                                    |
| -h node1   | Specifies the name of the primary node to master the disk set.  |
| node2      | Specifies the name of the secondary node to master the disk set |

---

**Note** – When you run the `metaset` command to configure a Solaris Volume Manager device group on a cluster, the command designates one secondary node by default. You can change the desired number of secondary nodes in the device group by using the `clsetup` utility after the device group is created. Refer to “Administering Device Groups” in *Sun Cluster System Administration Guide for Solaris OS* for more information about how to change the `numsecondaries` property.

---

**13 If you are configuring a replicated Solstice DiskSuite or Solaris Volume Manager device group, set the replication property for the device group.**

```
phys-schost# cldevicegroup sync device-group-name
```

For more information about data replication, see Chapter 4, “Data Replication Approaches,” in *Sun Cluster System Administration Guide for Solaris OS*.

**14 Verify the status of the new disk set.**

```
phys-schost# metaset -s setname
```

**15 As needed, set device group properties.**

```
phys-schost# cldevicegroup set -p name=value devicegroup
```

|             |                                                                                                 |
|-------------|-------------------------------------------------------------------------------------------------|
| -p          | Specifies a device-group property.                                                              |
| name        | Specifies the name of a property.                                                               |
| value       | Specifies the value or setting of the property.                                                 |
| devicegroup | Specifies the name of the device group. The device-group name is the same as the disk-set name. |

See the `cldevicegroup(1CL)` for information about device-group properties.

**Example 4–6** Creating a Disk Set

The following command creates two disk sets, `dg-schost-1` and `dg-schost-2`, with the nodes `phys-schost-1` and `phys-schost-2` specified as the potential primaries.

```
phys-schost# metaset -s dg-schost-1 -a -h phys-schost-1 phys-schost-2
phys-schost# metaset -s dg-schost-2 -a -h phys-schost-1 phys-schost-2
```

**Next Steps** Add drives to the disk set. Go to [“Adding Drives to a Disk Set” on page 150](#).

## Adding Drives to a Disk Set

When you add a drive to a disk set, the volume management software repartitions the drive as follows so that the state database for the disk set can be placed on the drive.

- A small portion of each drive is reserved for use by Solaris Volume Manager software. In volume table of contents (VTOC) labeled devices, slice 7 is used. In Extensible Firmware Interface (EFI) labeled devices, slice 6 is used. The remainder of the space on each drive is placed into slice 0.
- Drives are repartitioned when they are added to the disk set only if the target slice is not configured correctly.
- Any existing data on the drives is lost by the repartitioning.
- If the target slice starts at cylinder 0, and the drive partition is large enough to contain a state database replica, the drive is not repartitioned.

### ▼ How to Add Drives to a Disk Set

**Before You Begin** Ensure that the disk set has been created. For instructions, see [“How to Create a Disk Set” on page 147](#).

**1** Become superuser.

**2** List the DID mappings.

```
phys-schost# cldevice show | grep Device
```

- Choose drives that are shared by the cluster nodes that will master or potentially master the disk set.
- Use the full DID device name, which has the form `/dev/did/rdisk/dN`, when you add a drive to a disk set.

In the following example, the entries for DID device `/dev/did/rdisk/d3` indicate that the drive is shared by `phys-schost-1` and `phys-schost-2`.

```
=== DID Device Instances ===
DID Device Name: /dev/did/rdisk/d1
```

```

Full Device Path: phys-schost-1:/dev/rdisk/c0t0d0
DID Device Name: /dev/did/rdisk/d2
Full Device Path: phys-schost-1:/dev/rdisk/c0t6d0
DID Device Name: /dev/did/rdisk/d3
Full Device Path: phys-schost-1:/dev/rdisk/c1t1d0
Full Device Path: phys-schost-2:/dev/rdisk/c1t1d0
...

```

### 3 Become owner of the disk set.

```
phys-schost# cldegivegroup switch -n node devicegroup
```

*-n node* Specifies the node to take ownership of the device group.

*devicegroup* Specifies the device group name, which is the same as the disk set name.

### 4 Add the drives to the disk set.

Use the full DID path name.

```
phys-schost# metaset -s setname -a /dev/did/rdisk/dN
```

*-s setname* Specifies the disk set name, which is the same as the device group name.

*-a* Adds the drive to the disk set.

---

**Note** – Do *not* use the lower-level device name (*cNtXdY*) when you add a drive to a disk set. Because the lower-level device name is a local name and not unique throughout the cluster, using this name might prevent the metaset from being able to switch over.

---

### 5 Verify the status of the disk set and drives.

```
phys-schost# metaset -s setname
```

## Example 4–7 Adding Drives to a Disk Set

The `metaset` command adds the drives `/dev/did/rdisk/d1` and `/dev/did/rdisk/d2` to the disk set `dg-schost-1`.

```
phys-schost# metaset -s dg-schost-1 -a /dev/did/rdisk/d1 /dev/did/rdisk/d2
```

**Next Steps** To repartition drives for use in volumes, go to [“How to Repartition Drives in a Disk Set”](#) on page 152.

Otherwise, go to [“How to Create an `md.tab` File”](#) on page 152 to define metadevices or volumes by using an `md.tab` file.

## ▼ How to Repartition Drives in a Disk Set

The `metaset(1M)` command repartitions drives in a disk set so that a small portion of each drive is reserved for use by Solaris Volume Manager software. In volume table of contents (VTOC) labeled devices, slice 7 is used. In Extensible Firmware Interface (EFI) labeled devices, slice 6 is used. The remainder of the space on each drive is placed into slice 0. To make more effective use of the drive, use this procedure to modify the disk layout. If you allocate space to VTOC slices 1 through 6 or EFI slices 1 through 5, you can use these slices when you set up Solaris Volume Manager volumes.

### 1 Become superuser.

### 2 Use the `format` command to change the disk partitioning for each drive in the disk set.

When you repartition a drive, you must meet the following conditions to prevent the `metaset(1M)` command from repartitioning the drive.

- Create slice 7 for VTOC or slice 6 for EFI starting at cylinder 0, large enough to hold a state database replica. See your Solaris Volume Manager administration guide to determine the size of a state database replica for your version of the volume-manager software.
- Set the `Flag` field in the target slice to `wu` (read-write, unmountable). Do not set it to read-only.
- Do not allow the target slice to overlap any other slice on the drive.

See the `format(1M)` man page for details.

**Next Steps** Define 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 Solaris Volume Manager volumes for the disk sets that you created.

---

**Note** – If you are using local volumes, ensure that local volume names are distinct from the device-ID names that are used to form disk sets. For example, if the device-ID name `/dev/did/dsk/d3` is used in a disk set, do not use the name `/dev/md/dsk/d3` for a local volume. This requirement does not apply to shared volumes, which use the naming convention `/dev/md/setname/{r}dsk/d#`.

---

### 1 Become superuser.



## 2 List the DID mappings for reference when you create your `md.tab` file.

Use the full DID device names in the `md.tab` file in place of the lower-level device names (`cN tXdY`). The DID device name takes the form `/dev/did/rdisk/dN`.

```
phys-schost# cldevice show | grep Device
```

```
=== DID Device Instances ===
DID Device Name: /dev/did/rdisk/d1
 Full Device Path: phys-schost-1:/dev/rdisk/c0t0d0
DID Device Name: /dev/did/rdisk/d2
 Full Device Path: phys-schost-1:/dev/rdisk/c0t6d0
DID Device Name: /dev/did/rdisk/d3
 Full Device Path: phys-schost-1:/dev/rdisk/c1t1d0
 Full Device Path: phys-schost-2:/dev/rdisk/c1t1d0
...
```

## 3 Create an `/etc/lvm/md.tab` file and edit it with your preferred text editor.

---

**Note** – If you have existing data on the drives that will be used for the submirrors, you must back up the data before volume setup. Then restore the data onto the mirror.

---

To avoid possible confusion between local volumes on different nodes in a cluster environment, use a naming scheme that makes each local volume name unique throughout the cluster. For example, for node 1 choose names from `d100` to `d199`. And for node 2 use `d200` to `d299`.

See your Solaris Volume Manager documentation and the `md.tab(4)` man page for details about how to create an `md.tab` file.

### Example 4–8 Sample `md.tab` File

The following sample `md.tab` file defines the disk set that is named `dg-schost-1`. The ordering of lines in the `md.tab` file is not important.

```
dg-schost-1/d0 -m dg-schost-1/d10 dg-schost-1/d20
 dg-schost-1/d10 1 1 /dev/did/rdisk/d1s0
 dg-schost-1/d20 1 1 /dev/did/rdisk/d2s0
```

The sample `md.tab` file is constructed as follows.

1. The first line defines the device `d0` as a mirror of volumes `d10` and `d20`. The `-m` signifies that this device is a mirror device.

```
dg-schost-1/d0 -m dg-schost-1/d0 dg-schost-1/d20
```

2. The second line defines volume `d10`, the first submirror of `d0`, as a one-way stripe.

```
dg-schost-1/d10 1 1 /dev/did/rdisk/d1s0
```

- The third line defines volume `d20`, the second submirror of `d0`, as a one-way stripe.

```
dg-schost-1/d20 1 1 /dev/did/rdisk/d2s0
```

**Next Steps** Activate the volumes that are defined in the `md.tab` files. Go to [“How to Activate Volumes”](#) on page 154.

## ▼ How to Activate Volumes

Perform this procedure to activate Solaris Volume Manager volumes that are defined in `md.tab` files.

- Become superuser.**
- Ensure that `md.tab` files are located in the `/etc/lvm` directory.**
- Ensure that you have ownership of the disk set on the node where the command will be executed.**
- Take ownership of the disk set.**

```
phys-schost# cldevicegroup switch -n node devicegroup
```

`-n node` Specifies the node that takes ownership.

`devicegroup` Specifies the disk set name.

- Activate the disk set's volumes, which are defined in the `md.tab` file.**

```
phys-schost# metainit -s setname -a
```

`-s setname` Specifies the disk set name.

`-a` Activates all volumes in the `md.tab` file.

- Repeat [Step 3](#) through [Step 5](#) for each disk set in the cluster.**

If necessary, run the `metainit(1M)` command from another node that has connectivity to the drives. This step is required for cluster-pair topologies, where the drives are not accessible by all nodes.

- Check the status of the volumes.**

```
phys-schost# metastat -s setname
```

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

- (Optional) Capture the disk partitioning information for future reference.**

```
phys-schost# prtvtoc /dev/rdisk/cNtXdYsZ > filename
```

Store the file in a location outside the cluster. If you make any disk configuration changes, run this command again to capture the changed configuration. If a disk fails and needs replacement, you can use this information to restore the disk partition configuration. For more information, see the `prtvtoc(1M)` man page.

### 9 (Optional) Make a backup of your cluster configuration.

An archived backup of your cluster configuration facilitates easier recovery of the your cluster configuration,

For more information, see “How to Back Up the Cluster Configuration” in *Sun Cluster System Administration Guide for Solaris OS*.

### Example 4–9 Activating Volumes in the `md.tab` File

In the following example, all volumes that are defined in the `md.tab` file for disk set `dg-schost-1` are activated.

```
phys-schost# metainit -s dg-schost-1 -a
```

**Next Steps** If your cluster contains disk sets that are configured with exactly two disk enclosures and two nodes, add dual-string mediators. Go to “[Configuring Dual-String Mediators](#)” on page 155.

Otherwise, go to “[How to Create Cluster File Systems](#)” on page 175 to create a cluster file system.

## Configuring Dual-String Mediators

This section provides information and procedures to configure dual-string mediator hosts. Dual-string mediators are required for all Solaris Volume Manager disk sets that are configured with exactly two disk strings and two cluster nodes. 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.

A *dual-string mediator*, or mediator host, is a cluster node that stores mediator data. Mediator data provides information about the location of other mediators and contains a commit count that is identical to the commit count that is 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.

A *disk string* consists of a disk enclosure, its physical drives, cables from the enclosure to the node or nodes, and the interface adapter cards.

The following table lists the tasks that you perform to configure dual-string mediator hosts.

TABLE 4-3 Task Map: Installing and Configuring Solaris Volume Manager Software

| Task                                     | Instructions                                                                                                                    |
|------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------|
| 1. Configure dual-string mediator hosts. | <a href="#">“Requirements for Dual-String Mediators” on page 156</a><br><a href="#">“How to Add Mediator Hosts” on page 156</a> |
| 2. Check the status of mediator data.    | <a href="#">“How to Check the Status of Mediator Data” on page 157</a>                                                          |
| 3. If necessary, fix bad mediator data.  | <a href="#">“How to Fix Bad Mediator Data” on page 157</a>                                                                      |

## Requirements for Dual-String Mediators

The following rules apply to dual-string configurations that use mediators.

- Disk sets must be configured with exactly two mediator hosts. Those two mediator hosts must be the same two cluster nodes that are used for the disk set.
- A disk set cannot have more than two mediator hosts.
- Mediators cannot be configured for disk sets that do not meet the two-string and two-host criteria.

These rules do not require that the entire cluster must have exactly two nodes. Rather, only those disk sets that have two disk strings must be connected to exactly two nodes. An N+1 cluster and many other topologies are permitted under these rules.

### ▼ How to Add Mediator Hosts

Perform this procedure if your configuration requires dual-string mediators.

- 1 **Become superuser on the node that currently masters the disk set to which you intend to add mediator hosts.**
- 2 **Add each node with connectivity to the disk set as a mediator host for that disk set.**

```
phys-schost# metaset -s setname -a -m mediator-host-list
```

- s setname                      Specifies the disk set name.
- a                                Adds to the disk set.
- m mediator-host-list        Specifies the name of the node to add as a mediator host for the disk set.

See the `mediator(7D)` man page for details about mediator-specific options to the `metaset` command.

**Example 4–10** Adding Mediator Hosts

The following example adds the nodes `phys-schost-1` and `phys-schost-2` as mediator hosts for the disk set `dg-schost-1`. Both commands are run from the node `phys-schost-1`.

```
phys-schost# metaset -s dg-schost-1 -a -m phys-schost-1
phys-schost# metaset -s dg-schost-1 -a -m phys-schost-2
```

**Next Steps** Check the status of mediator data. Go to [“How to Check the Status of Mediator Data”](#) on page 157.

## ▼ How to Check the Status of Mediator Data

**Before You Begin** Ensure that you have added mediator hosts as described in [“How to Add Mediator Hosts”](#) on page 156.

**1** Display the status of the mediator data.

```
phys-schost# medstat -s setname
```

`-s setname` Specifies the disk set name.

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

**2** If `Bad` is the value in the `Status` field of the `medstat` output, repair the affected mediator host.

Go to [“How to Fix Bad Mediator Data”](#) on page 157.

**Next Steps** Go to [“How to Create Cluster File Systems”](#) on page 175 to create a cluster file system.

## ▼ How to Fix Bad Mediator Data

Perform this procedure to repair bad mediator data.

**1** Identify all mediator hosts with bad mediator data as described in the procedure [“How to Check the Status of Mediator Data”](#) on page 157.

**2** Become superuser on the node that owns the affected disk set.

**3** Remove all mediator hosts with bad mediator data from all affected disk sets.

```
phys-schost# metaset -s setname -d -m mediator-host-list
```

`-s setname` Specifies the disk set name.

`-d` Deletes from the disk set.

`-m mediator-host-list` Specifies the name of the node to remove as a mediator host for the disk set.

#### 4 Restore each mediator host that you removed in [Step 3](#).

```
phys-schost# metaset -s setname -a -m mediator-host-list
```

`-a` Adds to the disk set.

`-m mediator-host-list` Specifies the name of the node to add as a mediator host for the disk set.

See the `mediator(7D)` man page for details about mediator-specific options to the `metaset` command.

**Next Steps** Determine from the following list the next task to perform that applies to your cluster configuration. If you need to perform more than one task from this list, go to the first of those tasks in this list.

- To create cluster file systems, go to [“How to Create Cluster File Systems”](#) on page 175.
- To create non-global zones on a node, go to [“How to Create a Non-Global Zone on a Cluster Node”](#) on page 179.
- SPARC: To configure Sun Management Center to monitor the cluster, go to [“SPARC: Installing the Sun Cluster Module for Sun Management Center”](#) on page 183.
- Install third-party applications, register resource types, set up resource groups, and configure data services. See the documentation that is supplied with the application software and the *Sun Cluster Data Services Planning and Administration Guide for Solaris OS*.

# Installing and Configuring VERITAS Volume Manager

---

Install and configure your local and multihost disks for VERITAS Volume Manager (VxVM) by using the procedures in this chapter, along with the planning information in [“Planning Volume Management” on page 36](#). See your VxVM documentation for additional details.

The following sections are in this chapter:

- [“Installing and Configuring VxVM Software” on page 159](#)
- [“Creating Disk Groups in a Cluster” on page 166](#)
- [“Unencapsulating the Root Disk” on page 172](#)

## Installing and Configuring VxVM Software

This section provides information and procedures to install and configure VxVM software on a Sun Cluster configuration.

The following table lists the tasks to perform to install and configure VxVM software for Sun Cluster configurations.

**TABLE 5-1** Task Map: Installing and Configuring VxVM Software

| Task                                                                                 | Instructions                                                                                                    |
|--------------------------------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------|
| 1. Plan the layout of your VxVM configuration.                                       | <a href="#">“Planning Volume Management” on page 36</a>                                                         |
| 2. <i>(Optional)</i> Determine how you will create the root disk group on each node. | <a href="#">“Setting Up a Root Disk Group Overview” on page 160</a>                                             |
| 3. Install VxVM software.                                                            | <a href="#">“How to Install VERITAS Volume Manager Software” on page 160</a><br>VxVM installation documentation |

TABLE 5-1 Task Map: Installing and Configuring VxVM Software (Continued)

| Task                                                                                                                                    | Instructions                                                                                                                        |
|-----------------------------------------------------------------------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------|
| 4. (Optional) Create a root disk group. You can either encapsulate the root disk or create the root disk group on local, nonroot disks. | <p>“SPARC: How to Encapsulate the Root Disk” on page 162</p> <p>“How to Create a Root Disk Group on a Nonroot Disk” on page 163</p> |
| 5. (Optional) Mirror the encapsulated root disk.                                                                                        | “How to Mirror the Encapsulated Root Disk” on page 164                                                                              |
| 6. Create disk groups.                                                                                                                  | “Creating Disk Groups in a Cluster” on page 166                                                                                     |

## Setting Up a Root Disk Group Overview

The creation of a root disk group is optional. If you do not intend to create a root disk group, proceed to “How to Install VERITAS Volume Manager Software” on page 160.

- Access to a node’s root disk group must be restricted to only that node.
- Remote nodes must never access data stored in another node’s root disk group.
- Do not use the `cldevicegroup` command to register the root disk group as a device group.
- Whenever possible, configure the root disk group for each node on a nonshared disk.

Sun Cluster software supports the following methods to configure the root disk group.

- **Encapsulate the node’s root disk** – This method enables the root disk to be mirrored, which provides a boot alternative if the root disk is corrupted or damaged. To encapsulate the root disk you need two free disk slices as well as free cylinders, preferably at the beginning or the end of the disk.
- **Use local nonroot disks** – This method provides an alternative to encapsulating the root disk. If a node’s root disk is encapsulated, certain tasks you might later perform, such as upgrade the Solaris OS or perform disaster recovery procedures, could be more complicated than if the root disk is not encapsulated. To avoid this potential added complexity, you can instead initialize or encapsulate local nonroot disks for use as root disk groups.

A root disk group that is created on local nonroot disks is local to that node, neither globally accessible nor highly available. As with the root disk, to encapsulate a nonroot disk you need two free disk slices as well as free cylinders at the beginning or the end of the disk.

See your VxVM installation documentation for more information.

## ▼ How to Install VERITAS Volume Manager Software

Perform this procedure to install VERITAS Volume Manager (VxVM) software on each node that you want to install with VxVM. You can install VxVM on all nodes of the cluster, or install VxVM just on the nodes that are physically connected to the storage devices that VxVM will manage.



**Before You Begin** Perform the following tasks:

- Ensure that all nodes in the cluster are running in cluster mode.
- Obtain any VERITAS Volume Manager (VxVM) license keys that you need to install.
- Have available your VxVM installation documentation.

**1 Become superuser on a cluster node that you intend to install with VxVM.**

**2 Insert the VxVM CD-ROM in the CD-ROM drive on the node.**

**3 Follow procedures in your VxVM installation guide to install and configure VxVM software and licenses.**

**4 Run the `clvxxvm` utility in noninteractive mode.**

```
phys-schost# clvxxvm initialize
```

The `clvxxvm` utility performs necessary postinstallation tasks. The `clvxxvm` utility also selects and configures a cluster-wide `vxio` driver major number. See the `clvxxvm(1CL)` man page for more information.

**5 SPARC: To enable the VxVM cluster feature, supply the cluster feature license key, if you did not already do so.**

See your VxVM documentation for information about how to add a license.

**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 to support Sun Cluster software.**

See “Patches and Required Firmware Levels” in *Sun Cluster 3.2 Release Notes for Solaris OS* for the location of patches and installation instructions.

**9 Repeat [Step 1](#) through [Step 8](#) to install VxVM on any additional nodes.**

---

**Note** – SPARC: To enable the VxVM cluster feature, you *must* install VxVM on all nodes of the cluster.

---

**10 If you do *not* install one or more nodes with VxVM, modify the `/etc/name_to_major` file on each non-VxVM node.**

**a. On a node that is installed with VxVM, determine the `vxio` major number setting.**

```
phys-schost# 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](#).**

```
phys-schost# vi /etc/name_to_major
vxio NNN
```

- d. **Initialize the `vxio` entry.**

```
phys-schost# drvconfig -b -i vxio -m NNN
```

- e. **Repeat [Step a](#) 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.

- 11 **To create a root disk group, go to [“SPARC: How to Encapsulate the Root Disk”](#) on page 162 or [“How to Create a Root Disk Group on a Nonroot Disk”](#) on page 163.**

Otherwise, proceed to [Step 12](#).

---

**Note** – A root disk group is optional.

---

- 12 **Reboot each node on which you installed VxVM.**

```
phys-schost# shutdown -g0 -y -i6
```

**Next Steps** To create a root disk group, go to [“SPARC: How to Encapsulate the Root Disk”](#) on page 162 or [“How to Create a Root Disk Group on a Nonroot Disk”](#) on page 163.

Otherwise, create disk groups. Go to [“Creating Disk Groups in a Cluster”](#) on page 166.

## ▼ **SPARC: How to Encapsulate the Root Disk**

Perform this procedure to create a root disk group by encapsulating the root disk. Root disk groups are optional. See your VxVM documentation for more information.

---

**Note** – If you want to create a root disk group on nonroot disks, instead perform procedures in [“How to Create a Root Disk Group on a Nonroot Disk”](#) on page 163.

---

**Before You Begin** Ensure that you have installed VxVM as described in [“How to Install VERITAS Volume Manager Software”](#) on page 160.

- 1 **Become superuser on a node that you installed with VxVM.**

- 2 **Encapsulate the root disk.**

```
phys-schost# clvxdm encapsulate
```

See the `clvxdm(1CL)` man page for more information.

### 3 Repeat for any other node on which you installed VxVM.

**Next Steps** To mirror the encapsulated root disk, go to “[How to Mirror the Encapsulated Root Disk](#)” on page 164.

Otherwise, go to “[Creating Disk Groups in a Cluster](#)” on page 166.

## ▼ How to Create a Root Disk Group on a Nonroot Disk

Use this procedure to create a root disk group by encapsulating or initializing local disks other than the root disk. The creation of a root disk group is optional.

---

**Note** – If you want to create a root disk group on the root disk, instead perform procedures in “[SPARC: How to Encapsulate the Root Disk](#)” on page 162.

---

**Before You Begin** If the disks are to be encapsulated, ensure that each disk has at least two slices with 0 cylinders. If necessary, use the `format(1M)` command to assign 0 cylinders to each VxVM slice.

### 1 Become superuser.

### 2 Start the `vxinstall` utility.

```
phys-schost# vxinstall
```

### 3 When prompted by the `vxinstall` utility, make the following choices or entries.

- SPARC: To enable the VxVM cluster feature, supply the cluster feature license key.
- Choose Custom Installation.
- Do not encapsulate the boot disk.
- Choose any disks to add to the root disk group.
- Do not accept automatic reboot.

### 4 If the root disk group that you created contains one or more disks that connect to more than one node, enable the `localonly` property.

Use the following command to enable the `localonly` property of the raw-disk device group for each shared disk in the root disk group.

```
phys-schost# cldevicegroup set -p localonly=true disk/dN
```

-p Specifies a property of the device group.

localonly=true Sets the device group to be mastered only by the single node in the node list.

When the `localonly` property is enabled, the raw-disk device group is used exclusively by the node in its node list. This usage prevents unintentional fencing of the node from the disk that is used by the root disk group if that disk is connected to multiple nodes.

For more information about the `localonly` property, see the `scconf_dg_rawdisk(1M)` man page.

#### 5 Evacuate any resource groups or device groups from the node.

```
phys-schost# clnode evacuate from-node
```

*from-node* Specifies the name of the node from which to move resource or device groups.

#### 6 Reboot the node.

```
phys-schost# shutdown -g0 -y -i6
```

#### 7 Use the `vxdiskadm` command to add multiple disks to the root disk group.

The root disk group becomes tolerant of a disk failure when it contains multiple disks. See VxVM documentation for procedures.

**Next Steps** Create disk groups. Go to “[Creating Disk Groups in a Cluster](#)” on page 166.

## ▼ 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.

**Before You Begin** Ensure that you have encapsulated the root disk as described in “[SPARC: How to Encapsulate the Root Disk](#)” on page 162.

#### 1 Become superuser.

#### 2 List the devices.

```
phys-schost# cldevice list -v
```

Output looks similar to the following:

| DID Device | Full Device Path                |
|------------|---------------------------------|
| -----      | -----                           |
| d1         | phys-schost-1:/dev/rdisk/c0t0d0 |
| d2         | phys-schost-1:/dev/rdisk/c0t6d0 |
| d3         | phys-schost-2:/dev/rdisk/c1t1d0 |
| d3         | phys-schost-1:/dev/rdisk/c1t1d0 |

**3 Mirror the encapsulated root disk.**

Follow the procedures in your VxVM documentation.

For maximum availability and simplified administration, use a local disk for the mirror. See “Guidelines for Mirroring the Root Disk” on page 42 for additional guidelines.




---

**Caution** – Do not use a quorum device to mirror a root disk. Using a quorum device to mirror a root disk might prevent the node from booting from the root-disk mirror under certain circumstances.

---

**4 View the node list of the raw-disk device group for the device that you used to mirror the root disk.**

The name of the device group is the form `dsk/dN`, where `dN` is the DID device name.

```
phys-schost# cldevicegroup list -v dsk/dN
```

`-v` Displays verbose output.

Output looks similar to the following.

| Device group | Type       | Node list                    |
|--------------|------------|------------------------------|
| -----        | ----       | -----                        |
| dsk/dN       | Local_Disk | phys-schost-1, phys-schost-3 |

**5 If the node list contains more than one node name, remove from the node list all nodes except the node whose root disk you mirrored.**

Only the node whose root disk you mirrored should remain in the node list for the raw-disk device group.

```
phys-schost# cldevicegroup remove-node -n node dsk/dN
```

`-n node` Specifies the node to remove from the device-group node list.

**6 Enable the `localonly` property of the raw-disk device group, if it is not already enabled.**

When the `localonly` property is enabled, the raw-disk device group is used exclusively by the node in its node list. This usage prevents unintentional fencing of the node from its boot device if the boot device is connected to multiple nodes.

```
phys-schost# cldevicegroup set -p localonly=true dsk/dN
```

`-p` Sets the value of a device-group property.

`localonly=true` Enables the `localonly` property of the device group.

For more information about the `localonly` property, see the `scconf_dg_rawdisk(1M)` man page.

**7 Repeat this procedure for each node in the cluster whose encapsulated root disk you want to mirror.**

**Example 5-1** 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 `c0t0d0`, whose raw-disk device-group name is `dsk/d2`. Disk `c1t1d0` is a multihost disk, so the node `phys-schost-3` is removed from the disk's node list and the `localonly` property is enabled.

```
phys-schost# cldevice list -v
DID Device Full Device Path

d2 pcircinus1:/dev/rdisk/c0t0d0
...
 Create the mirror by using VxVM procedures
phys-schost# cldevicegroup list -v dsk/d2
Device group Type Node list

dsk/d2 Local_Disk phys-schost-1, phys-schost-3
phys-schost# cldevicegroup remove-node -n phys-schost-3 dsk/d2
phys-schost# cldevicegroup set -p localonly=true dsk/d2
```

**Next Steps** Create disk groups. Go to [“Creating Disk Groups in a Cluster”](#) on page 166.

## Creating Disk Groups in a Cluster

This section describes how to create VxVM disk groups in a cluster. The following table describes the types of VxVM disk groups you can configure in a Sun Cluster configuration and their characteristics.

| Disk Group Type        | Use                                                                                           | Registered with Sun Cluster? | Storage Requirement        |
|------------------------|-----------------------------------------------------------------------------------------------|------------------------------|----------------------------|
| VxVM disk group        | Device groups for failover or scalable data services, global devices, or cluster file systems | Yes                          | Shared storage             |
| Local VxVM disk group  | Applications that are not highly available and are confined to a single node                  | No                           | Shared or unshared storage |
| VxVM shared disk group | Oracle RAC (also requires the VxVM cluster feature)                                           | No                           | Shared storage             |

The following table lists the tasks to perform to create VxVM disk groups in a Sun Cluster configuration.

TABLE 5-2 Task Map: Creating VxVM Disk Groups

| Task                                                                                                                        | Instructions                                                     |
|-----------------------------------------------------------------------------------------------------------------------------|------------------------------------------------------------------|
| 1. Create disk groups and volumes.                                                                                          | “How to Create a Disk Group” on page 167                         |
| 2. Register as Sun Cluster device groups those disk groups that are not local and that do not use the VxVM cluster feature. | “How to Register a Disk Group” on page 168                       |
| 3. If necessary, resolve any minor-number conflicts between device groups by assigning a new minor number.                  | “How to Assign a New Minor Number to a Device Group” on page 169 |
| 4. Verify the disk groups and volumes.                                                                                      | “How to Verify the Disk Group Configuration” on page 170         |

## ▼ How to Create a Disk Group

Use this procedure to create your VxVM disk groups and volumes.

Perform this procedure from a node that is physically connected to the disks that make the disk group that you add.

**Before You Begin** Perform the following tasks:

- Make mappings of your storage disk drives. See the appropriate manual in the *Sun Cluster Hardware Administration Collection* to perform an initial installation of your storage device.
  - Complete the following configuration planning worksheets.
    - “Local File System Layout Worksheet” on page 275
    - “Device Group Configurations Worksheet” on page 279
    - “Volume-Manager Configurations Worksheet” on page 281
- See “Planning Volume Management” on page 36 for planning guidelines.
- If you did not create root disk groups, ensure that you have rebooted each node on which you installed VxVM, as instructed in Step 12 of “How to Install VERITAS Volume Manager Software” on page 160.

**1 Become superuser on the node that will own the disk group.**

**2 Create the VxVM disk groups and volumes.**

Observe the following special instructions:

- SPARC: If you are installing Oracle RAC, create shared VxVM disk groups by using the cluster feature of VxVM. Observe guidelines and instructions in “How to Create a VxVM Shared-Disk Group for the Oracle RAC Database” in *Sun Cluster Data Service for Oracle RAC Guide for Solaris OS* and in the *VERITAS Volume Manager Administrator’s Reference Guide*.

- Otherwise, create VxVM disk groups by using the standard procedures that are documented in the VxVM documentation.

---

**Note** – You can use Dirty Region Logging (DRL) to decrease volume recovery time if a node failure occurs. However, DRL might decrease I/O throughput.

---

### 3 For local disk groups, set the `localonly` property and add a single node to the disk group's node list.

---

**Note** – A disk group that is configured to be local only is not highly available or globally accessible.

---

**a. Start the `clsetup` utility.**

```
phys-schost# clsetup
```

**b. Choose the menu item, Device groups and volumes.**

**c. Choose the menu item, Set localonly on a VxVM disk group.**

**d. Follow the instructions to set the `localonly` property and to specify the single node that will exclusively master the disk group.**

Only one node at any time is permitted to master the disk group. You can later change which node is the configured master.

**e. When finished, quit the `clsetup` utility.**

**Next Steps** Determine your next step:

- SPARC: If the VxVM cluster feature is enabled, go to [“How to Verify the Disk Group Configuration” on page 170.](#)
- If you created disk groups that are not local and the VxVM cluster feature is *not* enabled, register the disk groups as Sun Cluster device groups. Go to [“How to Register a Disk Group” on page 168.](#)
- If you created only local disk groups, go to [“How to Verify the Disk Group Configuration” on page 170.](#)

## ▼ How to Register a Disk Group

If the VxVM cluster feature is *not* enabled, perform this procedure to register disk groups that are not local as Sun Cluster device groups.



---

**Note** – SPARC: If the VxVM cluster feature is enabled or you created a local disk group, do not perform this procedure. Instead, proceed to [“How to Verify the Disk Group Configuration”](#) on page 170.

---

- 1 **Become superuser on a node of the cluster.**
- 2 **Register the global disk group as a Sun Cluster device group.**
  - a. **Start the `clsetup` utility.**  
`phys-schost# clsetup`
  - b. **Choose the menu item, Device groups and volumes.**
  - c. **Choose the menu item, Register a VxVM disk group.**
  - d. **Follow the instructions to specify the VxVM disk group that you want to register as a Sun Cluster device group.**
  - e. **When finished, quit the `clsetup` utility.**
- 3 **Verify that the device group is registered.**  
 Look for the disk device information for the new disk that is displayed by the following command.  
`phys-schost# cldevicegroup status`

**Next Steps** Go to [“How to Verify the Disk Group Configuration”](#) on page 170.

**Troubleshooting** **Stack overflow** – If a stack overflows when the device group is brought online, the default value of the thread stack size might be insufficient. On each node, add the entry `set cl_comm:rm_thread_stacksize=0xsize` to the `/etc/system` file, where `size` is a number greater than 8000, which is the default setting.

**Configuration changes** – If you change any configuration information for a VxVM device group or its volumes, you must register the configuration changes by using the `clsetup` utility. Configuration changes that you must register include adding or removing volumes and changing the group, owner, or permissions of existing volumes. See [“Administering Device Groups”](#) in *Sun Cluster System Administration Guide for Solaris OS* for procedures to register configuration changes that are made to a VxVM device group.

## ▼ How to Assign a New Minor Number to a Device Group

If 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.**  

```
phys-schost# 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.**  

```
phys-schost# vxkg remminor diskgroup base-minor-number
```

### Example 5-2 How to Assign a New Minor Number to a Device Group

This example uses the minor numbers 16000-16002 and 4000-4001. The vxkg remminor command reminars the new device group to use the base minor number 5000.

```
phys-schost# 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
phys-schost# vxkg remminor dg3 5000
```

**Next Steps** Register the disk group as a Sun Cluster device group. Go to [“How to Register a Disk Group”](#) on page 168.

## ▼ How to Verify the Disk Group Configuration

Perform this procedure on each node of the cluster.

- 1 **Become superuser.**
- 2 **List the disk groups.**  

```
phys-schost# vxkg list
```
- 3 **List the device groups.**  

```
phys-schost# cldevicegroup list -v
```
- 4 **Verify that all disk groups are correctly configured.**  
 Ensure that the following requirements are met:

- The root disk group includes only local disks.
- All disk groups and any local disk groups are imported on the current primary node only.

#### 5 Verify that all volumes have been started.

```
phys-schost# vxprint
```

#### 6 Verify that all disk groups have been registered as Sun Cluster device groups and are online.

```
phys-schost# cldevicegroup status
```

Output should not display any local disk groups.

#### 7 (Optional) Capture the disk partitioning information for future reference.

```
phys-schost# prtvtoc /dev/rdisk/cNtXdYsZ > filename
```

Store the file in a location outside the cluster. If you make any disk configuration changes, run this command again to capture the changed configuration. If a disk fails and needs replacement, you can use this information to restore the disk partition configuration. For more information, see the `prtvtoc(1M)` man page.

#### 8 (Optional) Make a backup of your cluster configuration.

An archived backup of your cluster configuration facilitates easier recovery of the your cluster configuration,

For more information, see “How to Back Up the Cluster Configuration” in *Sun Cluster System Administration Guide for Solaris OS*.

### More Information

## Guidelines for Administering VxVM Disk Groups

Observe the following guidelines for administering VxVM disk groups in a Sun Cluster configuration:

- **VxVM device groups** - VxVM disk groups that have been registered as device groups are managed by Sun Cluster software. After a disk group is registered as a device group, you should never import or deport that VxVM disk group by using VxVM commands. The Sun Cluster software can handle all cases where device groups need to be imported or deported. See “Administering Device Groups” in *Sun Cluster System Administration Guide for Solaris OS* for procedures about how to manage device groups.
- **Local disk groups** - Local VxVM disk groups are not managed by Sun Cluster software. Use VxVM commands to administer local disk groups as you would in a nonclustered system.

### Troubleshooting

If the output of the `cldevicegroup status` command includes any local disk groups, the displayed disk groups are not configured correctly for local-only access. Return to “[How to Create a Disk Group](#)” on page 167 to reconfigure the local disk group.

### Next Steps

Determine from the following list the next task to perform that applies to your cluster configuration. If you need to perform more than one task from this list, go to the first of those tasks in this list.

- To create cluster file systems, go to “[How to Create Cluster File Systems](#)” on page 175.
- To create non-global zones on a node, go to “[How to Create a Non-Global Zone on a Cluster Node](#)” on page 179.
- SPARC: To configure Sun Management Center to monitor the cluster, go to “[SPARC: Installing the Sun Cluster Module for Sun Management Center](#)” on page 183.
- Install third-party applications, register resource types, set up resource groups, and configure data services. See the documentation that is supplied with the application software and the *Sun Cluster Data Services Planning and Administration Guide for Solaris OS*.

## Unencapsulating the Root Disk

This section describes how to unencapsulate the root disk in a Sun Cluster configuration.

### ▼ How to Unencapsulate the Root Disk

Perform this procedure to unencapsulate the root disk.

**Before You Begin** Perform the following tasks:

- Ensure that only Solaris root file systems are present on the root disk. The Solaris root file systems are root (/), swap, the global devices namespace, /usr, /var, /opt, and /home.
- Back up and remove from the root disk any file systems other than Solaris root file systems that reside on the root disk.

**1 Become superuser on the node that you intend to unencapsulate.**

**2 Evacuate all resource groups and device groups from the node.**

```
phys-schost# clnode evacuate from-node
```

*from-node* Specifies the name of the node from which to move resource or device groups.

**3 Determine the node-ID number of the node.**

```
phys-schost# clinfo -n
```

**4 Unmount the global-devices file system for this node, where *N* is the node ID number that is returned in [Step 3](#).**

```
phys-schost# umount /global/.devices/node@N
```

**5 View the /etc/vfstab file and determine which VxVM volume corresponds to the global-devices file system.**

```
phys-schost# 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
```

## 6 Remove from the root disk group the VxVM volume that corresponds to the global-devices file system.

```
phys-schost# vxedit -g rootdiskgroup -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 that is related to global devices entries is restored after the root disk is unencapsulated.

---

## 7 Unencapsulate the root disk.

---

**Note** – Do **not** accept the shutdown request from the command.

---

```
phys-schost# /etc/vx/bin/vxunroot
```

See your VxVM documentation for details.

## 8 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.

---

## 9 Set up a file system on the partition that you created in [Step 8](#).

```
phys-schost# newfs /dev/rdisk/cNtXdYsZ
```

## 10 Determine the DID name of the root disk.

```
phys-schost# cldevice list cNtXdY
dN
```

## 11 In the `/etc/vfstab` file, replace the path names in the global-devices file system entry with the DID path that you identified in [Step 10](#).

The original entry would look similar to the following.

```
phys-schost# 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
```

**12 Mount the global-devices file system.**

```
phys-schost# mount /global/.devices/node@N
```

**13 From one node of the cluster, repopulate the global-devices file system with device nodes for any raw-disk devices and Solaris Volume Manager devices.**

```
phys-schost# cldevice populate
```

VxVM devices are recreated during the next reboot.

**14 On each node, verify that the cldevice populate command has completed processing before you proceed to the next step.**

The cldevice populate command executes remotely on all nodes, even though the command is issued from just one node. To determine whether the cldevice populate command has completed processing, run the following command on each node of the cluster.

```
phys-schost# ps -ef | grep scgdevs
```

**15 Reboot the node.**

```
phys-schost# shutdown -g0 -y -i6
```

**16 Repeat this procedure on each node of the cluster to unencapsulate the root disk on those nodes.**

# Creating Cluster File Systems and Non-Global Zones

---

This chapter provides the following procedures:

- “How to Create Cluster File Systems” on page 175
- “How to Create a Non-Global Zone on a Cluster Node” on page 179

## Creating Cluster File Systems

This section provides procedures to create cluster file systems to support data services.

### ▼ How to Create Cluster File Systems

Perform this procedure for each cluster file system that you want to create. Unlike a local file system, a cluster file system is accessible from any node in the cluster.

---

**Note** – Alternatively, you can use a highly available local file system to support a data service. For information about choosing between creating a cluster file system or a highly available local file system to support a particular data service, see the manual for that data service. For general information about creating a highly available local file system, see “Enabling Highly Available Local File Systems” in *Sun Cluster Data Services Planning and Administration Guide for Solaris OS*.

---

**Before You Begin** Perform the following tasks:

- Ensure that you installed software packages for the Solaris OS, Sun Cluster framework, and other products as described in “Installing the Software” on page 45.
- Ensure that you established the new cluster or cluster node as described in “Establishing a New Cluster or New Cluster Node” on page 63.
- If you are using a volume manager, ensure that volume-management software is installed and configured. For volume-manager installation procedures, see “Configuring Solaris Volume Manager Software” on page 127 or “Installing and Configuring VxVM Software” on page 159.

**Note** – If you added a new node to a cluster that uses VxVM, you must do one of the following tasks:

- Install VxVM on that node.
- Modify that node’s `/etc/name_to_major` file to support coexistence with VxVM.

Follow procedures in “[How to Install VERITAS Volume Manager Software](#)” on page 160 to perform one of these required tasks.

- Determine the mount options to use for each cluster file system that you want to create. See “[Choosing Mount Options for Cluster File Systems](#)” on page 33.

**1 Become superuser on any node in the cluster.**

For Solaris, you must perform this procedure from the global zone if non-global zones are configured in the cluster.

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

**2 Create a file system.**



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

- **For a UFS file system, use the `newfs(1M)` command.**

`phys-schost# newfs raw-disk-device`

The following table shows examples of names for the `raw-disk-device` argument. Note that naming conventions differ for each volume manager.

| Volume Manager         | Sample Disk Device Name               | Description                                       |
|------------------------|---------------------------------------|---------------------------------------------------|
| Solaris Volume Manager | <code>/dev/md/nfs/rdisk/d1</code>     | Raw disk device d1 within the nfs disk set        |
| VERITAS Volume Manager | <code>/dev/vx/rdsk/oradg/vol01</code> | Raw disk device vol01 within the oradg disk group |
| None                   | <code>/dev/global/rdsk/d1s3</code>    | Raw disk device d1s3                              |

- **For a Sun StorEdge QFS file system, follow the procedures for defining the configuration in the *Sun StorEdge QFS Installation and Upgrade Guide*.**
- **SPARC: For a VERITAS File System (VxFS) file system, follow the procedures that are provided in your VxFS documentation.**



### 3 On each node in the cluster, create a mount-point directory for the cluster file system.

A mount point is required *on each node*, even if the cluster file system is not accessed on that node.

---

**Tip** – For ease of administration, create the mount point in the `/global/device-group/` directory. This location enables you to easily distinguish cluster file systems, which are globally available, from local file systems.

---

```
phys-schost# 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.

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

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

---

**Note** – If non-global zones are configured in the cluster, ensure that you mount cluster file systems in the global zone on a path in the global zone's root directory.

---

#### a. In each entry, specify the required mount options for the type of file system that you use.

---

**Note** – Do *not* use the logging mount option for Solaris Volume Manager transactional volumes. Transactional volumes provide their own logging.

In addition, Solaris Volume Manager transactional-volume logging is removed from the Solaris 10 OS. Solaris UFS logging provides the same capabilities but superior performance, as well as lower system administration requirements and overhead.

---

#### 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/`.

### 5 On any node in the cluster, run the configuration check utility.

```
phys-schost# scheck
```

The configuration check utility verifies that the mount points exist. The utility also verifies that `/etc/vfstab` file entries are correct on all nodes of the cluster. If no errors occur, nothing is returned.

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

## 6 Mount the cluster file system.

```
phys-schost# mount /global/device-group/mountpoint/
```

- For UFS and QFS, mount the cluster file system from any node in the cluster.
- SPARC: For VxFS, mount the cluster file system from the current master of *device-group* to ensure that the file system mounts successfully.

In addition, unmount a VxFS file system from the current master of *device-group* to ensure that the file system unmounts successfully.

---

**Note** – To manage a VxFS cluster file system in a Sun Cluster environment, run administrative commands only from the primary node on which the VxFS cluster file system is mounted.

---

## 7 On each node of the cluster, verify that the cluster file system is mounted.

You can use either the `df` command or `mount` command to list mounted file systems. For more information, see the `df(1M)` man page or `mount(1M)` man page.

For the Solaris 10 OS, cluster file systems are accessible from both the global zone and the non-global zone.

### Example 6–1 Creating a Cluster File System

The following example creates a UFS cluster file system on the Solaris Volume Manager volume `/dev/md/oracle/rdisk/d1`. An entry for the cluster file system is added to the `vfstab` file on each node. Then from one node the `sccheck` command is run. After configuration check processing is completed successfully, the cluster file system is mounted from one node and verified on all nodes.

```
phys-schost# newfs /dev/md/oracle/rdisk/d1
...
phys-schost# mkdir -p /global/oracle/d1
phys-schost# vi /etc/vfstab
#device device mount FS fsck mount mount
#to mount to fsck point type pass at boot options
#
/dev/md/oracle/dsk/d1 /dev/md/oracle/rdisk/d1 /global/oracle/d1 ufs 2 yes global,logging
...
phys-schost# sccheck
phys-schost# mount /global/oracle/d1
phys-schost# mount
...
```

```
/global/oracle/d1 on /dev/md/oracle/dsk/d1 read/write/setuid/global/logging/largefiles
on Sun Oct 3 08:56:16 2005
```

**Next Steps** Determine from the following list the next task to perform that applies to your cluster configuration. If you need to perform more than one task from this list, go to the first of those tasks in this list.

- To create non-global zones on a node, go to [“How to Create a Non-Global Zone on a Cluster Node” on page 179](#).
- SPARC: To configure Sun Management Center to monitor the cluster, go to [“SPARC: Installing the Sun Cluster Module for Sun Management Center” on page 183](#).
- Install third-party applications, register resource types, set up resource groups, and configure data services. See the documentation that is supplied with the application software and the *Sun Cluster Data Services Planning and Administration Guide for Solaris OS*.

## Configuring a Non-Global Zone on a Cluster Node

This section provides procedures to create a non-global zone on a cluster node.

### ▼ How to Create a Non-Global Zone on a Cluster Node

Perform this procedure for each non-global zone that you create in the cluster.

---

**Note** – For complete information about installing a zone, refer to *System Administration Guide: Solaris Containers-Resource Management and Solaris Zones*.

---

You can configure a Solaris 10 non-global zone, simply referred to as a zone, on a cluster node while the node is booted in either cluster mode or in noncluster mode.

- If you create a zone while the node is booted in noncluster mode, the cluster software discovers the zone when the node joins the cluster.
- If you create or remove a zone while the node is in cluster mode, the cluster software dynamically changes its list of zones that can master resource groups.

**Before You Begin** Perform the following tasks:

- Plan your non-global zone configuration. Observe the requirements and restrictions in [“Guidelines for Non-Global Zones in a Cluster” on page 19](#).
- Have available the following information:
  - The total number of non-global zones that you will create.
  - The public adapter and public IP address that each zone will use.

- The zone path for each zone. This path must be a local file system, not a cluster file system or a highly available local file system.
- One or more devices that should appear in each zone.
- (Optional) The name that you will assign each zone.
- If you will assign the zone a private IP address, ensure that the cluster IP address range can support the additional private IP addresses that you will configure. Use the `cluster show-netprops` command to display the current private-network configuration.  
If the current IP address range is not sufficient to support the additional private IP addresses that you will configure, follow the procedures in [“How to Change the Private Network Configuration When Adding Nodes or Private Networks”](#) on page 98 to reconfigure the private IP address range.

For additional information, see “Zone Components” in *System Administration Guide: Solaris Containers-Resource Management and Solaris Zones*.

### 1 Become superuser on the node on which you are creating the non-global zone.

You must be in the global zone.

### 2 For the Solaris 10 OS, verify on each node that multi-user services for the Service Management Facility (SMF) are online.

If services are not yet online for a node, wait until the state becomes online before you proceed to the next step.

```
phys-schost# svcs multi-user-server
STATE STIME FMRI
online 17:52:55 svc:/milestone/multi-user-server:default
```

### 3 Configure, install, and boot the new zone.

---

**Note** – You must set the `autoboot` property to `true` to support resource-group functionality in the non-global zone.

---

Follow procedures in the following documentation:

- a. Perform procedures in Chapter 18, “Planning and Configuring Non-Global Zones (Tasks),” in *System Administration Guide: Solaris Containers-Resource Management and Solaris Zones*.
- b. Perform procedures in “Installing and Booting Zones” in *System Administration Guide: Solaris Containers-Resource Management and Solaris Zones*.
- c. Perform procedures in “How to Boot a Zone” in *System Administration Guide: Solaris Containers-Resource Management and Solaris Zones*.

**4 Verify that the zone is in the ready state.**

```
phys-schost# zoneadm list -v
ID NAME STATUS PATH
0 global running /
1 my-zone ready /zone-path
```

**5 (Optional) Assign a private IP address and a private hostname to the zone.**

The following command chooses and assigns an available IP address from the cluster's private IP address range. The command also assigns the specified private hostname, or host alias, to the zone and maps it to the assigned private IP address.

```
phys-schost# clnode set -p zprivatehostname=hostalias node:zone
```

|                                         |                                                     |
|-----------------------------------------|-----------------------------------------------------|
| <code>-p</code>                         | Specifies a property.                               |
| <code>zprivatehostname=hostalias</code> | Specifies the zone private hostname, or host alias. |
| <code>node</code>                       | The name of the node.                               |
| <code>zone</code>                       | The name of the non-global zone.                    |

**6 Perform the initial internal zone configuration.**

Follow the procedures in “Performing the Initial Internal Zone Configuration” in *System Administration Guide: Solaris Containers-Resource Management and Solaris Zones*. Choose either of the following methods:

- Log in to the zone
- Use an `/etc/sysidcfg` file

**7 In the non-global zone, modify the `nsswitch.conf` file.**

You must make these changes to enable the zone to resolve searches for cluster-specific hostnames and IP addresses.

**a. Log in to the zone.**

```
phys-schost# zlogin -c zonename
```

**b. Open the `/etc/nsswitch.conf` file for editing.**

```
phys-schost# vi /etc/nsswitch.conf
```

**c. Add the `cluster switch` to the beginning of the lookups for the `hosts` and `netmasks` entries.**

The modified entries would appear similar to the following:

```
...
hosts: cluster files nis [NOTFOUND=return]
...
netmasks: cluster files nis [NOTFOUND=return]
...

```

**Next Steps** To install an application in a non-global zone, use the same procedure as for a standalone system. See your application's installation documentation for procedures to install the software in a non-global zone. Also see "Adding and Removing Packages and Patches on a Solaris System With Zones Installed (Task Map)" in *System Administration Guide: Solaris Containers-Resource Management and Solaris Zones*.

To install and configure a data service in a non-global zone, see the Sun Cluster manual for the individual data service.

# Installing and Upgrading the Sun Cluster Module to Sun Management Center

---

This chapter provides guidelines and procedures to install the Sun Cluster module to the Sun Management Center graphical user interface (GUI).

## SPARC: Installing the Sun Cluster Module for Sun Management Center

This section provides information and procedures to install software for the Sun Cluster module to Sun Management Center.

The Sun Cluster module for Sun Management Center enables you to use Sun Management Center to monitor the cluster. The following table lists the tasks to perform to install the Sun Cluster-module software for Sun Management Center.

TABLE 7-1 Task Map: Installing the Sun Cluster Module for Sun Management Center

| Task                                                                   | Instructions                                                                                                     |
|------------------------------------------------------------------------|------------------------------------------------------------------------------------------------------------------|
| 1. Install Sun Management Center server, agent, and console packages.  | Sun Management Center documentation<br>“SPARC: Installation Requirements for Sun Cluster Monitoring” on page 184 |
| 2. Install Sun Cluster-module packages.                                | “SPARC: How to Install the Sun Cluster Module for Sun Management Center” on page 185                             |
| 3. Start Sun Management Center server, console, and agent processes.   | “SPARC: How to Start Sun Management Center” on page 186                                                          |
| 4. Add each cluster node as a Sun Management Center agent host object. | “SPARC: How to Add a Cluster Node as a Sun Management Center Agent Host Object” on page 186                      |

TABLE 7-1 Task Map: Installing the Sun Cluster Module for Sun Management Center (Continued)

| Task                                                            | Instructions                                            |
|-----------------------------------------------------------------|---------------------------------------------------------|
| 5. Load the Sun Cluster module to begin to monitor the cluster. | "SPARC: How to Load the Sun Cluster Module" on page 187 |

## SPARC: Installation Requirements for Sun Cluster Monitoring

The Sun Cluster module for Sun Management Center is used to monitor a Sun Cluster configuration. Perform the following tasks before you install the Sun Cluster module packages.

- **Space requirements** - Ensure that 25 Mbytes of space is available on each cluster node for Sun Cluster-module packages.
- **Sun Management Center installation** - Follow procedures in your Sun Management Center installation documentation to install Sun Management Center software.

The following are additional requirements for a Sun Cluster configuration:

- Install the Sun Management Center agent package on each cluster node.
- When you install Sun Management Center on an agent machine (cluster node), choose whether to use the default of 161 for the agent (SNMP) communication port or another number. This port number enables the server to communicate with this agent. Record the port number that you choose for reference later when you configure the cluster nodes for monitoring.  
See your Sun Management Center installation documentation for information about choosing an SNMP port number.
- If you have an administrative console or other dedicated machine, you can run the console process on the administrative console and the server process on a separate machine. This installation approach improves Sun Management Center performance.
- For best results, install the Sun Management Center server and console packages on noncluster machines.
- If you choose to install the server or console packages on a cluster node, you might experience the following adverse performance impacts:
  - Cluster performance might be noticeably slowed due to the increased load from Sun Management Center processes. This is especially likely when the Sun Management Center server runs on a cluster node.
  - Sun Management Center is not made highly available when you install the server on a cluster node. Sun Management Center service will fail if the node goes down, such as during a failover to another node.
- **Web browser** - Ensure that the web browser that you use to connect to Sun Management Center is supported by Sun Management Center. Certain features, such as online help, might not be available on unsupported web browsers. See your Sun Management Center documentation for information about supported web browsers and any configuration requirements.



## ▼ SPARC: How to Install the Sun Cluster Module for Sun Management Center

Perform this procedure to install the Sun Cluster–module server package.

---

**Note** – The Sun Cluster–module agent packages, `SUNWscsa1` and `SUNWscsam`, are already added to cluster nodes during Sun Cluster software installation.

---

**Before You Begin** Ensure that all Sun Management Center core packages are installed on the appropriate machines. This task includes installing Sun Management Center agent packages on each cluster node. See your Sun Management Center documentation for installation instructions.

**1 On the Sun Management Center server machine, install the Sun Cluster–module server package `SUNWscssv`.**

a. Become superuser.

b. Load the Sun Java Availability Suite DVD-ROM into the DVD-ROM drive.

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

c. Change to the `Solaris_sparc/Product/sun_cluster/Solaris_ver/Packages/` directory, where `ver` is 9 for Solaris 9 or 10 for Solaris 10.

```
phys-schost# cd /cdrom/cdrom0/Solaris_sparc/Product/sun_cluster/Solaris_ver/Packages/
```

d. Install the Sun Cluster–module server package.

```
phys-schost# pkgadd -d . SUNWscssv
```

e. Unload the Sun Java Availability Suite DVD-ROM from the DVD-ROM drive.

i. To ensure that the DVD-ROM is not being used, change to a directory that does *not* reside on the DVD-ROM.

ii. Eject the DVD-ROM.

```
phys-schost# eject cdrom
```

**2 Install any Sun Cluster–module patches.**

See “Patches and Required Firmware Levels” in *Sun Cluster 3.2 Release Notes for Solaris OS* for the location of patches and installation instructions.

**Next Steps** Start Sun Management Center. Go to “[SPARC: How to Start Sun Management Center](#)” on page 186.

## ▼ SPARC: How to Start Sun Management Center

Perform this procedure to start the Sun Management Center server, agent, and console processes.

- 1 **As superuser, on the Sun Management Center server machine, start the Sun Management Center server process.**

The *install-dir* is the directory on which you installed the Sun Management Center software. The default directory is */opt*.

```
server# /install-dir/SUNWsymon/sbin/es-start -S
```

- 2 **As superuser, on each Sun Management Center agent machine (cluster node), start the Sun Management Center agent process.**

```
phys-schost# /install-dir/SUNWsymon/sbin/es-start -a
```

- 3 **On each Sun Management Center agent machine (cluster node), ensure that the `scsymon_srv` daemon is running.**

```
phys-schost# ps -ef | grep scsymon_srv
```

If any cluster node is not already running the `scsymon_srv` daemon, start the daemon on that node.

```
phys-schost# /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.

```
adminconsole% /install-dir/SUNWsymon/sbin/es-start -c
```

**Next Steps** Add a cluster node as a monitored host object. Go to [“SPARC: How to Add a Cluster Node as a Sun Management Center Agent Host Object”](#) on page 186.

## ▼ SPARC: How to Add a Cluster Node as a Sun Management Center Agent Host Object

Perform this procedure to create a Sun Management Center agent host object for a cluster node.

- 1 **Log in to Sun Management Center.**  
See your Sun Management Center documentation.
- 2 **From the Sun Management Center main window, select a domain from the Sun Management Center Administrative Domains pull-down list.**

This domain contains the Sun Management Center agent host object that you create. During Sun Management Center software installation, a Default Domain was automatically created for you. You can use this domain, select another existing domain, or create a new domain.

See your Sun Management Center documentation for information about how to create Sun Management Center domains.

- 3 Choose Edit⇒Create an Object from the pull-down menu.**
- 4 Click the Node tab.**
- 5 From the Monitor Via pull-down list, select Sun Management Center Agent - Host.**
- 6 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.
- 7 In the Port text field, type the port number that you chose when you installed the Sun Management Center agent machine.**
- 8 Click OK.**  
A Sun Management Center agent host object is created in the domain.

**Next Steps** Load the Sun Cluster module. Go to [“SPARC: How to Load the Sun Cluster Module”](#) on page 187.

**Troubleshooting** 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.

## ▼ **SPARC: How to Load the Sun Cluster Module**

Perform this procedure to start cluster monitoring.

- 1 In the Sun Management Center main window, right click the icon of a cluster node.**  
The pull-down menu is displayed.
- 2 Choose Load Module.**  
The Load Module window lists each available Sun Management Center module and whether the module is currently loaded.
- 3 Choose Sun Cluster: Not Loaded and click OK.**  
The Module Loader window shows the current parameter information for the selected module.
- 4 Click OK.**  
After a few moments, the module is loaded. A Sun Cluster icon is then displayed in the Details window.

**5 Verify that the Sun Cluster module is loaded.**

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.**

**See Also** See the Sun Cluster module online help for information about how to use Sun Cluster module features.

- To view online help for a specific Sun Cluster module item, place the cursor over the item. Then click the right mouse button and select Help from the pop-up menu.
- To access the home page for the Sun Cluster module online help, place the cursor over the Cluster Info icon. Then click the right mouse button and select Help from the pop-up menu.
- To directly access the home page for the Sun Cluster module online help, click the Sun Management Center Help button to launch the help browser. Then go to the following URL, where *install-dir* is the directory on which you installed the Sun Management Center software:

```
file:/install-dir/SUNWsymon/lib/locale/C/help/main.top.html
```

---

**Note** – The Help button in the Sun Management Center browser accesses online help for Sun Management Center, not the topics specific to the Sun Cluster module.

---

See Sun Management Center online help and your Sun Management Center documentation for information about how to use Sun Management Center.

**Next Steps** Install third-party applications, register resource types, set up resource groups, and configure data services. See the documentation that is supplied with the application software and the *Sun Cluster Data Services Planning and Administration Guide for Solaris OS*.

## SPARC: Upgrading Sun Management Center Software

This section provides the following procedures to upgrade the Sun Cluster module for Sun Management Center:

- [“SPARC: How to Upgrade Sun Cluster Module Software for Sun Management Center” on page 189](#)
- [“SPARC: How to Upgrade Sun Management Center Software” on page 190](#)

## ▼ SPARC: How to Upgrade Sun Cluster Module Software for Sun Management Center

Perform the following steps to upgrade Sun Cluster module software on the Sun Management Center server machine and console machine.

---

**Note** – To upgrade the Sun Management Center software itself, do not perform this procedure. Instead, go to “[SPARC: How to Upgrade Sun Management Center Software](#)” on page 190 to upgrade the Sun Management Center software and the Sun Cluster module.

---

**Before You Begin** Have available the Sun Java Availability Suite DVD-ROM for the SPARC platform or the path to the DVD-ROM image.

### 1 As superuser, remove the existing Sun Cluster module packages from each machine.

Use the `pkgrm(1M)` command to remove all Sun Cluster module packages from all locations that are listed in the following table.

| Location                              | Module Package to Remove |
|---------------------------------------|--------------------------|
| Sun Management Center console machine | SUNWscscn                |
| Sun Management Center server machine  | SUNWscssv, SUNWscshl     |

```
machine# pkgrm module-package
```

---

**Note** – Sun Cluster module software on the cluster nodes was already upgraded during the cluster-framework upgrade.

---

### 2 As superuser, reinstall Sun Cluster module packages on each machine.

a. Insert the Sun Java Availability Suite DVD-ROM for the SPARC platform into the DVD-ROM drive of the machine.

b. Change to the `Solaris_sparc/Product/sun_cluster/Solaris_ver/Packages/` directory, where `ver` is 9 for Solaris 9 or 10 for Solaris 10.

```
machine# cd Solaris_sparc/Product/sun_cluster/Solaris_ver/Packages/
```

c. Install the SUNWscssv package on the server machine.

Note that you do not upgrade to a new SUNWscscn package on the console machine or a new SUNWscshl package on the server machine.

```
machine# pkgadd -d . SUNWscssv
```

- d. **Unload the Sun Java Availability Suite DVD-ROM from the DVD-ROM drive.**
  - i. **To ensure that the DVD-ROM is not being used, change to a directory that does *not* reside on the DVD-ROM.**
  - ii. **Eject the DVD-ROM.**

```
machine# eject cdrom
```

## ▼ **SPARC: How to Upgrade Sun Management Center Software**

Perform the following steps to upgrade from Sun Management Center 2.1.1, 3.0, 3.5, or 3.6 software to Sun Management Center 3.6.1 software.

**Before You Begin** Have available the following items:

- Sun Java Availability Suite DVD-ROM for the SPARC platform and, if applicable, for the x86 platform, or the paths to the DVD-ROM images. You use the DVD-ROM to reinstall the Sun Cluster 3.2 version of the Sun Cluster module packages after you upgrade Sun Management Center software.

---

**Note** – The agent packages to install on the cluster nodes are available for both SPARC based systems and x86 based systems. The package for the server machine is available for SPARC based systems only.

---

- Sun Management Center documentation.
- Sun Management Center patches and Sun Cluster module patches, if any.  
See “Patches and Required Firmware Levels” in *Sun Cluster 3.2 Release Notes for Solaris OS* for the location of patches and installation instructions.

### **1 Stop any Sun Management Center processes.**

- a. **If the Sun Management Center console is running, exit the console.**  
In the console window, choose File⇒Exit.
- b. **On each Sun Management Center agent machine (cluster node), stop the Sun Management Center agent process.**

```
phys-schost# /opt/SUNWsymon/sbin/es-stop -a
```
- c. **On the Sun Management Center server machine, stop the Sun Management Center server process.**

```
server# /opt/SUNWsymon/sbin/es-stop -S
```

**2 As superuser, remove Sun Cluster–module packages.**

Use the `pkgrm(1M)` command to remove all Sun Cluster module packages from all locations that are listed in the following table.

| Location                                                                                         | Module Package to Remove |
|--------------------------------------------------------------------------------------------------|--------------------------|
| Each cluster node                                                                                | SUNWscsam, SUNWscsal     |
| Sun Management Center console machine                                                            | SUNWscscn                |
| Sun Management Center server machine                                                             | SUNWscssv                |
| Sun Management Center 3.0 help-server machine or Sun Management Center 3.5 or 3.6 server machine | SUNWscshl                |

```
machine# pkgrm module-package
```

If you do not remove the listed packages, the Sun Management Center software upgrade might fail because of package dependency problems. You reinstall these packages in [Step 4](#), after you upgrade Sun Management Center software.

**3 Upgrade the Sun Management Center software.**

Follow the upgrade procedures in your Sun Management Center documentation.

**4 As superuser, reinstall Sun Cluster module packages from the Sun Java Availability Suite DVD-ROM to the locations that are listed in the following table.**

| Location                             | Module Package to Install |
|--------------------------------------|---------------------------|
| Each cluster node                    | SUNWscsam, SUNWscsal      |
| Sun Management Center server machine | SUNWscssv                 |

- a. **Insert the Sun Java Availability Suite DVD-ROM for the appropriate platform in the DVD-ROM drive of the machine.**
- b. **Change to the `Solaris_arch/Product/sun_cluster/Solaris_ver/Packages/` directory, where `arch` is `sparc` or `x86`, and `ver` is `9` for Solaris 9 or `10` for Solaris 10.**

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

---

**Note** – The agent packages to install on the cluster nodes are available for both SPARC based systems and x86 based systems. The package for the server machine is available for SPARC based systems only.

---

**c. Install the appropriate module package on the machine.**

- **For cluster nodes that run the Solaris 10 OS, use the following command:**

```
phys-schost# pkgadd -G -d . module-package
```

The -G option adds packages to the current zone only. You must add these packages only to the global zone. Therefore, this option also specifies that the packages are *not* propagated to any existing non-global zone or to any non-global zone that is created later.

- **SPARC: For cluster nodes that run the Solaris 9 OS and for the server machine, use the following command:**

```
phys-schost# pkgadd -d . module-package
```

**5 Apply any Sun Management Center patches and any Sun Cluster module patches to each node of the cluster.**

**6 Restart Sun Management Center agent, server, and console processes.**

Follow procedures in [“SPARC: How to Start Sun Management Center”](#) on page 186.

**7 Load the Sun Cluster module.**

Follow procedures in [“SPARC: How to Load the Sun Cluster Module”](#) on page 187.

If the Sun Cluster module was previously loaded, unload the module and then reload it to clear all cached alarm definitions on the server. To unload the module, choose Unload Module from the Module menu on the console’s Details window.



# Upgrading Sun Cluster Software

---

This chapter provides the following information and procedures to upgrade a Sun Cluster 3.0 or 3.1 configuration to Sun Cluster 3.2 software:

- “Upgrade Requirements and Software Support Guidelines” on page 193
- “Choosing a Sun Cluster Upgrade Method” on page 195
- “Performing a Standard Upgrade to Sun Cluster 3.2 Software” on page 196
- “Performing a Dual-Partition Upgrade to Sun Cluster 3.2 Software” on page 214
- “Performing a Live Upgrade to Sun Cluster 3.2 Software” on page 236
- “Completing the Upgrade” on page 249
- “Recovering From an Incomplete Upgrade” on page 257

## Upgrade Requirements and Software Support Guidelines

Observe the following requirements and software-support guidelines when you upgrade to Sun Cluster 3.2 software:

- **Upgrade of x86 based systems** - On x86 based systems, you cannot upgrade from Sun Cluster 3.1 8/05 software on the Solaris 9 OS directly to Sun Cluster 3.2 software on the Solaris 10 OS. You must first upgrade Sun Cluster 3.1 8/05 software to at least the Solaris 10 1/06 OS.  
You can upgrade from Sun Cluster 3.1 8/05 software on the Solaris 10 OS directly to Sun Cluster 3.2 software on the Solaris 10 OS.
- **Minimum Sun Cluster software version** - Sun Cluster 3.2 software supports the following direct upgrade paths:
  - **SPARC: From version 3.0 including update releases to version 3.2** - Use the standard upgrade method only.
  - **SPARC: From version 3.1, 3.1 10/03, 3.1 4/04, or 3.1 9/04 to version 3.2** - Use the standard, dual-partition, or live upgrade method.
  - **From version 3.1 8/05 to version 3.2** - Use the standard, dual-partition, or live upgrade method.

See “[Choosing a Sun Cluster Upgrade Method](#)” on page 195 for additional requirements and restrictions for each upgrade method.

- **Minimum Solaris OS** - The cluster must run on or be upgraded to at least Solaris 9 9/05 software or Solaris 10 11/06 software, including the most current required patches. The Solaris 9 OS is supported only on SPARC based platforms.
- **Supported hardware** - The cluster hardware must be a supported configuration for Sun Cluster 3.2 software. Contact your Sun representative for information about current supported Sun Cluster configurations.
- **Architecture changes during upgrade** - Sun Cluster 3.2 software does not support upgrade between architectures.
- **Software migration** - Do not migrate from one type of software product to another product during Sun Cluster upgrade. For example, migration from Solaris Volume Manager disk sets to VxVM disk groups or from UFS file systems to VxFS file systems is not supported during Sun Cluster upgrade. Perform only software configuration changes that are specified by upgrade procedures of an installed software product.
- **Global-devices partition size** - If the size of your `/global/.devices/node@nodeid` partition is less than 512 Mbytes but it provides sufficient space for existing device nodes, you do not need to change the file-system size. The 512-Mbyte minimum applies to new installations of Sun Cluster 3.2 software. However, you must still ensure that the global-devices file system has ample space and ample inode capacity for existing devices and for any new devices that you intend to configure. Certain configuration changes, such as adding disks, disk volumes, or metadevices, might require increasing the partition size to provide sufficient additional inodes.
- **Data services** - You must upgrade all Sun Cluster data service software to version 3.2 and migrate resources to the new resource-type version. Sun Cluster 3.0 and 3.1 data services are not supported on Sun Cluster 3.2 software.
- **Upgrading to compatible versions** - You must upgrade all software on the cluster nodes to a version that is supported by Sun Cluster 3.2 software. For example, if a version of a data service is supported on Sun Cluster 3.1 software but is not supported on Sun Cluster 3.2 software, you must upgrade that data service to the version that is supported on Sun Cluster 3.2 software, if such a version exists. See “Supported Products” in *Sun Cluster 3.2 Release Notes for Solaris OS* for information about supported products.
- **Converting from NAFO to IPMP groups** - For upgrade from a Sun Cluster 3.0 release, have available the test IP addresses to use with your public-network adapters when NAFO groups are converted to IP network multipathing groups. The `scinstall` upgrade utility prompts you for a test IP address for each public-network adapter in the cluster. A test IP address must be on the same subnet as the primary IP address for the adapter.  
  
See “IPMP,” in *System Administration Guide: IP Services* (Solaris 9 or Solaris 10) for information about test IP addresses for IPMP groups.
- **Downgrade** - Sun Cluster 3.2 software does not support any downgrade of Sun Cluster software.
- **Limitation of `scinstall` for data-service upgrades** - The `scinstall` upgrade utility only upgrades those data services that are provided with Sun Cluster 3.2 software. You must manually upgrade any custom or third-party data services.

# Choosing a Sun Cluster Upgrade Method

Choose from the following methods to upgrade your cluster to Sun Cluster 3.2 software:

- **Standard upgrade** – In a standard upgrade, you shut down the cluster before you upgrade the cluster nodes. You return the cluster to production after all nodes are fully upgraded. Use this method if you are upgrading from a Sun Cluster 3.0 release.
- **Dual-partition upgrade** - In a *dual-partition* upgrade, you divide the cluster into two groups of nodes. You bring down one group of nodes and upgrade those nodes. The other group of nodes continues to provide services. After you complete upgrade of the first group of nodes, you switch services to those upgraded nodes. You then upgrade the remaining nodes and boot them back into the rest of the cluster. The cluster outage time is limited to the amount of time needed for the cluster to switch over services to the upgraded partition.

Observe the following additional restrictions and requirements for the dual-partition upgrade method:

- **Sun Cluster HA for Sun Java System Application Server EE (HADB)** - If you are running the Sun Cluster HA for Sun Java System Application Server EE (HADB) data service with Sun Java System Application Server EE (HADB) software as of version 4.4, you must shut down the database before you begin the dual-partition upgrade. The HADB database does not tolerate the loss of membership that would occur when a partition of nodes is shut down for upgrade. This requirement does not apply to versions before version 4.4.
- **Data format changes** - Do not use the dual-partition upgrade method if you intend to upgrade an application that requires that you change its data format during the application upgrade. The dual-partition upgrade method is not compatible with the extended downtime that is needed to perform data transformation.
- **Location of application software** - Applications must be installed on nonshared storage. Shared storage is not accessible to a partition that is in noncluster mode. Therefore, it is not possible to upgrade application software that is located on shared storage.
- **Division of storage** - Each shared storage device must be connected to a node in each group.
- **Single-node clusters** - Dual-partition upgrade is not available to upgrade a single-node cluster. Use the standard upgrade or live upgrade method instead.
- **Minimum Sun Cluster version** - The cluster must be running a Sun Cluster 3.1 release before you begin the dual-partition upgrade.
- **Configuration changes** - Do not make cluster configuration changes that are not documented in the upgrade procedures. Such changes might not be propagated to the final cluster configuration. Also, validation attempts of such changes would fail because not all nodes are reachable during a dual-partition upgrade.
- **Live upgrade** - A live upgrade maintains your previous cluster configuration until you have upgraded all nodes and you commit to the upgrade. If the upgraded configuration causes a problem, you can revert to your previous cluster configuration until you can rectify the problem.

Observe the following additional restrictions and requirements for the live upgrade method:

- **Minimum Sun Cluster version** - The cluster must be running a Sun Cluster 3.1 release before you begin the dual-partition upgrade.

- **Minimum version of Live Upgrade software** - To use the live upgrade method, you must use the Solaris Live Upgrade packages from at least the Solaris 9 9/04 or Solaris 10 release. This requirement applies to clusters running on all Solaris OS versions, including Solaris 8 software. The live upgrade procedures provide instructions for upgrading these packages.
- **Dual-partition upgrade** - The live upgrade method cannot be used in conjunction with a dual-partition upgrade.
- **HASStoragePlus** - The live upgrade method does not support the upgrade of clusters that run HASStoragePlus. Instead, use the standard or dual-partition upgrade method.
- **Non-global zones** - The live upgrade method does not support the upgrade of clusters that have non-global zones configured on any of the cluster nodes. Instead, use the standard upgrade or dual-partition upgrade method.
- **Disk space** - To use the live upgrade method, you must have enough spare disk space available to make a copy of each node's boot environment. You reclaim this disk space after the upgrade is complete and you have verified and committed the upgrade. For information about space requirements for an inactive boot environment, refer to "Solaris Live Upgrade Disk Space Requirements" in *Solaris 9 9/04 Installation Guide* or "Allocating Disk and Swap Space" in *Solaris 10 Installation Guide: Solaris Live Upgrade and Upgrade Planning*.

For overview information about planning your Sun Cluster 3.2 configuration, see [Chapter 1](#).

## Performing a Standard Upgrade to Sun Cluster 3.2 Software

This section provides the following information to upgrade to Sun Cluster 3.2 software by using the standard upgrade method:

- ["How to Prepare the Cluster for Upgrade \(Standard\)" on page 197](#)
- ["How to Upgrade the Solaris OS and Volume Manager Software \(Standard\)" on page 202](#)
- ["How to Upgrade Sun Cluster 3.2 Software \(Standard\)" on page 207](#)

The following table lists the tasks to perform to upgrade from Sun Cluster 3.1 software to Sun Cluster 3.2 software. You also perform these tasks to upgrade only the version of the Solaris OS. If you upgrade the Solaris OS from Solaris 9 to Solaris 10 software, you must also upgrade the Sun Cluster software and dependency software to the version that is compatible with the new version of the Solaris OS.

**TABLE 8-1** Task Map: Performing a Standard Upgrade to Sun Cluster 3.2 Software

| Task                                                                                                                     | Instructions                                                                                                                                              |
|--------------------------------------------------------------------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------|
| 1. Read the upgrade requirements and restrictions. Determine the proper upgrade method for your configuration and needs. | <a href="#">"Upgrade Requirements and Software Support Guidelines" on page 193</a><br><a href="#">"Choosing a Sun Cluster Upgrade Method" on page 195</a> |
| 2. Remove the cluster from production and back up shared data.                                                           | <a href="#">"How to Prepare the Cluster for Upgrade (Standard)" on page 197</a>                                                                           |

TABLE 8-1 Task Map: Performing a Standard Upgrade to Sun Cluster 3.2 Software (Continued)

| Task                                                                                                                                                                                                                                                                                                                                                    | Instructions                                                                              |
|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------------------------------------------------------------------------------|
| 3. Upgrade the Solaris software, if necessary, to a supported Solaris update. If the cluster uses dual-string mediators for Solaris Volume Manager software, unconfigure the mediators. As needed, upgrade VERITAS Volume Manager (VxVM) and VERITAS File System (VxFS). Solaris Volume Manager software is automatically upgraded with the Solaris OS. | “How to Upgrade the Solaris OS and Volume Manager Software (Standard)” on page 202        |
| 4. Upgrade to Sun Cluster 3.2 framework and data-service software. If necessary, upgrade applications. If the cluster uses dual-string mediators and you upgraded the Solaris OS, reconfigure the mediators. If you upgraded VxVM, upgrade disk groups.                                                                                                 | “How to Upgrade Sun Cluster 3.2 Software (Standard)” on page 207                          |
| 5. Verify successful completion of upgrade to Sun Cluster 3.2 software.                                                                                                                                                                                                                                                                                 | “How to Verify Upgrade of Sun Cluster 3.2 Software” on page 249                           |
| 6. Enable resources and bring resource groups online. Migrate existing resources to new resource types.                                                                                                                                                                                                                                                 | “How to Finish Upgrade to Sun Cluster 3.2 Software” on page 251                           |
| 7. (Optional) SPARC: Upgrade the Sun Cluster module for Sun Management Center, if needed.                                                                                                                                                                                                                                                               | “SPARC: How to Upgrade Sun Cluster Module Software for Sun Management Center” on page 189 |

## ▼ How to Prepare the Cluster for Upgrade (Standard)

Perform this procedure to remove the cluster from production before you perform a standard upgrade. On the Solaris 10 OS, perform all steps from the global zone only.

**Before You Begin** Perform the following tasks:

- Ensure that the configuration meets the requirements for upgrade. See “Upgrade Requirements and Software Support Guidelines” on page 193.
- Have available the installation media, documentation, and patches for all software products that you are upgrading, including the following software:
  - Solaris OS
  - Sun Cluster 3.2 framework
  - Sun Cluster 3.2 data services (agents)
  - Applications that are managed by Sun Cluster 3.2 data-services
  - VERITAS Volume Manager, if applicable

See “Patches and Required Firmware Levels” in *Sun Cluster 3.2 Release Notes for Solaris OS* for the location of patches and installation instructions.

- If you use role-based access control (RBAC) instead of superuser to access the cluster nodes, ensure that you can assume an RBAC role that provides authorization for all Sun Cluster commands. This series of upgrade procedures requires the following Sun Cluster RBAC authorizations if the user is not superuser:

- `solaris.cluster.modify`
- `solaris.cluster.admin`
- `solaris.cluster.read`

See “Role-Based Access Control (Overview)” in *System Administration Guide: Security Services* for more information about using RBAC roles. See the Sun Cluster man pages for the RBAC authorization that each Sun Cluster subcommand requires.

**1 Ensure that the cluster is functioning normally.**

- a. View the current status of the cluster by running the following command from any node.**

```
phys-schost% scstat
```

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

- b. Search the `/var/adm/messages` log on the same node for unresolved error messages or warning messages.**
- c. Check the volume-manager status.**

**2 Notify users that cluster services will be unavailable during the upgrade.**

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

**4 Take each resource group offline and disable all resources.**

Take offline all resource groups in the cluster, including those that are in non-global zones. Then disable all resources, to prevent the cluster from bringing the resources online automatically if a node is mistakenly rebooted into cluster mode.

- **If you are upgrading from Sun Cluster 3.1 software and want to use the `scsetup` utility, perform the following steps:**

- a. Start the `scsetup` utility.**

```
phys-schost# scsetup
```

The `scsetup` Main Menu is displayed.

- b. Type the number that corresponds to the option for Resource groups and press the Return key.**

The Resource Group Menu is displayed.

- c. Type the number that corresponds to the option for Online/Offline or Switchover a resource group and press the Return key.**

- d. Follow the prompts to take offline all resource groups and to put them in the unmanaged state.**

e. **When all resource groups are offline, type q to return to the Resource Group Menu.**

f. **Exit the `scsetup` utility.**

Type **q** to back out of each submenu or press **Ctrl-C**.

■ **To use the command line, perform the following steps:**

a. **Take each resource offline.**

```
phys-schost# scswitch -F -g resource-group
```

-F Switches a resource group offline.

-g *resource-group* Specifies the name of the resource group to take offline.

b. **From any node, list all enabled resources in the cluster.**

```
phys-schost# scrgadm -pv | grep "Res enabled"
(resource-group:resource) Res enabled: True
```

c. **Identify those resources that depend on other resources.**

You must disable dependent resources first before you disable the resources that they depend on.

d. **Disable each enabled resource in the cluster.**

```
phys-schost# scswitch -n -j resource
```

-n Disables.

-j *resource* Specifies the resource.

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

e. **Verify that all resources are disabled.**

```
phys-schost# scrgadm -pv | grep "Res enabled"
(resource-group:resource) Res enabled: False
```

f. **Move each resource group to the unmanaged state.**

```
phys-schost# scswitch -u -g resource-group
```

-u Moves the specified resource group to the unmanaged state.

-g *resource-group* Specifies the name of the resource group to move into the unmanaged state.

**5 Verify that all resources on all nodes are Offline and that all resource groups are in the Unmanaged state.**

```
phys-schost# scstat
```

- 6 For a two-node cluster that uses Sun StorEdge Availability Suite software or Sun StorageTek™ Availability Suite software, ensure that the configuration data for availability services resides on the quorum disk.**

The configuration data must reside on a quorum disk to ensure the proper functioning of Availability Suite after you upgrade the cluster software.

- a. Become superuser on a node of the cluster that runs Availability Suite software.**
- b. Identify the device ID and the slice that is used by the Availability Suite configuration file.**

```
phys-schost# /usr/opt/SUNWscm/sbin/dscfg
/dev/did/rdisk/dNsS
```

In this example output, *N* is the device ID and *S* the slice of device *N*.

- c. Identify the existing quorum device.**

```
phys-schost# scstat -q
-- Quorum Votes by Device --
 Device Name Present Possible Status

Device votes: /dev/did/rdisk/dQsS 1 1 Online
```

In this example output, dQsS is the existing quorum device.

- d. If the quorum device is not the same as the Availability Suite configuration-data device, move the configuration data to an available slice on the quorum device.**

```
phys-schost# dd if='/usr/opt/SUNWesm/sbin/dscfg' of=/dev/did/rdisk/dQsS
```

---

**Note** – You must use the name of the raw DID device, /dev/did/rdisk/, not the block DID device, /dev/did/dsk/.

---

- e. If you moved the configuration data, configure Availability Suite software to use the new location.**

As superuser, issue the following command on each node that runs Availability Suite software.

```
phys-schost# /usr/opt/SUNWesm/sbin/dscfg -s /dev/did/rdisk/dQsS
```

- 7 (Optional) If you are upgrading from a version of Sun Cluster 3.0 software and do not want your ntp.conf file renamed to ntp.conf.cluster, create an ntp.conf.cluster file.**

On each node, copy /etc/inet/ntp.cluster as ntp.conf.cluster.

```
phys-schost# cp /etc/inet/ntp.cluster /etc/inet/ntp.conf.cluster
```

The existence of an ntp.conf.cluster file prevents upgrade processing from renaming the ntp.conf file. The ntp.conf file will still be used to synchronize NTP among the cluster nodes.

- 8 Stop all applications that are running on each node of the cluster.**



**9 Ensure that all shared data is backed up.****10 If you will upgrade the Solaris OS and your cluster uses dual-string mediators for Solaris Volume Manager software, unconfigure your mediators.**

See “[Configuring Dual-String Mediators](#)” on page 155 for more information about mediators.

**a. Run the following command to verify that no mediator data problems exist.**

```
phys-schost# medstat -s setname
```

-s *setname* Specifies the disk set name.

If the value in the Status field is Bad, repair the affected mediator host. Follow the procedure “[How to Fix Bad Mediator Data](#)” on page 157.

**b. List all mediators.**

Save this information for when you restore the mediators during the procedure “[How to Finish Upgrade to Sun Cluster 3.2 Software](#)” on page 251.

**c. For a disk set that uses mediators, take ownership of the disk set if no node already has ownership.**

```
phys-schost# scswitch -z -D setname -h node
```

-z Changes mastery.

-D *devicegroup* Specifies the name of the disk set.

-h *node* Specifies the name of the node to become primary of the disk set.

**d. Unconfigure all mediators for the disk set.**

```
phys-schost# metaset -s setname -d -m mediator-host-list
```

-s *setname* Specifies the disk set name.

-d Deletes from the disk set.

-m *mediator-host-list* Specifies the name of the node to remove as a mediator host for the disk set.

See the `mediator(7D)` man page for further information about mediator-specific options to the `metaset` command.

**e. Repeat Step c through Step d for each remaining disk set that uses mediators.****11 Ensure that each system disk is backed up.**

**Next Steps** Upgrade software on each node.

- To upgrade Solaris software before you perform Sun Cluster software upgrade, go to [“How to Upgrade the Solaris OS and Volume Manager Software \(Standard\)”](#) on page 202.
  - You must upgrade the Solaris software to a supported release if Sun Cluster 3.2 software does not support the release of the Solaris OS that your cluster currently runs. See “Supported Products” in *Sun Cluster 3.2 Release Notes for Solaris OS* for more information.
  - If Sun Cluster 3.2 software supports the release of the Solaris OS that you currently run on your cluster, further Solaris software upgrade is optional.
- Otherwise, upgrade to Sun Cluster 3.2 software. Go to [“How to Upgrade Sun Cluster 3.2 Software \(Standard\)”](#) on page 207.

## ▼ How to Upgrade the Solaris OS and Volume Manager Software (Standard)

Perform this procedure on each node in the cluster to upgrade the Solaris OS. On the Solaris 10 OS, perform all steps from the global zone only. If the cluster already runs on a version of the Solaris OS that supports Sun Cluster 3.2 software, further upgrade of the Solaris OS is optional. If you do not intend to upgrade the Solaris OS, proceed to [“How to Upgrade Sun Cluster 3.2 Software \(Standard\)”](#) on page 207.

---

**Note** – The cluster must already run on, or be upgraded to, at least the minimum required level of the Solaris OS to support upgrade to Sun Cluster 3.2 software. See Supported Products in *Sun Cluster 3.2 Release Notes for Solaris OS* for more information.

---

**Before You Begin** Ensure that all steps in [“How to Prepare the Cluster for Upgrade \(Standard\)”](#) on page 197 are completed.

**1 Become superuser on the cluster node to upgrade.**

If you are performing a dual-partition upgrade, the node must be a member of the partition that is in noncluster mode.

**2 If Sun Cluster Geographic Edition software is installed, uninstall it.**

For uninstallation procedures, see the documentation for your version of Sun Cluster Geographic Edition software.

**3 Determine whether the following Apache run-control scripts exist and are enabled or disabled:**

```
/etc/rc0.d/K16apache
/etc/rc1.d/K16apache
/etc/rc2.d/K16apache
/etc/rc3.d/S50apache
/etc/rc5.d/K16apache
```

Some applications, such as Sun Cluster HA for Apache, require that Apache run control scripts be disabled.

- If these scripts exist and contain an uppercase K or S in the file name, the scripts are enabled. No further action is necessary for these scripts.
- If these scripts do not exist, in [Step 8](#) you must ensure that any Apache run control scripts that are installed during the Solaris OS upgrade are disabled.
- If these scripts exist but the file names contain a lowercase k or s, the scripts are disabled. In [Step 8](#) you must ensure that any Apache run control scripts that are installed during the Solaris OS upgrade are disabled.

#### 4 Comment out all entries for globally mounted file systems in the node's `/etc/vfstab` file.

- a. For later reference, make a record of all entries that are already commented out.
- b. Temporarily comment out all entries for globally mounted file systems in the `/etc/vfstab` file.  
Entries for globally mounted file systems contain the `global` mount option. Comment out these entries to prevent the Solaris upgrade from attempting to mount the global devices.

#### 5 Determine which procedure to follow to upgrade the Solaris OS.

| Volume Manager         | Procedure                                                        | Location of Instructions                          |
|------------------------|------------------------------------------------------------------|---------------------------------------------------|
| Solaris Volume Manager | Any Solaris upgrade method <i>except</i> the Live Upgrade method | Solaris installation documentation                |
| VERITAS Volume Manager | "Upgrading VxVM and Solaris"                                     | VERITAS Volume Manager installation documentation |

**Note** – If your cluster has VxVM installed, you must reinstall the existing VxVM software or upgrade to the Solaris 9 or 10 version of VxVM software as part of the Solaris upgrade process.

#### 6 Upgrade the Solaris software, following the procedure that you selected in [Step 5](#).

**Note** – *Do not* perform the final reboot instruction in the Solaris software upgrade. Instead, do the following:

- a. Return to this procedure to perform [Step 7](#) and [Step 8](#).
  - b. Reboot into noncluster mode in [Step 9](#) to complete Solaris software upgrade.
- When prompted, choose the manual reboot option.

- When you are instructed to reboot a node during the upgrade process, *always* reboot into noncluster mode. For the boot and reboot commands, add the `-x` option to the command. The `-x` option ensures that the node reboots into noncluster mode. For example, either of the following two commands boot a node into single-user noncluster mode:

- On SPARC based systems, perform either of the following commands:**

```
phys-schost# reboot -- -xs
 or
ok boot -xs
```

If the instruction says to run the `init 5` command, use the `reboot -- -xs` command instead.

- On x86 based systems running the Solaris 9 OS, perform either of the following commands:**

```
phys-schost# reboot -- -xs
 or
...
```

```

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

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

```

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

- On x86 based systems running the Solaris 10 OS, perform the following command:**

```
phys-schost# shutdown -g -y -i0Press any key to continue
```

- In the GRUB menu, use the arrow keys to select the appropriate Solaris entry and type `e` to edit its commands.**

The GRUB menu appears similar to the following:

```
GNU GRUB version 0.95 (631K lower / 2095488K upper memory)
+-----+
| Solaris 10 /sol_10_x86 |
| Solaris failsafe |
| |
+-----+
Use the ^ and v keys to select which entry is highlighted.
Press enter to boot the selected OS, 'e' to edit the
commands before booting, or 'c' for a command-line.
```

Use the `^` and `v` keys to select which entry is highlighted. Press enter to boot the selected OS, `'e'` to edit the commands before booting, or `'c'` for a command-line.

For more information about GRUB based booting, see Chapter 11, “GRUB Based Booting (Tasks),” in *System Administration Guide: Basic Administration*.

- b. In the boot parameters screen, use the arrow keys to select the kernel entry and type e to edit the entry.**

The GRUB boot parameters screen appears similar to the following:

```
GNU GRUB version 0.95 (615K lower / 2095552K upper memory)
```

```
+-----+
| root (hd0,0,a) |
| kernel /platform/i86pc/multiboot |
| module /platform/i86pc/boot_archive |
+-----+
```

Use the ^ and v keys to select which entry is highlighted. Press 'b' to boot, 'e' to edit the selected command in the boot sequence, 'c' for a command-line, 'o' to open a new line after ('O' for before) the selected line, 'd' to remove the selected line, or escape to go back to the main menu.

- c. Add -x to the command to specify that the system boot into noncluster mode.**

[ Minimal BASH-like line editing is supported. For the first word, TAB lists possible command completions. Anywhere else TAB lists the possible completions of a device/filename. ESC at any time exits. ]

```
grub edit> kernel /platform/i86pc/multiboot -x
```

- d. Press Enter to accept the change and return to the boot parameters screen.**

The screen displays the edited command.

```
GNU GRUB version 0.95 (615K lower / 2095552K upper memory)
```

```
+-----+
| root (hd0,0,a) |
| kernel /platform/i86pc/multiboot -x |
| module /platform/i86pc/boot_archive |
+-----+
```

Use the ^ and v keys to select which entry is highlighted. Press 'b' to boot, 'e' to edit the selected command in the boot sequence, 'c' for a command-line, 'o' to open a new line after ('O' for before) the selected line, 'd' to remove the selected line, or escape to go back to the main menu.-

- e. Type b to boot the node into noncluster mode.**

---

**Note** – This change to the kernel boot parameter command does not persist over the system boot. The next time you reboot the node, it will boot into cluster mode. To boot into noncluster mode instead, perform these steps to again to add the -x option to the kernel boot parameter command.

---

If the instruction says to run the `init S` command, shut down the system then change the GRUB kernel boot command to `/platform/i86pc/multiboot -sx` instead.

**7 In the `/a/etc/vfstab` file, uncomment those entries for globally mounted file systems that you commented out in [Step 4](#).**

**8 If Apache run control scripts were disabled or did not exist before you upgraded the Solaris OS, ensure that any scripts that were installed during Solaris upgrade are disabled.**

To disable Apache run control scripts, use the following commands to rename the files with a lowercase k or s.

```
phys-schost# mv /a/etc/rc0.d/K16apache /a/etc/rc0.d/k16apache
phys-schost# mv /a/etc/rc1.d/K16apache /a/etc/rc1.d/k16apache
phys-schost# mv /a/etc/rc2.d/K16apache /a/etc/rc2.d/k16apache
phys-schost# mv /a/etc/rc3.d/S50apache /a/etc/rc3.d/s50apache
phys-schost# mv /a/etc/rcS.d/K16apache /a/etc/rcS.d/k16apache
```

Alternatively, you can rename the scripts to be consistent with your normal administration practices.

**9 Reboot the node into noncluster mode.**

Include the double dashes (`--`) in the following command:

```
phys-schost# reboot -- -x
```

**10 If your cluster runs VxVM, perform the remaining steps in the procedure “Upgrading VxVM and Solaris” to reinstall or upgrade VxVM.**

Make the following changes to the procedure:

- **After VxVM upgrade is complete but before you reboot, verify the entries in the `/etc/vfstab` file.**  
If any of the entries that you uncommented in [Step 7](#) were commented out, make those entries uncommented again.
- **When the VxVM procedures instruct you to perform a final reconfiguration reboot, *do not* use the `-r` option alone. Instead, reboot into noncluster mode by using the `-rx` options.**
  - **On SPARC based systems, perform the following command:**  

```
phys-schost# reboot -- -rx
```
  - **On x86 based systems, perform the shutdown and boot procedures that are described in [Step 6](#) except add `-rx` to the kernel boot command instead of `-sx`.**

---

**Note** – If you see a message similar to the following, type the root password to continue upgrade processing. Do *not* run the fsck command nor type Ctrl-D.

```
WARNING - Unable to repair the /global/.devices/node@1 filesystem.
Run fsck manually (fsck -F ufs /dev/vx/rdisk/rootdisk_13vol). Exit the
shell when done to continue the boot process.
```

```
Type control-d to proceed with normal startup,
(or give root password for system maintenance): Type the root password
```

---

### 11 (Optional) SPARC: Upgrade VxFS.

Follow procedures that are provided in your VxFS documentation.

### 12 Install any required Solaris software patches and hardware-related patches, and download any needed firmware that is contained in the hardware patches.

---

**Note** – Do not reboot after you add patches. Wait to reboot the node until after you upgrade the Sun Cluster software.

See “Patches and Required Firmware Levels” in *Sun Cluster 3.2 Release Notes for Solaris OS* for the location of patches and installation instructions.

**Next Steps** Upgrade to Sun Cluster 3.2 software. Go to “[How to Upgrade Sun Cluster 3.2 Software \(Standard\)](#)” on page 207.

---

**Note** – To complete the upgrade to a new marketing release of the Solaris OS, such as from Solaris 8 to Solaris 10 software, you must also upgrade the Sun Cluster software and dependency software to the version that is compatible with the new version of the Solaris OS.

---

## ▼ How to Upgrade Sun Cluster 3.2 Software (Standard)

Perform this procedure to upgrade each node of the cluster to Sun Cluster 3.2 software. This procedure also upgrades required Sun Java Enterprise System shared components.

You must also perform this procedure after you upgrade to a different marketing release of the Solaris OS, such as from Solaris 8 to Solaris 10 software.

On the Solaris 10 OS, perform all steps from the global zone only.

---

**Tip** – You can perform this procedure on more than one node at the same time.

---

**Before You Begin** Perform the following tasks:

- Ensure that all steps in “How to Prepare the Cluster for Upgrade (Standard)” on page 197 are completed.
- If you upgraded to a new marketing release of the Solaris OS, such as from Solaris 8 to Solaris 10 software, ensure that all steps in “How to Upgrade the Solaris OS and Volume Manager Software (Standard)” on page 202 are completed.
- Ensure that you have installed all required Solaris software patches and hardware-related patches.

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

### 2 Ensure that the `/usr/java/` directory is a symbolic link to the minimum or latest version of Java software.

Sun Cluster software requires at least version 1.5.0\_06 of Java software. If you upgraded to a version of Solaris that installs an earlier version of Java, the upgrade might have changed the symbolic link to point to a version of Java that does not meet the minimum requirement for Sun Cluster 3.2 software.

#### a. Determine what directory the `/usr/java/` directory is symbolically linked to.

```
phys-schost# ls -l /usr/java
lrwxrwxrwx 1 root other 9 Apr 19 14:05 /usr/java -> /usr/j2se/
```

#### b. Determine what version or versions of Java software are installed.

The following are examples of commands that you can use to display the version of their related releases of Java software.

```
phys-schost# /usr/j2se/bin/java -version
phys-schost# /usr/java1.2/bin/java -version
phys-schost# /usr/jdk/jdk1.5.0_06/bin/java -version
```

#### c. If the `/usr/java/` directory is not symbolically linked to a supported version of Java software, recreate the symbolic link to link to a supported version of Java software.

The following example shows the creation of a symbolic link to the `/usr/j2se/` directory, which contains Java 1.5.0\_06 software.

```
phys-schost# rm /usr/java
phys-schost# ln -s /usr/j2se /usr/java
```

### 3 Load the Sun Java Availability Suite DVD-ROM into the DVD-ROM drive.

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

### 4 Change to the installation wizard directory of the DVD-ROM.

- If you are installing the software packages on the SPARC platform, type the following command:

```
phys-schost# cd /cdrom/cdrom0//Solaris_sparc
```

- If you are installing the software packages on the x86 platform, type the following command:



```
phys-schost# cd /cdrom/cdrom0//Solaris_x86
```

**5 Start the installation wizard program.**

```
phys-schost# ./installer
```

**6 Follow the instructions on the screen to select and upgrade Shared Components software packages on the node.**

---

**Note** – Do *not* use the installation wizard program to upgrade Sun Cluster software packages.

---

The installation wizard program displays the status of the installation. When the installation is complete, the program displays an installation summary and the installation logs.

**7 Exit the installation wizard program.**

**8 Change to the Solaris\_arch/Product/sun\_cluster/Solaris\_ver/Tools/ directory, where arch is sparc or x86 (Solaris 10 only) and where ver is 9 for Solaris 9 or 10 for Solaris 10.**

```
phys-schost# cd /cdrom/cdrom0/Solaris_arch/Product/sun_cluster/Solaris_ver/Tools
```

**9 Start the scinstall utility.**

```
phys-schost# ./scinstall
```

---

**Note** – Do not use the /usr/cluster/bin/scinstall command that is already installed on the node. You must use the scinstall command that is located on the Sun Java Availability Suite DVD-ROM.

---

The scinstall Main Menu is displayed.

**10 Type the number that corresponds to the option for Upgrade this cluster node and press the Return key.**

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

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

- \* 1) Create a new cluster or add a cluster node
- 2) Configure a cluster to be JumpStarted from this install server
- \* 3) Manage a dual-partition upgrade
- \* 4) Upgrade this cluster node
- \* 5) Print release information for this cluster node
  
- \* ?) Help with menu options
- \* q) Quit

```
Option: 4
```

The Upgrade Menu is displayed.

**11 Type the number that corresponds to the option for Upgrade Sun Cluster framework on this cluster node and press the Return key.**

**12 Follow the menu prompts to upgrade the cluster framework.**

During the Sun Cluster upgrade, `scinstall` might make one or more of the following configuration changes:

- Convert NAFO groups to IPMP groups but keep the original NAFO-group name. See one of the following manuals for information about test addresses for IPMP:
  - “Configuring Test Addresses” in “Administering Multipathing Groups With Multiple Physical Interfaces” in *System Administration Guide: IP Services* (Solaris 9)
  - “Test Addresses” in *System Administration Guide: IP Services* (Solaris 10)

See the `scinstall(1M)` man page for more information about the conversion of NAFO groups to IPMP during Sun Cluster software upgrade.

- Rename the `ntp.conf` file to `ntp.conf.cluster`, if `ntp.conf.cluster` does not already exist on the node.
- Set the `local-mac-address?` variable to `true`, if the variable is not already set to that value.

Upgrade processing is finished when the system displays the message `Completed Sun Cluster framework upgrade` and prompts you to press `Enter` to continue.

**13 Quit the `scinstall` utility.**

**14 Unload the Sun Java Availability Suite DVD-ROM from the DVD-ROM drive.**

a. To ensure that the DVD-ROM is not being used, change to a directory that does *not* reside on the DVD-ROM.

b. Eject the DVD-ROM.

```
phys-schost# eject cdrom
```

**15 Upgrade data service packages.**

You must upgrade all data services to the Sun Cluster 3.2 version.

---

**Note** – For Sun Cluster HA for SAP Web Application Server, if you are using a J2EE engine resource or a web application server component resource or both, you must delete the resource and recreate it with the new web application server component resource. Changes in the new web application server component resource includes integration of the J2EE functionality. For more information, see *Sun Cluster Data Service for SAP Web Application Server Guide for Solaris OS*.

---

a. **Start the upgraded interactive `scinstall` utility.**

```
phys-schost# /usr/cluster/bin/scinstall
```

---

**Note** – Do not use the `scinstall` utility that is on the installation media to upgrade data service packages.

---

The `scinstall` Main Menu is displayed.

- b. Type the number that corresponds to the option for Upgrade this cluster node and press the Return key.**

The Upgrade Menu is displayed.

- c. Type the number that corresponds to the option for Upgrade Sun Cluster data service agents on this node and press the Return key.**

- d. Follow the menu prompts to upgrade Sun Cluster data service agents that are installed on the node.**

You can choose from the list of data services that are available to upgrade or choose to upgrade all installed data services.

Upgrade processing is finished when the system displays the message `Completed upgrade of Sun Cluster data services agents` and prompts you to press `Enter` to continue.

- e. Press Enter.**

The Upgrade Menu is displayed.

**16 Quit the `scinstall` utility.**

**17 If you have Sun Cluster HA for NFS configured on a highly available local file system, ensure that the loopback file system (LOFS) is disabled.**

---

**Note** – If you have non-global zones configured, LOFS must remain enabled. For guidelines about using LOFS and alternatives to disabling it, see [“Cluster File Systems” on page 32](#).

---

As of the Sun Cluster 3.2 release, LOFS is no longer disabled by default during Sun Cluster software installation or upgrade. To disable LOFS, ensure that the `/etc/system` file contains the following entry:

```
exclude:lofs
```

This change becomes effective at the next system reboot.

**18 As needed, manually upgrade any custom data services that are not supplied on the product media.**

**19 Verify that each data-service update is installed successfully.**

View the upgrade log file that is referenced at the end of the upgrade output messages.

**20 Install any Sun Cluster 3.2 framework and data-service software patches.**

See “Patches and Required Firmware Levels” in *Sun Cluster 3.2 Release Notes for Solaris OS* for the location of patches and installation instructions.

**21 Upgrade software applications that are installed on the cluster.**

Ensure that application levels are compatible with the current versions of Sun Cluster and Solaris software. See your application documentation for installation instructions.

**22 (Optional) Reconfigure the private-network address range.**

Perform this step if you want to increase or decrease the size of the IP address range that is used by the private interconnect. The IP address range that you configure must minimally support the number of nodes and private networks in the cluster. See “Private Network” on page 26 for more information.

**a. From one node, start the `clsetup` utility.**

When run in noncluster mode, the `clsetup` utility displays the Main Menu for noncluster-mode operations.

**b. Type the number that corresponds to the option for Change IP Address Range and press the Return key.**

The `clsetup` utility displays the current private-network configuration, then asks if you would like to change this configuration.

**c. To change either the private-network IP address or the IP address range, type `yes` and press the Return key.**

The `clsetup` utility displays the default private-network IP address, `172.16.0.0`, and asks if it is okay to accept this default.

**d. Change or accept the private-network IP address.**

▪ **To accept the default private-network IP address and proceed to changing the IP address range, type `yes` and press the Return key.**

The `clsetup` utility will ask if it is okay to accept the default netmask. Skip to the next step to enter your response.

▪ **To change the default private-network IP address, perform the following substeps.**

i. **Type `no` in response to the `clsetup` utility question about whether it is okay to accept the default address, then press the Return key.**

The `clsetup` utility will prompt for the new private-network IP address.

ii. **Type the new IP address and press the Return key.**

The `clsetup` utility displays the default netmask and then asks if it is okay to accept the default netmask.

**e. Change or accept the default private-network IP address range.**

The default netmask is 255 . 255 . 248 . 0. This default IP address range supports up to 64 nodes and up to 10 private networks in the cluster.

- **To accept the default IP address range, type `yes` and press the Return key.**

Then skip to the next step.

- **To change the IP address range, perform the following substeps.**

- i. **Type `no` in response to the `clsetup` utility's question about whether it is okay to accept the default address range, then press the Return key.**

When you decline the default netmask, the `clsetup` utility prompts you for the number of nodes and private networks that you expect to configure in the cluster.

- ii. **Enter the number of nodes and private networks that you expect to configure in the cluster.**

From these numbers, the `clsetup` utility calculates two proposed netmasks:

- The first netmask is the minimum netmask to support the number of nodes and private networks that you specified.
- The second netmask supports twice the number of nodes and private networks that you specified, to accommodate possible future growth.

- iii. **Specify either of the calculated netmasks, or specify a different netmask that supports the expected number of nodes and private networks.**

- f. **Type `yes` in response to the `clsetup` utility's question about proceeding with the update.**

- g. **When finished, exit the `clsetup` utility.**

**23 After all nodes in the cluster are upgraded, reboot the upgraded nodes.**

- a. **Shut down each node.**

```
phys-schost# shutdown -g0 -y
```

- b. **Boot each node into cluster mode.**

- **On SPARC based systems, do the following:**

```
ok boot
```

- **On x86 based systems, do the following:**

When the GRUB menu is displayed, select the appropriate Solaris entry and press Enter. The GRUB menu appears similar to the following:

```
GNU GRUB version 0.95 (631K lower / 2095488K upper memory)
+-----+
| Solaris 10 /sol_10_x86 |
| Solaris failsafe |
| |
+-----+
```

Use the ^ and v keys to select which entry is highlighted.

Press enter to boot the selected OS, 'e' to edit the commands before booting, or 'c' for a command-line.

For more information about GRUB based booting, see Chapter 11, “GRUB Based Booting (Tasks),” in *System Administration Guide: Basic Administration*.

**Next Steps** Go to “How to Verify Upgrade of Sun Cluster 3.2 Software” on page 249

## Performing a Dual-Partition Upgrade to Sun Cluster 3.2 Software

This section provides the following information to upgrade from a Sun Cluster 3.1 release to Sun Cluster 3.2 software by using the dual-partition upgrade method:

- “How to Prepare the Cluster for Upgrade (Dual-Partition)” on page 215
- “How to Upgrade the Solaris OS and Volume Manager Software (Dual-Partition)” on page 224
- “How to Upgrade Sun Cluster 3.2 Software (Dual-Partition)” on page 229

The following table lists the tasks to perform to upgrade from Sun Cluster 3.1 software to Sun Cluster 3.2 software. You also perform these tasks to upgrade only the version of the Solaris OS. If you upgrade the Solaris OS from Solaris 9 to Solaris 10 software, you must also upgrade the Sun Cluster software and dependency software to the version that is compatible with the new version of the Solaris OS.

**TABLE 8–2** Task Map: Performing a Dual-Partition Upgrade to Sun Cluster 3.2 Software

| Task                                                                                                                     | Instructions                                                                                                                         |
|--------------------------------------------------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------|
| 1. Read the upgrade requirements and restrictions. Determine the proper upgrade method for your configuration and needs. | <p>“Upgrade Requirements and Software Support Guidelines” on page 193</p> <p>“Choosing a Sun Cluster Upgrade Method” on page 195</p> |
| 2. Partition the cluster into two groups of nodes.                                                                       | <p>“How to Prepare the Cluster for Upgrade (Dual-Partition)” on page 215</p>                                                         |

**TABLE 8-2** Task Map: Performing a Dual-Partition Upgrade to Sun Cluster 3.2 Software (Continued)

| Task                                                                                                                                                                                                                                                                                                                                                    | Instructions                                                                              |
|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------------------------------------------------------------------------------|
| 3. Upgrade the Solaris software, if necessary, to a supported Solaris update. If the cluster uses dual-string mediators for Solaris Volume Manager software, unconfigure the mediators. As needed, upgrade VERITAS Volume Manager (VxVM) and VERITAS File System (VxFS). Solaris Volume Manager software is automatically upgraded with the Solaris OS. | “How to Upgrade the Solaris OS and Volume Manager Software (Dual-Partition)” on page 224  |
| 4. Upgrade to Sun Cluster 3.2 framework and data-service software. If necessary, upgrade applications. If the cluster uses dual-string mediators and you upgraded the Solaris OS, reconfigure the mediators. If you upgraded VxVM, upgrade disk groups.                                                                                                 | “How to Upgrade Sun Cluster 3.2 Software (Dual-Partition)” on page 229                    |
| 5. Verify successful completion of upgrade to Sun Cluster 3.2 software.                                                                                                                                                                                                                                                                                 | “How to Verify Upgrade of Sun Cluster 3.2 Software” on page 249                           |
| 6. Enable resources and bring resource groups online. Optionally, migrate existing resources to new resource types.                                                                                                                                                                                                                                     | “How to Finish Upgrade to Sun Cluster 3.2 Software” on page 251                           |
| 7. (Optional) SPARC: Upgrade the Sun Cluster module for Sun Management Center, if needed.                                                                                                                                                                                                                                                               | “SPARC: How to Upgrade Sun Cluster Module Software for Sun Management Center” on page 189 |

## ▼ How to Prepare the Cluster for Upgrade (Dual-Partition)

Perform this procedure to prepare the cluster for a dual-partition upgrade. These procedures will refer to the two groups of nodes as the first partition and the second partition. The nodes that you assign to the second partition will continue cluster services while you upgrade the nodes in the first partition. After all nodes in the first partition are upgraded, you switch cluster services to the first partition and upgrade the second partition. After all nodes in the second partition are upgraded, you boot the nodes into cluster mode to rejoin the nodes from the first partition.

---

**Note** – If you are upgrading a single-node cluster, do not use this upgrade method. Instead, go to “How to Prepare the Cluster for Upgrade (Standard)” on page 197 or “How to Prepare the Cluster for Upgrade (Live Upgrade)” on page 237.

---

On the Solaris 10 OS, perform all steps from the global zone only.

**Before You Begin** Perform the following tasks:

- Ensure that the configuration meets the requirements for upgrade. See “Upgrade Requirements and Software Support Guidelines” on page 193.
- Have available the installation media, documentation, and patches for all software products that you are upgrading, including the following software:

- Solaris OS
- Sun Cluster 3.2 framework
- Sun Cluster 3.2 data services (agents)
- Applications that are managed by Sun Cluster 3.2 data-services
- VERITAS Volume Manager, if applicable

See “Patches and Required Firmware Levels” in *Sun Cluster 3.2 Release Notes for Solaris OS* for the location of patches and installation instructions.

- If you use role-based access control (RBAC) instead of superuser to access the cluster nodes, ensure that you can assume an RBAC role that provides authorization for all Sun Cluster commands. This series of upgrade procedures requires the following Sun Cluster RBAC authorizations if the user is not superuser:
  - `solaris.cluster.modify`
  - `solaris.cluster.admin`
  - `solaris.cluster.read`

See “Role-Based Access Control (Overview)” in *System Administration Guide: Security Services* for more information about using RBAC roles. See the Sun Cluster man pages for the RBAC authorization that each Sun Cluster subcommand requires.

**1 Ensure that the cluster is functioning normally.**

- a. View the current status of the cluster by running the following command from any node.**

```
% scstat
```

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

- b. Search the `/var/adm/messages` log on the same node for unresolved error messages or warning messages.**

- c. Check the volume-manager status.**

**2 If necessary, notify users that cluster services might be temporarily interrupted during the upgrade.**

Service interruption will be approximately the amount of time that your cluster normally takes to switch services to another node.

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

**4 For a two-node cluster that uses Sun StorEdge Availability Suite software or Sun StorageTek Availability Suite software, ensure that the configuration data for availability services resides on the quorum disk.**

The configuration data must reside on a quorum disk to ensure the proper functioning of Availability Suite after you upgrade the cluster software.

- a. Become superuser on a node of the cluster that runs Availability Suite software.**



**b. Identify the device ID and the slice that is used by the Availability Suite configuration file.**

```
phys-schost# /usr/opt/SUNWscm/sbin/dscfg
/dev/did/rdisk/dNsS
```

In this example output, *N* is the device ID and *S* the slice of device *N*.

**c. Identify the existing quorum device.**

```
phys-schost# scstat -q
-- Quorum Votes by Device --
 Device Name Present Possible Status

Device votes: /dev/did/rdisk/dQsS 1 1 Online
```

In this example output, *dQsS* is the existing quorum device.

**d. If the quorum device is not the same as the Availability Suite configuration-data device, move the configuration data to an available slice on the quorum device.**

```
phys-schost# dd if='/usr/opt/SUNWesm/sbin/dscfg' of=/dev/did/rdisk/dQsS
```

---

**Note** – You must use the name of the raw DID device, `/dev/did/rdisk/`, not the block DID device, `/dev/did/dsk/`.

---

**e. If you moved the configuration data, configure Availability Suite software to use the new location.**

As superuser, issue the following command on each node that runs Availability Suite software.

```
phys-schost# /usr/opt/SUNWesm/sbin/dscfg -s /dev/did/rdisk/dQsS
```

**5 If you will upgrade the Solaris OS and your cluster uses dual-string mediators for Solaris Volume Manager software, unconfigure your mediators.**

See [“Configuring Dual-String Mediators” on page 155](#) for more information about mediators.

**a. Run the following command to verify that no mediator data problems exist.**

```
phys-schost# medstat -s setname
```

`-s setname` Specifies the disk set name.

If the value in the Status field is `Bad`, repair the affected mediator host. Follow the procedure [“How to Fix Bad Mediator Data” on page 157](#).

**b. List all mediators.**

Save this information for when you restore the mediators during the procedure [“How to Finish Upgrade to Sun Cluster 3.2 Software” on page 251](#).

- c. **For a disk set that uses mediators, take ownership of the disk set if no node already has ownership.**

```
phys-schost# scswitch -z -D setname -h node
```

- z                      Changes mastery.
- D *devicegroup*      Specifies the name of the disk set.
- h *node*                Specifies the name of the node to become primary of the disk set.

- d. **Unconfigure all mediators for the disk set.**

```
phys-schost# metaset -s setname -d -m mediator-host-list
```

- s *setname*             Specifies the disk set name.
- d                      Deletes from the disk set.
- m *mediator-host-list* Specifies the name of the node to remove as a mediator host for the disk set.

See the `mediator(7D)` man page for further information about mediator-specific options to the `metaset` command.

- e. **Repeat Step c through Step d for each remaining disk set that uses mediators.**

- 6 If you are running the Sun Cluster HA for Sun Java System Application Server EE (HADB) data service with Sun Java System Application Server EE (HADB) software as of version 4.4, disable the HADB resource and shut down the HADB database.**

If you are running a version of Sun Java System Application Server EE (HADB) software before 4.4, you can skip this step.

When one cluster partition is out of service during upgrade, there are not enough nodes in the active partition to meet HADB membership requirements. Therefore, you must stop the HADB database and disable the HADB resource before you begin to partition the cluster.

```
phys-schost# hadbm stop database-name
phys-schost# clresource disable hadb-resource
```

For more information, see the `hadbm(1m)` man page.

- 7 If you are upgrading a two-node cluster, skip to Step 16.**

Otherwise, proceed to [Step 8](#) to determine the partitioning scheme to use. You will determine which nodes each partition will contain, but interrupt the partitioning process. You will then compare the node lists of all resource groups against the node members of each partition in the scheme that you will use. If any resource group does not contain a member of each partition, you must change the node list.

**8 Load the Sun Java Availability Suite DVD-ROM into the DVD-ROM drive.**

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

**9 Change to the `Solaris_arch/Product/sun_cluster/Solaris_ver/Tools/` directory, where `arch` is `sparc` or `x86` (Solaris 10 only) and where `ver` is 9 for Solaris 9 or 10 for Solaris 10 .**

```
phys-schost# cd /cdrom/cdrom0/Solaris_arch/Product/sun_cluster/Solaris_ver/Tools
```

**10 Start the `scinstall` utility in interactive mode.**

```
phys-schost# ./scinstall
```

---

**Note** – Do not use the `/usr/cluster/bin/scinstall` command that is already installed on the node. You must use the `scinstall` command on the Sun Java Availability Suite DVD-ROM.

---

The `scinstall` Main Menu is displayed.

**11 Type the number that corresponds to the option for Manage a dual-partition upgrade and press the Return key.**

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

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

- ```

1) Create a new cluster or add a cluster node
2) Configure a cluster to be JumpStarted from this install server
* 3) Manage a dual-partition upgrade
* 4) Upgrade this cluster node
* 5) Print release information for this cluster node

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

```
Option: 3
```

The Manage a Dual-Partition Upgrade Menu is displayed.

12 Type the number that corresponds to the option for Display and select possible partitioning schemes and press the Return key.**13 Follow the prompts to perform the following tasks:**

- a. Display the possible partitioning schemes for your cluster.
- b. Choose a partitioning scheme.
- c. Choose which partition to upgrade first.

Note – Stop and do *not* respond yet when prompted, Do you want to begin the dual-partition upgrade?, but do not exit the `scinstall` utility. You will respond to this prompt in [Step 18](#) of this procedure.

14 Make note of which nodes belong to each partition in the partition scheme.

15 On another node of the cluster, become superuser.

16 Ensure that any critical data services can switch over between partitions.

For a two-node cluster, each node will be the only node in its partition.

When the nodes of a partition are shut down in preparation for dual-partition upgrade, the resource groups that are hosted on those nodes switch over to a node in the other partition. If a resource group does not contain a node from each partition in its node list, the resource group cannot switch over. To ensure successful switchover of all critical data services, verify that the node list of the related resource groups contains a member of each upgrade partition.

a. Display the node list of each resource group that you require to remain in service during the entire upgrade.

```
phys-schost# scrgadm -pv -g resourcegroup | grep "Res Group Nodelist"
```

- p Displays configuration information.
- v Displays in verbose mode.
- g *resourcegroup* Specifies the name of the resource group.

b. If the node list of a resource group does not contain at least one member of each partition, redefine the node list to include a member of each partition as a potential primary node.

```
phys-schost# scrgadm -a -g resourcegroup -h nodelist
```

- a Adds a new configuration.
- h Specifies a comma-separated list of node names.

17 Determine your next step.

- **If you are upgrading a two-node cluster, return to [Step 8](#) through [Step 13](#) to designate your partitioning scheme and upgrade order.**

When you reach the prompt Do you want to begin the dual-partition upgrade?, skip to [Step 18](#).

- **If you are upgrading a cluster with three or more nodes, return to the node that is running the interactive `scinstall` utility.**

Proceed to [Step 18](#).

- 18 At the interactive `scinstall` prompt** Do you want to begin the dual-partition upgrade?, type **Yes**.

The command verifies that a remote installation method is available.

- 19 When prompted, press Enter to continue each stage of preparation for dual-partition upgrade.**

The command switches resource groups to nodes in the second partition, and then shuts down each node in the first partition.

- 20 After all nodes in the first partition are shut down, boot each node in that partition into noncluster mode.**

- **On SPARC based systems, perform the following command:**

```
ok boot -x
```

- **On x86 based systems running the Solaris 9 OS, perform either of the following commands:**

```
phys-schost# reboot -- -xs
```

or

...

```
<<< Current Boot Parameters >>>
```

```
Boot path: /pci@0,0/pci-ide@7,1/ata@1/cmdk@0,0:b
```

```
Boot args:
```

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

```
<<< timeout in 5 seconds >>>
```

```
Select (b)oot or (i)nterpreter: b -xs
```

- **On x86 based systems running the Solaris 10 OS, perform the following commands:**

- a. **In the GRUB menu, use the arrow keys to select the appropriate Solaris entry and type `e` to edit its commands.**

The GRUB menu appears similar to the following:

```
GNU GRUB version 0.95 (631K lower / 2095488K upper memory)
```

```
+-----+
| Solaris 10 /sol_10_x86 |
| Solaris failsafe |
| |
+-----+
```

Use the `^` and `v` keys to select which entry is highlighted.

Press enter to boot the selected OS, `'e'` to edit the commands before booting, or `'c'` for a command-line.

For more information about GRUB based booting, see Chapter 11, “GRUB Based Booting (Tasks),” in *System Administration Guide: Basic Administration*.

- b. In the boot parameters screen, use the arrow keys to select the kernel entry and type e to edit the entry.**

The GRUB boot parameters screen appears similar to the following:

```
GNU GRUB version 0.95 (615K lower / 2095552K upper memory)
+-----+
| root (hd0,0,a)                                     |
| kernel /platform/i86pc/multiboot                  |
| module /platform/i86pc/boot_archive                |
+-----+
```

Use the ^ and v keys to select which entry is highlighted. Press 'b' to boot, 'e' to edit the selected command in the boot sequence, 'c' for a command-line, 'o' to open a new line after ('O' for before) the selected line, 'd' to remove the selected line, or escape to go back to the main menu.

- c. Add -x to the command to specify that the system boot into noncluster mode.**

[Minimal BASH-like line editing is supported. For the first word, TAB lists possible command completions. Anywhere else TAB lists the possible completions of a device/filename. ESC at any time exits.]

```
grub edit> kernel /platform/i86pc/multiboot -x
```

- d. Press Enter to accept the change and return to the boot parameters screen.**

The screen displays the edited command.

```
GNU GRUB version 0.95 (615K lower / 2095552K upper memory)
+-----+
| root (hd0,0,a)                                     |
| kernel /platform/i86pc/multiboot -x                |
| module /platform/i86pc/boot_archive                |
+-----+
```

Use the ^ and v keys to select which entry is highlighted. Press 'b' to boot, 'e' to edit the selected command in the boot sequence, 'c' for a command-line, 'o' to open a new line after ('O' for before) the selected line, 'd' to remove the selected line, or escape to go back to the main menu.-

- e. Type b to boot the node into noncluster mode.**

Note – This change to the kernel boot parameter command does not persist over the system boot. The next time you reboot the node, it will boot into cluster mode. To boot into noncluster mode instead, perform these steps to again to add the -x option to the kernel boot parameter command.

21 If any applications that are running in the second partition are not under control of the Resource Group Manager (RGM), create scripts to halt the applications before you begin to upgrade those nodes.

During dual-partition upgrade processing, these scripts would be called to stop applications such as Oracle RAC before the nodes in the second partition are halted.

a. Create the scripts that you need to stop applications that are not under RGM control.

- Create separate scripts for those applications that you want stopped before applications under RGM control are stopped and for those applications that you want stop afterwards.
- To stop applications that are running on more than one node in the partition, write the scripts accordingly.
- Use any name and directory path for your scripts that you prefer.

b. Ensure that each node in the cluster has its own copy of your scripts.

c. On each node, modify the following Sun Cluster scripts to call the scripts that you placed on that node.

- `/etc/cluster/ql/cluster_pre_halt_apps` - Use this file to call those scripts that you want to run *before* applications that are under RGM control are shut down.
- `/etc/cluster/ql/cluster_post_halt_apps` - Use this file to call those scripts that you want to run *after* applications that are under RGM control are shut down.

The Sun Cluster scripts are issued from one arbitrary node in the partition during post-upgrade processing of the partition. Therefore, ensure that the scripts on any node of the partition will perform the necessary actions for all nodes in the partition.

Next Steps Upgrade software on each node in the first partition.

- To upgrade Solaris software before you perform Sun Cluster software upgrade, go to [“How to Upgrade the Solaris OS and Volume Manager Software \(Dual-Partition\)”](#) on page 224.
- If Sun Cluster 3.2 software does not support the release of the Solaris OS that you currently run on your cluster, you must upgrade the Solaris software to a supported release. See “Supported Products” in *Sun Cluster 3.2 Release Notes for Solaris OS* for more information.
- If Sun Cluster 3.2 software supports the release of the Solaris OS that you currently run on your cluster, further Solaris software upgrade is optional.

- Otherwise, upgrade to Sun Cluster 3.2 software. Go to [“How to Upgrade Sun Cluster 3.2 Software \(Dual-Partition\)”](#) on page 229.

▼ How to Upgrade the Solaris OS and Volume Manager Software (Dual-Partition)

Perform this procedure on each node in the cluster to upgrade the Solaris OS. On the Solaris 10 OS, perform all steps from the global zone only. If the cluster already runs on a version of the Solaris OS that supports Sun Cluster 3.2 software, further upgrade of the Solaris OS is optional. If you do not intend to upgrade the Solaris OS, proceed to [“How to Upgrade Sun Cluster 3.2 Software \(Standard\)”](#) on page 207.

Note – The cluster must already run on, or be upgraded to, at least the minimum required level of the Solaris OS to support upgrade to Sun Cluster 3.2 software. See Supported Products in *Sun Cluster 3.2 Release Notes for Solaris OS* for more information.

Before You Begin Ensure that all steps in [“How to Prepare the Cluster for Upgrade \(Standard\)”](#) on page 197 are completed.

1 Become superuser on the cluster node to upgrade.

The node must be a member of the partition that is in noncluster mode.

2 If Sun Cluster Geographic Edition software is installed, uninstall it.

For uninstallation procedures, see the documentation for your version of Sun Cluster Geographic Edition software.

3 Determine whether the following Apache run-control scripts exist and are enabled or disabled:

```
/etc/rc0.d/K16apache  
/etc/rc1.d/K16apache  
/etc/rc2.d/K16apache  
/etc/rc3.d/S50apache  
/etc/rc5.d/K16apache
```

Some applications, such as Sun Cluster HA for Apache, require that Apache run control scripts be disabled.

- If these scripts exist and contain an uppercase K or S in the file name, the scripts are enabled. No further action is necessary for these scripts.
- If these scripts do not exist, in [Step 8](#) you must ensure that any Apache run control scripts that are installed during the Solaris OS upgrade are disabled.

- If these scripts exist but the file names contain a lowercase k or s, the scripts are disabled. In [Step 8](#) you must ensure that any Apache run control scripts that are installed during the Solaris OS upgrade are disabled.

4 Comment out all entries for globally mounted file systems in the node's `/etc/vfstab` file.

- a. For later reference, make a record of all entries that are already commented out.
- b. Temporarily comment out all entries for globally mounted file systems in the `/etc/vfstab` file.
Entries for globally mounted file systems contain the `global` mount option. Comment out these entries to prevent the Solaris upgrade from attempting to mount the global devices.

5 Determine which procedure to follow to upgrade the Solaris OS.

Volume Manager	Procedure	Location of Instructions
Solaris Volume Manager	Any Solaris upgrade method <i>except</i> the Live Upgrade method	Solaris installation documentation
VERITAS Volume Manager	“Upgrading VxVM and Solaris”	VERITAS Volume Manager installation documentation

Note – If your cluster has VxVM installed, you must reinstall the existing VxVM software or upgrade to the Solaris 9 or 10 version of VxVM software as part of the Solaris upgrade process.

6 Upgrade the Solaris software, following the procedure that you selected in [Step 5](#).

- a. When prompted, choose the manual reboot option.
- b. When prompted to reboot, always reboot into noncluster mode.

Note – *Do not* perform the final reboot instruction in the Solaris software upgrade. Instead, do the following:

- a. Return to this procedure to perform [Step 7](#) and [Step 8](#).
 - b. Reboot into noncluster mode in [Step 9](#) to complete Solaris software upgrade.
-

Execute the following commands to boot a node into noncluster mode during Solaris upgrade:

- On SPARC based systems, perform either of the following commands:

```
phys-schost# reboot -- -xs
or
ok boot -xs
```

If the instruction says to run the `init S` command, use the `reboot -- -xs` command instead.

- **On x86 based systems, perform the following command:**

```
phys-schost# shutdown -g -y -i0
```

Press any key to continue

- i. **In the GRUB menu, use the arrow keys to select the appropriate Solaris entry and type `e` to edit its commands.**

The GRUB menu appears similar to the following:

```
GNU GRUB version 0.95 (631K lower / 2095488K upper memory)
```

```
+-----+
| Solaris 10 /sol_10_x86                               |
| Solaris failsafe                                     |
|                                                       |
+-----+
```

Use the `^` and `v` keys to select which entry is highlighted.

Press `enter` to boot the selected OS, `'e'` to edit the commands before booting, or `'c'` for a command-line.

For more information about GRUB based booting, see Chapter 11, “GRUB Based Booting (Tasks),” in *System Administration Guide: Basic Administration*.

- ii. **In the boot parameters screen, use the arrow keys to select the kernel entry and type `e` to edit the entry.**

The GRUB boot parameters screen appears similar to the following:

```
GNU GRUB version 0.95 (615K lower / 2095552K upper memory)
```

```
+-----+
| root (hd0,0,a)                                       |
| kernel /platform/i86pc/multiboot                   |
| module /platform/i86pc/boot_archive                 |
+-----+
```

Use the `^` and `v` keys to select which entry is highlighted.

Press `'b'` to boot, `'e'` to edit the selected command in the boot sequence, `'c'` for a command-line, `'o'` to open a new line after (`'O'` for before) the selected line, `'d'` to remove the selected line, or `escape` to go back to the main menu.

- iii. **Add `-x` to the command to specify that the system boot into noncluster mode.**

[Minimal BASH-like line editing is supported. For the first word, `TAB` lists possible command completions. Anywhere else `TAB` lists the possible completions of a device/filename. `ESC` at any time exits.]

```
grub edit> kernel /platform/i86pc/multiboot -x
```

iv. Press Enter to accept the change and return to the boot parameters screen.

The screen displays the edited command.

```
GNU GRUB version 0.95 (615K lower / 2095552K upper memory)
+-----+
| root (hd0,0,a)                                     |
| kernel /platform/i86pc/multiboot -x                |
| module /platform/i86pc/boot_archive                |
+-----+
```

Use the ^ and v keys to select which entry is highlighted. Press 'b' to boot, 'e' to edit the selected command in the boot sequence, 'c' for a command-line, 'o' to open a new line after ('O' for before) the selected line, 'd' to remove the selected line, or escape to go back to the main menu.-

v. Type b to boot the node into noncluster mode.

Note – This change to the kernel boot parameter command does not persist over the system boot. The next time you reboot the node, it will boot into cluster mode. To boot into noncluster mode instead, perform these steps to again to add the -x option to the kernel boot parameter command.

If the instruction says to run the `init 5` command, shut down the system then change the GRUB kernel boot command to `/platform/i86pc/multiboot -sx` instead.

7 In the `/a/etc/vfstab` file, uncomment those entries for globally mounted file systems that you commented out in [Step 4](#).**8 If Apache run control scripts were disabled or did not exist before you upgraded the Solaris OS, ensure that any scripts that were installed during Solaris upgrade are disabled.**

To disable Apache run control scripts, use the following commands to rename the files with a lowercase k or s.

```
phys-schost# mv /a/etc/rc0.d/K16apache /a/etc/rc0.d/k16apache
phys-schost# mv /a/etc/rc1.d/K16apache /a/etc/rc1.d/k16apache
phys-schost# mv /a/etc/rc2.d/K16apache /a/etc/rc2.d/k16apache
phys-schost# mv /a/etc/rc3.d/S50apache /a/etc/rc3.d/s50apache
phys-schost# mv /a/etc/rc5.d/K16apache /a/etc/rc5.d/k16apache
```

Alternatively, you can rename the scripts to be consistent with your normal administration practices.

9 Reboot the node into noncluster mode.

- **On SPARC based systems, perform the following command.**

Include the double dashes (--) in the command:

```
phys-schost# reboot -- -x
```

- On x86 based systems, perform the shutdown and boot procedures that are described in [Step 6](#) except add `-x` to the kernel boot command instead of `-sx`.

10 If your cluster runs VxVM, perform the remaining steps in the procedure “Upgrading VxVM and Solaris” to reinstall or upgrade VxVM.

Make the following changes to the procedure:

- After VxVM upgrade is complete but before you reboot, verify the entries in the `/etc/vfstab` file. If any of the entries that you uncommented in [Step 7](#) were commented out, make those entries uncommented again.
- When the VxVM procedures instruct you to perform a final reconfiguration reboot, *do not* use the `-r` option alone. Instead, reboot into noncluster mode by using the `-rx` options.
 - On SPARC based systems, perform the following command:

```
phys-schost# reboot -- -rx
```
 - On x86 based systems, perform the shutdown and boot procedures that are described in [Step 6](#) except add `-rx` to the kernel boot command instead of `-sx`.

Note – If you see a message similar to the following, type the root password to continue upgrade processing. Do *not* run the `fsck` command nor type `Ctrl-D`.

```
WARNING - Unable to repair the /global/.devices/node@1 filesystem.  
Run fsck manually (fsck -F ufs /dev/vx/rdisk/rootdisk_13vol). Exit the  
shell when done to continue the boot process.
```

```
Type control-d to proceed with normal startup,  
(or give root password for system maintenance): Type the root password
```

11 (Optional) SPARC: Upgrade VxFS.

Follow procedures that are provided in your VxFS documentation.

12 Install any required Solaris software patches and hardware-related patches, and download any needed firmware that is contained in the hardware patches.

Note – Do not reboot after you add patches. Wait to reboot the node until after you upgrade the Sun Cluster software.

See “Patches and Required Firmware Levels” in *Sun Cluster 3.2 Release Notes for Solaris OS* for the location of patches and installation instructions.

Next Steps Upgrade to Sun Cluster 3.2 software. Go to “[How to Upgrade Sun Cluster 3.2 Software \(Dual-Partition\)](#)” on page 229.

Note – To complete the upgrade to a new marketing release of the Solaris OS, such as from Solaris 9 to Solaris 10 software, you must also upgrade the Sun Cluster software and dependency software to the version that is compatible with the new version of the Solaris OS.

▼ How to Upgrade Sun Cluster 3.2 Software (Dual-Partition)

Perform this procedure to upgrade each node of the cluster to Sun Cluster 3.2 software. This procedure also upgrades required Sun Java Enterprise System shared components. You must also perform this procedure after you upgrade to a different marketing release of the Solaris OS, such as from Solaris 9 to Solaris 10 software.

On the Solaris 10 OS, perform all steps from the global zone only.

Tip – You can perform this procedure on more than one node of the partition at the same time.

Before You Begin Perform the following tasks:

- Ensure that all steps in “[How to Prepare the Cluster for Upgrade \(Dual-Partition\)](#)” on page 215 are completed.
- Ensure that the node you are upgrading belongs to the partition that is not active in the cluster and that the node is in noncluster mode.
- If you upgraded to a new marketing release of the Solaris OS, such as from Solaris 9 to Solaris 10 software, ensure that all steps in “[How to Upgrade the Solaris OS and Volume Manager Software \(Dual-Partition\)](#)” on page 224 are completed.
- Ensure that you have installed all required Solaris software patches and hardware-related patches.

1 Become superuser on a node that is a member of the partition that is in noncluster mode.

2 Ensure that the `/usr/java/` directory is a symbolic link to the minimum or latest version of Java software.

Sun Cluster software requires at least version 1.5.0_06 of Java software. If you upgraded to a version of Solaris that installs an earlier version of Java, the upgrade might have changed the symbolic link to point to a version of Java that does not meet the minimum requirement for Sun Cluster 3.2 software.

a. Determine what directory the `/usr/java/` directory is symbolically linked to.

```
phys-schost# ls -l /usr/java
lrwxrwxrwx 1 root  other   9 Apr 19 14:05 /usr/java -> /usr/j2se/
```

b. Determine what version or versions of Java software are installed.

The following are examples of commands that you can use to display the version of their related releases of Java software.

```
phys-schost# /usr/j2se/bin/java -version
phys-schost# /usr/java1.2/bin/java -version
phys-schost# /usr/jdk/jdk1.5.0_06/bin/java -version
```

c. If the /usr/java/ directory is not symbolically linked to a supported version of Java software, recreate the symbolic link to link to a supported version of Java software.

The following example shows the creation of a symbolic link to the /usr/j2se/ directory, which contains Java 1.5.0_06 software.

```
phys-schost# rm /usr/java
phys-schost# ln -s /usr/j2se /usr/java
```

3 Load the Sun Java Availability Suite DVD-ROM into the DVD-ROM drive.

If the volume management daemon `vol1d(1M)` is running and is configured to manage CD-ROM or DVD devices, the daemon automatically mounts the media on the `/cdrom/cdrom0/` directory.

4 Change to the installation wizard directory of the DVD-ROM.

- If you are installing the software packages on the SPARC platform, type the following command:

```
phys-schost# cd /cdrom/cdrom0//Solaris_sparc
```

- If you are installing the software packages on the x86 platform, type the following command:

```
phys-schost# cd /cdrom/cdrom0//Solaris_x86
```

5 Start the installation wizard program.

```
phys-schost# ./installer
```

6 Follow the instructions on the screen to select and upgrade Shared Components software packages on the node.

Note – Do *not* use the installation wizard program to upgrade Sun Cluster software packages.

The installation wizard program displays the status of the installation. When the installation is complete, the program displays an installation summary and the installation logs.

7 Exit the installation wizard program.**8 Change to the `Solaris_arch/Product/sun_cluster/Solaris_ver/Tools/` directory, where `arch` is `sparc` or `x86` (Solaris 10 only) and where `ver` is 9 for Solaris 9 or 10 for Solaris 10.**

```
phys-schost# cd /cdrom/cdrom0/Solaris_arch/Product/sun_cluster/Solaris_ver/Tools
```

9 Start the `scinstall` utility.

```
phys-schost# ./scinstall
```

Note – Do not use the `/usr/cluster/bin/scinstall` command that is already installed on the node. You must use the `scinstall` command that is located on the Sun Java Availability Suite DVD-ROM.

The `scinstall` Main Menu is displayed.

10 Type the number that corresponds to the option for Upgrade this cluster node and press the Return key.

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

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

- ```
* 1) Create a new cluster or add a cluster node
 2) Configure a cluster to be JumpStarted from this install server
* 3) Manage a dual-partition upgrade
* 4) Upgrade this cluster node
* 5) Print release information for this cluster node

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

```
Option: 4
```

The Upgrade Menu is displayed.

**11 Type the number that corresponds to the option for Upgrade Sun Cluster framework on this cluster node and press the Return key.****12 Follow the menu prompts to upgrade the cluster framework.**

During the Sun Cluster upgrade, `scinstall` might make one or more of the following configuration changes:

- Convert NAFO groups to IPMP groups but keep the original NAFO-group name. See one of the following manuals for information about test addresses for IPMP:
  - “Configuring Test Addresses” in “Administering Multipathing Groups With Multiple Physical Interfaces” in *System Administration Guide: IP Services* (Solaris 9)
  - “Test Addresses” in *System Administration Guide: IP Services* (Solaris 10)

See the `scinstall(1M)` man page for more information about the conversion of NAFO groups to IPMP during Sun Cluster software upgrade.

- Rename the `ntp.conf` file to `ntp.conf.cluster`, if `ntp.conf.cluster` does not already exist on the node.
- Set the `local-mac-address?` variable to `true`, if the variable is not already set to that value.

Upgrade processing is finished when the system displays the message `Completed Sun Cluster framework upgrade` and prompts you to press `Enter` to continue.

**13 Quit the `scinstall` utility.**

**14 Unload the Sun Java Availability Suite DVD-ROM from the DVD-ROM drive.**

a. To ensure that the DVD-ROM is not being used, change to a directory that does *not* reside on the DVD-ROM.

b. Eject the DVD-ROM.

```
phys-schost# eject cdrom
```

**15 Upgrade data service packages.**

You must upgrade all data services to the Sun Cluster 3.2 version.

---

**Note** – For Sun Cluster HA for SAP Web Application Server, if you are using a J2EE engine resource or a web application server component resource or both, you must delete the resource and recreate it with the new web application server component resource. Changes in the new web application server component resource includes integration of the J2EE functionality. For more information, see *Sun Cluster Data Service for SAP Web Application Server Guide for Solaris OS*.

---

a. **Start the upgraded interactive `scinstall` utility.**

```
phys-schost# /usr/cluster/bin/scinstall
```

---

**Note** – Do not use the `scinstall` utility that is on the installation media to upgrade data service packages.

---

The `scinstall` Main Menu is displayed.

b. **Type the number that corresponds to the option for Upgrade this cluster node and press the Return key.**

The Upgrade Menu is displayed.

c. **Type the number that corresponds to the option for Upgrade Sun Cluster data service agents on this node and press the Return key.**

d. **Follow the menu prompts to upgrade Sun Cluster data service agents that are installed on the node.**

You can choose from the list of data services that are available to upgrade or choose to upgrade all installed data services.

Upgrade processing is finished when the system displays the message `Completed upgrade of Sun Cluster data services agents` and prompts you to press `Enter` to continue.



**e. Press Enter.**

The Upgrade Menu is displayed.

**16 Quit the `scinstall` utility.****17 If you have Sun Cluster HA for NFS configured on a highly available local file system, ensure that the loopback file system (LOFS) is disabled.**


---

**Note** – If you have non-global zones configured, LOFS must remain enabled. For guidelines about using LOFS and alternatives to disabling it, see “Cluster File Systems” on page 32.

---

As of the Sun Cluster 3.2 release, LOFS is no longer disabled by default during Sun Cluster software installation or upgrade. To disable LOFS, ensure that the `/etc/system` file contains the following entry:

```
exclude:lofs
```

This change becomes effective at the next system reboot.

**18 As needed, manually upgrade any custom data services that are not supplied on the product media.****19 Verify that each data-service update is installed successfully.**

View the upgrade log file that is referenced at the end of the upgrade output messages.

**20 Install any Sun Cluster 3.2 framework and data-service software patches.**

See “Patches and Required Firmware Levels” in *Sun Cluster 3.2 Release Notes for Solaris OS* for the location of patches and installation instructions.

**21 Upgrade software applications that are installed on the cluster.**

Ensure that application levels are compatible with the current versions of Sun Cluster and Solaris software. See your application documentation for installation instructions.

**22 After all nodes in a partition are upgraded, apply the upgrade changes.****a. From one node in the partition that you are upgrading, start the interactive `scinstall` utility.**

```
phys-schost# /usr/cluster/bin/scinstall
```

---

**Note** – Do not use the `scinstall` command that is located on the installation media. Only use the `scinstall` command that is located on the cluster node.

---

The `scinstall` Main Menu is displayed.

**b. Type the number that corresponds to the option for Apply dual-partition upgrade changes to the partition and press the Return key.**

**c. Follow the prompts to continue each stage of the upgrade processing.**

The command performs the following tasks, depending on which partition the command is run from:

- **First partition** - The command halts each node in the second partition, one node at a time. When a node is halted, any services on that node are automatically switched over to a node in the first partition, provided that the node list of the related resource group contains a node in the first partition. After all nodes in the second partition are halted, the nodes in the first partition are booted into cluster mode and take over providing cluster services.
- **Second partition** - The command boots the nodes in the second partition into cluster mode, to join the active cluster that was formed by the first partition. After all nodes have rejoined the cluster, the command performs final processing and reports on the status of the upgrade.

**d. Exit the `scinstall` utility, if it is still running.**

**e. If you are finishing upgrade of the first partition, perform the following substeps to prepare the second partition for upgrade.**

Otherwise, if you are finishing upgrade of the second partition, proceed to [“How to Verify Upgrade of Sun Cluster 3.2 Software”](#) on page 249.

**i. Boot each node in the second partition into noncluster mode.**

- **On SPARC based systems, perform the following command:**  
`ok boot -x`
- **On x86 based systems, perform the following commands:**

**In the GRUB menu, use the arrow keys to select the appropriate Solaris entry and type `e` to edit its commands.**

The GRUB menu appears similar to the following:

```
GNU GRUB version 0.95 (631K lower / 2095488K upper memory)
+-----+
| Solaris 10 /sol_10_x86 |
| Solaris failsafe |
| |
+-----+

```

Use the `^` and `v` keys to select which entry is highlighted.  
 Press enter to boot the selected OS, `'e'` to edit the  
 commands before booting, or `'c'` for a command-line.

For more information about GRUB based booting, see Chapter 11, “GRUB Based Booting (Tasks),” in *System Administration Guide: Basic Administration*.

**In the boot parameters screen, use the arrow keys to select the kernel entry and type e to edit the entry.**

The GRUB boot parameters screen appears similar to the following:

```
GNU GRUB version 0.95 (615K lower / 2095552K upper memory)
+-----+
| root (hd0,0,a) |
| kernel /platform/i86pc/multiboot |
| module /platform/i86pc/boot_archive |
+-----+
```

Use the ^ and v keys to select which entry is highlighted. Press 'b' to boot, 'e' to edit the selected command in the boot sequence, 'c' for a command-line, 'o' to open a new line after ('O' for before) the selected line, 'd' to remove the selected line, or escape to go back to the main menu.

**Add -x to the command to specify that the system boot into noncluster mode.**

[ Minimal BASH-like line editing is supported. For the first word, TAB lists possible command completions. Anywhere else TAB lists the possible completions of a device/filename. ESC at any time exits. ]

```
grub edit> kernel /platform/i86pc/multiboot -x
```

**Press Enter to accept the change and return to the boot parameters screen.**

The screen displays the edited command.

```
GNU GRUB version 0.95 (615K lower / 2095552K upper memory)
+-----+
| root (hd0,0,a) |
| kernel /platform/i86pc/multiboot -x |
| module /platform/i86pc/boot_archive |
+-----+
```

Use the ^ and v keys to select which entry is highlighted. Press 'b' to boot, 'e' to edit the selected command in the boot sequence, 'c' for a command-line, 'o' to open a new line after ('O' for before) the selected line, 'd' to remove the selected line, or escape to go back to the main menu.-

**Type b to boot the node into noncluster mode.**

---

**Note** – This change to the kernel boot parameter command does not persist over the system boot. The next time you reboot the node, it will boot into cluster mode. To boot into noncluster mode instead, perform these steps to again to add the -x option to the kernel boot parameter command.

---

## ii. Upgrade the nodes in the second partition.

To upgrade Solaris software before you perform Sun Cluster software upgrade, go to “How to Upgrade the Solaris OS and Volume Manager Software (Dual-Partition)” on page 224.

Otherwise, upgrade Sun Cluster software on the second partition. Return to [Step 1](#).

**Next Steps** Go to [“How to Verify Upgrade of Sun Cluster 3.2 Software”](#) on page 249.

**Troubleshooting** If you experience an unrecoverable error during dual-partition upgrade, perform recovery procedures in [“How to Recover from a Failed Dual-Partition Upgrade”](#) on page 258.

## Performing a Live Upgrade to Sun Cluster 3.2 Software

This section provides the following information to upgrade from Sun Cluster 3.1 software to Sun Cluster 3.2 software by using the live upgrade method:

- [“How to Prepare the Cluster for Upgrade \(Live Upgrade\)”](#) on page 237
- [“How to Upgrade the Solaris OS and Sun Cluster 3.2 Software \(Live Upgrade\)”](#) on page 239

The following table lists the tasks to perform to upgrade from Sun Cluster 3.1 software to Sun Cluster 3.2 software. You also perform these tasks to upgrade only the version of the Solaris OS. If you upgrade the Solaris OS from Solaris 9 to Solaris 10 software, you must also upgrade the Sun Cluster software and dependency software to the version that is compatible with the new version of the Solaris OS.

**TABLE 8-3** Task Map: Performing a Live Upgrade to Sun Cluster 3.2 Software

| Task                                                                                                                                                                                                                                                                                                                                                              | Instructions                                                                                                                                              |
|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------|
| 1. Read the upgrade requirements and restrictions. Determine the proper upgrade method for your configuration and needs.                                                                                                                                                                                                                                          | <a href="#">“Upgrade Requirements and Software Support Guidelines”</a> on page 193<br><a href="#">“Choosing a Sun Cluster Upgrade Method”</a> on page 195 |
| 2. Remove the cluster from production, disable resources, and back up shared data and system disks. If the cluster uses dual-string mediators for Solaris Volume Manager software, unconfigure the mediators.                                                                                                                                                     | <a href="#">“How to Prepare the Cluster for Upgrade (Live Upgrade)”</a> on page 237                                                                       |
| 3. Upgrade the Solaris software, if necessary, to a supported Solaris update. Upgrade to Sun Cluster 3.2 framework and data-service software. If necessary, upgrade applications. If the cluster uses dual-string mediators, reconfigure the mediators. As needed, upgrade VERITAS Volume Manager (VxVM) software and disk groups and VERITAS File System (VxFS). | <a href="#">“How to Upgrade the Solaris OS and Sun Cluster 3.2 Software (Live Upgrade)”</a> on page 239                                                   |
| 4. Verify successful completion of upgrade to Sun Cluster 3.2 software.                                                                                                                                                                                                                                                                                           | <a href="#">“How to Verify Upgrade of Sun Cluster 3.2 Software”</a> on page 249                                                                           |
| 5. Enable resources and bring resource groups online. Migrate existing resources to new resource types.                                                                                                                                                                                                                                                           | <a href="#">“How to Finish Upgrade to Sun Cluster 3.2 Software”</a> on page 251                                                                           |

TABLE 8-3 Task Map: Performing a Live Upgrade to Sun Cluster 3.2 Software (Continued)

| Task                                                                                      | Instructions                                                                              |
|-------------------------------------------------------------------------------------------|-------------------------------------------------------------------------------------------|
| 6. (Optional) SPARC: Upgrade the Sun Cluster module for Sun Management Center, if needed. | “SPARC: How to Upgrade Sun Cluster Module Software for Sun Management Center” on page 189 |

## ▼ How to Prepare the Cluster for Upgrade (Live Upgrade)

Perform this procedure to prepare a cluster for live upgrade.

**Before You Begin** Perform the following tasks:

- Ensure that the configuration meets the requirements for upgrade. See “[Upgrade Requirements and Software Support Guidelines](#)” on page 193.
- Have available the installation media, documentation, and patches for all software products that you are upgrading, including the following software:
  - Solaris OS
  - Sun Cluster 3.2 framework
  - Sun Cluster 3.2 data services (agents)
  - Applications that are managed by Sun Cluster 3.2 data-services
  - VERITAS Volume Manager, if applicable

See “Patches and Required Firmware Levels” in *Sun Cluster 3.2 Release Notes for Solaris OS* for the location of patches and installation instructions.

- If you use role-based access control (RBAC) instead of superuser to access the cluster nodes, ensure that you can assume an RBAC role that provides authorization for all Sun Cluster commands. This series of upgrade procedures requires the following Sun Cluster RBAC authorizations if the user is not superuser:
  - `solaris.cluster.modify`
  - `solaris.cluster.admin`
  - `solaris.cluster.read`

See “Role-Based Access Control (Overview)” in *System Administration Guide: Security Services* for more information about using RBAC roles. See the Sun Cluster man pages for the RBAC authorization that each Sun Cluster subcommand requires.

### 1 Ensure that the cluster is functioning normally.

- a. **View the current status of the cluster by running the following command from any node.**

```
phys-schost% scstat
```

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

b. Search the `/var/adm/messages` log on the same node for unresolved error messages or warning messages.

c. Check the volume-manager status.

2 If necessary, notify users that cluster services will be temporarily interrupted during the upgrade.

Service interruption will be approximately the amount of time that your cluster normally takes to switch services to another node.

3 Become superuser on a node of the cluster.

4 If Sun Cluster Geographic Edition software is installed, uninstall it.

For uninstallation procedures, see the documentation for your version of Sun Cluster Geographic Edition software.

5 For a two-node cluster that uses Sun StorEdge Availability Suite software or Sun StorageTek Availability Suite software, ensure that the configuration data for availability services resides on the quorum disk.

The configuration data must reside on a quorum disk to ensure the proper functioning of Availability Suite after you upgrade the cluster software.

a. Become superuser on a node of the cluster that runs Availability Suite software.

b. Identify the device ID and the slice that is used by the Availability Suite configuration file.

```
phys-schost# /usr/opt/SUNWscm/sbin/dscfg
/dev/did/rdisk/dNsS
```

In this example output, *N* is the device ID and *S* the slice of device *N*.

c. Identify the existing quorum device.

```
phys-schost# scstat -q
-- Quorum Votes by Device --
 Device Name Present Possible Status

Device votes: /dev/did/rdisk/dQsS 1 1 Online
```

In this example output, *dQsS* is the existing quorum device.

d. If the quorum device is not the same as the Availability Suite configuration-data device, move the configuration data to an available slice on the quorum device.

```
phys-schost# dd if='/usr/opt/SUNWscm/sbin/dscfg' of=/dev/did/rdisk/dQsS
```

---

**Note** – You must use the name of the raw DID device, `/dev/did/rdisk/`, not the block DID device, `/dev/did/dsk/`.

---

- e. **If you moved the configuration data, configure Availability Suite software to use the new location.**

As superuser, issue the following command on each node that runs Availability Suite software.

```
phys-schost# /usr/opt/SUNWesm/sbin/dscfg -s /dev/did/rdisk/dQsS
```

- 6 Ensure that all shared data is backed up.**

- 7 Ensure that each system disk is backed up.**

**Next Steps** Perform a live upgrade of the Solaris OS, Sun Cluster 3.2 software, and other software. Go to [“How to Upgrade the Solaris OS and Sun Cluster 3.2 Software \(Live Upgrade\)”](#) on page 239.

## ▼ How to Upgrade the Solaris OS and Sun Cluster 3.2 Software (Live Upgrade)

Perform this procedure to upgrade the Solaris OS, Java ES shared components, volume-manager software, and Sun Cluster software by using the live upgrade method. The Sun Cluster live upgrade method uses the Solaris Live Upgrade feature. For information about live upgrade of the Solaris OS, refer to the documentation for the Solaris version that you are using:

- Chapter 32, “Solaris Live Upgrade (Topics),” in *Solaris 9 9/04 Installation Guide*
- *Solaris 10 Installation Guide: Solaris Live Upgrade and Upgrade Planning*

---

**Note** – The cluster must already run on, or be upgraded to, at least the minimum required level of the Solaris OS to support upgrade to Sun Cluster 3.2 software. See Supported Products in *Sun Cluster 3.2 Release Notes for Solaris OS* for more information.

---

Perform this procedure on each node in the cluster.

---

**Tip** – You can use the `cconsole` utility to perform this procedure on all nodes simultaneously. See [“How to Install Cluster Control Panel Software on an Administrative Console”](#) on page 47 for more information.

---

**Before You Begin** Ensure that all steps in [“How to Prepare the Cluster for Upgrade \(Live Upgrade\)”](#) on page 237 are completed.

**1 Ensure that a supported version of Solaris Live Upgrade software is installed on each node.**

If your operating system is already upgraded to Solaris 9 9/05 software or Solaris 10 11/06 software, you have the correct Solaris Live Upgrade software. If your operating system is an older version, perform the following steps:

**a. Insert the Solaris 9 9/05 software or Solaris 10 11/06 software media.****b. Become superuser.****c. Install the SUNWLu0 and SUNWLur packages.**

```
phys-schost# pkgadd -d path SUNWLu0 SUNWLur
```

*path* Specifies the absolute path to the software packages.

**d. Verify that the packages have been installed.**

```
phys-schost# pkgchk -v SUNWLu0 SUNWLur
```

**2 If you will upgrade the Solaris OS and your cluster uses dual-string mediators for Solaris Volume Manager software, unconfigure your mediators.**

See “[Configuring Dual-String Mediators](#)” on page 155 for more information about mediators.

**a. Run the following command to verify that no mediator data problems exist.**

```
phys-schost# medstat -s setname
```

-s *setname* Specifies the disk set name.

If the value in the Status field is Bad, repair the affected mediator host. Follow the procedure “[How to Fix Bad Mediator Data](#)” on page 157.

**b. List all mediators.**

Save this information for when you restore the mediators during the procedure “[How to Finish Upgrade to Sun Cluster 3.2 Software](#)” on page 251.

**c. For a disk set that uses mediators, take ownership of the disk set if no node already has ownership.**

```
phys-schost# scswitch -z -D setname -h node
```

-z Changes mastery.

-D *devicegroup* Specifies the name of the disk set.

-h *node* Specifies the name of the node to become primary of the disk set.

**d. Unconfigure all mediators for the disk set.**

```
phys-schost# metaset -s setname -d -m mediator-host-list
```

-s *setname* Specifies the disk set name.



- d Deletes from the disk set.
- m *mediator-host-list* Specifies the name of the node to remove as a mediator host for the disk set.

See the `mediator(7D)` man page for further information about mediator-specific options to the `metaset` command.

**e. Repeat Step c through Step d for each remaining disk set that uses mediators.**

### 3 Build an inactive boot environment (BE).

```
phys-schost# lucreate options-n BE-name
```

- n *BE-name* Specifies the name of the boot environment that is to be upgraded.

For information about important options to the `lucreate` command, see *Solaris 10 Installation Guide: Solaris Live Upgrade and Upgrade Planning* and the `lucreate(1M)` man page.

### 4 If necessary, upgrade the Solaris OS software in your inactive BE.

If the cluster already runs on a properly patched version of the Solaris OS that supports Sun Cluster 3.2 software, this step is optional.

- **If you use Solaris Volume Manager software, run the following command:**

```
phys-schost# luupgrade -u -n BE-name -s os-image-path
```

- u Upgrades an operating system image on a boot environment.
- s *os-image-path* Specifies the path name of a directory that contains an operating system image.

- **If you use VERITAS Volume Manager, follow live upgrade procedures in your VxVM installation documentation.**

### 5 Mount your inactive BE by using the `lumount` command.

```
phys-schost# lumount -n BE-name -m BE-mount-point
```

- m *BE-mount-point* Specifies the mount point of *BE-name*.

For more information, see *Solaris 10 Installation Guide: Solaris Live Upgrade and Upgrade Planning* and the `lumount(1M)` man page.

**6 Ensure that the `/BE-mount-point/usr/java/` directory is a symbolic link to the minimum or latest version of Java software.**

Sun Cluster software requires at least version 1.5.0\_06 of Java software. If you upgraded to a version of Solaris that installs an earlier version of Java, the upgrade might have changed the symbolic link to point to a version of Java that does not meet the minimum requirement for Sun Cluster 3.2 software.

**a. Determine what directory the `/BE-mount-point/usr/java/` directory is symbolically linked to.**

```
phys-schost# ls -l /BE-mount-point/usr/java
lrwxrwxrwx 1 root other 9 Apr 19 14:05 /BE-mount-point/usr/java -> /BE-mount-point/usr/j2se/
```

**b. Determine what version or versions of Java software are installed.**

The following are examples of commands that you can use to display the version of their related releases of Java software.

```
phys-schost# /BE-mount-point/usr/j2se/bin/java -version
phys-schost# /BE-mount-point/usr/java1.2/bin/java -version
phys-schost# /BE-mount-point/usr/jdk/jdk1.5.0_06/bin/java -version
```

**c. If the `/BE-mount-point/usr/java/` directory is not symbolically linked to a supported version of Java software, recreate the symbolic link to link to a supported version of Java software.**

The following example shows the creation of a symbolic link to the `/usr/j2se/` directory, which contains Java 1.5.0\_06 software.

```
phys-schost# rm /BE-mount-point/usr/java
phys-schost# cd /mnt/usr
phys-schost# ln -s j2se java
```

**7 Apply any necessary Solaris patches.**

You might need to patch your Solaris software to use the Live Upgrade feature. For details about the patches that the Solaris OS requires and where to download them, see “Managing Packages and Patches With Solaris Live Upgrade” in *Solaris 9 9/04 Installation Guide* or “Upgrading a System With Packages or Patches” in *Solaris 10 Installation Guide: Solaris Live Upgrade and Upgrade Planning*.

**8 If necessary and if your version of the VERITAS Volume Manager (VxVM) software supports it, upgrade your VxVM software.**

Refer to your VxVM software documentation to determine whether your version of VxVM can use the live upgrade method.

**9 (Optional) SPARC: Upgrade VxFS.**

Follow procedures that are provided in your VxFS documentation.

**10 If your cluster hosts software applications that require an upgrade and that you can upgrade by using the live upgrade method, upgrade those software applications.**

If your cluster hosts software applications that cannot use the live upgrade method, you will upgrade them later in [Step 25](#).

**11 Load the Sun Java Availability Suite DVD-ROM into the DVD-ROM drive.**

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

**12 Change to the installation wizard directory of the DVD-ROM.**

- If you are installing the software packages on the SPARC platform, type the following command:

```
phys-schost# cd /cdrom/cdrom0/Solaris_sparc
```

- If you are installing the software packages on the x86 platform, type the following command:

```
phys-schost# cd /cdrom/cdrom0/Solaris_x86
```

**13 Start the installation wizard program to direct output to a state file.**

Specify the name to give the state file and the absolute or relative path where the file should be created.

- To create a state file by using the graphical interface, use the following command:

```
phys-schost# ./installer -no -saveState statefile
```

- To create a state file by using the text-based interface, use the following command:

```
phys-schost# ./installer -no -nodisplay -saveState statefile
```

See “Generating the Initial State File” in *Sun Java Enterprise System 5 Installation Guide for UNIX* for more information.

**14 Follow the instructions on the screen to select and upgrade Shared Components software packages on the node.**

The installation wizard program displays the status of the installation. When the installation is complete, the program displays an installation summary and the installation logs.

**15 Exit the installation wizard program.****16 Run the `installer` program in silent mode and direct the installation to the alternate boot environment.**

---

**Note** – The `installer` program must be the same version that you used to create the state file.

---

```
phys-schost# ./installer -nodisplay -noconsole -state statefile -altroot BE-mount-point
```

See “To Run the Installer in Silent Mode” in *Sun Java Enterprise System 5 Installation Guide for UNIX* for more information.

- 17 Change to the `Solaris_arch/Product/sun_cluster/Solaris_ver/Tools/` directory, where `arch` is `sparc` or `x86` (Solaris 10 only) and where `ver` is 9 for Solaris 9 or 10 for Solaris 10.**

```
phys-schost# cd /cdrom/cdrom0/Solaris_arch/Product/sun_cluster/Solaris_ver/Tools
```

- 18 Upgrade your Sun Cluster software by using the `scinstall` command.**

```
phys-schost# ./scinstall -u update -R BE-mount-point
```

`-u update` Specifies that you are performing an upgrade of Sun Cluster software.

`-R BE-mount-point` Specifies the mount point for your alternate boot environment.

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

- 19 Upgrade your data services by using the `scinstall` command.**

```
phys-schost# ./scinstall -u -s all -d /cdrom/cdrom0/ -R BE-mount-point
```

- 20 Unload the Sun Java Availability Suite DVD-ROM from the DVD-ROM drive.**

- a. To ensure that the DVD-ROM is not being used, change to a directory that does *not* reside on the DVD-ROM.

- b. Eject the DVD-ROM.

```
phys-schost# eject cdrom
```

- 21 Unmount the inactive BE.**

```
phys-schost# luumount -n BE-name
```

- 22 Activate the upgraded inactive BE.**

```
phys-schost# luactivate BE-name
```

`BE-name` The name of the alternate BE that you built in [Step 3](#).

- 23 Repeat [Step 1](#) through [Step 22](#) for each node in the cluster.**

---

**Note** – Do not reboot any node until all nodes in the cluster are upgraded on their inactive BE.

---

- 24 Reboot all nodes.**

```
phys-schost# shutdown -y -g0 -i6
```

---

**Note** – Do not use the `reboot` or `halt` command. These commands do not activate a new BE. Use only `shutdown` or `init` to reboot into a new BE.

---

The nodes reboot into cluster mode using the new, upgraded BE.

**25 (Optional) If your cluster hosts software applications that require upgrade for which you cannot use the live upgrade method, perform the following steps.**

**Note** – Throughout the process of software-application upgrade, always reboot into noncluster mode until all upgrades are complete.

**a. Shut down the node.**

```
phys-schost# shutdown -y -g0 -i0
```

**b. Boot each node into noncluster mode.**

▪ **On SPARC based systems, perform the following command:**

```
ok boot -x
```

▪ **On x86 based systems, perform the following commands:**

**i. In the GRUB menu, use the arrow keys to select the appropriate Solaris entry and type e to edit its commands.**

The GRUB menu appears similar to the following:

```
GNU GRUB version 0.95 (631K lower / 2095488K upper memory)
```

```
+-----+
| Solaris 10 /sol_10_x86 |
| Solaris failsafe |
| |
+-----+
```

Use the ^ and v keys to select which entry is highlighted.

Press enter to boot the selected OS, 'e' to edit the commands before booting, or 'c' for a command-line.

For more information about GRUB based booting, see Chapter 11, “GRUB Based Booting (Tasks),” in *System Administration Guide: Basic Administration*.

**ii. In the boot parameters screen, use the arrow keys to select the kernel entry and type e to edit the entry.**

The GRUB boot parameters screen appears similar to the following:

```
GNU GRUB version 0.95 (615K lower / 2095552K upper memory)
```

```
+-----+
| root (hd0,0,a) |
| kernel /platform/i86pc/multiboot |
| module /platform/i86pc/boot_archive |
+-----+
```

Use the ^ and v keys to select which entry is highlighted.

Press 'b' to boot, 'e' to edit the selected command in the boot sequence, 'c' for a command-line, 'o' to open a new line

after ('O' for before) the selected line, 'd' to remove the selected line, or escape to go back to the main menu.

**iii. Add -x to the command to specify that the system boot into noncluster mode.**

[ Minimal BASH-like line editing is supported. For the first word, TAB lists possible command completions. Anywhere else TAB lists the possible completions of a device/filename. ESC at any time exits. ]

```
grub edit> kernel /platform/i86pc/multiboot -x
```

**iv. Press Enter to accept the change and return to the boot parameters screen.**

The screen displays the edited command.

```
GNU GRUB version 0.95 (615K lower / 2095552K upper memory)
```

```
+-----+
| root (hd0,0,a) |
| kernel /platform/i86pc/multiboot -x |
| module /platform/i86pc/boot_archive |
+-----+
```

Use the ^ and v keys to select which entry is highlighted. Press 'b' to boot, 'e' to edit the selected command in the boot sequence, 'c' for a command-line, 'O' to open a new line after ('O' for before) the selected line, 'd' to remove the selected line, or escape to go back to the main menu.-

**v. Type b to boot the node into noncluster mode.**

---

**Note** – This change to the kernel boot parameter command does not persist over the system boot. The next time you reboot the node, it will boot into cluster mode. To boot into noncluster mode instead, perform these steps to again to add the -x option to the kernel boot parameter command.

---

If the instruction says to run the init 5 command, shut down the system then change the GRUB kernel boot command to /platform/i86pc/multiboot -sx instead.

**c. Upgrade each software application that requires an upgrade.**

Remember to boot into noncluster mode if you are directed to reboot, until all applications have been upgraded.

**d. Boot each node into cluster mode.**

- **On SPARC based systems, perform the following command:**

```
ok boot
```

- **On x86 based systems, perform the following commands:**

When the GRUB menu is displayed, select the appropriate Solaris entry and press Enter. The GRUB menu appears similar to the following:

```
GNU GRUB version 0.95 (631K lower / 2095488K upper memory)
+-----+
| Solaris 10 /sol_10_x86 |
| Solaris failsafe |
| |
+-----+
```

Use the ^ and v keys to select which entry is highlighted.

Press enter to boot the selected OS, 'e' to edit the commands before booting, or 'c' for a command-line.

### Example 8-1 Live Upgrade to Sun Cluster 3.2 Software

This example shows a live upgrade of a cluster node. The example upgrades the SPARC based node to the Solaris 10 OS, Sun Cluster 3.2 framework, and all Sun Cluster data services that support the live upgrade method. In this example, sc31u2 is the original boot environment (BE). The new BE that is upgraded is named sc32 and uses the mount point /sc32. The directory /net/installmachine/export/solaris10/OS\_image/ contains an image of the Solaris 10 OS. The Java ES installer state file is named sc32state.

The following commands typically produce copious output. This output is shown only where necessary for clarity.

```
phys-schost# lucreate sc31u2 -m /:/dev/dsk/c0t4d0s0:ufs -n sc32
```

```
...
```

```
lucreate: Creation of Boot Environment sc32 successful.
```

```
phys-schost# luupgrade -u -n sc32 -s /net/installmachine/export/solaris10/OS_image/
```

```
The Solaris upgrade of the boot environment sc32 is complete.
```

*Apply patches*

```
phys-schost# lumount sc32 /sc32
```

```
phys-schost# ls -l /sc32/usr/java
```

```
lrwxrwxrwx 1 root other 9 Apr 19 14:05 /sc32/usr/java -> /sc32/usr/j2se/
```

*Insert the Sun Java Availability Suite DVD-ROM.*

```
phys-schost# cd /cdrom/cdrom0/Solaris_sparc
```

```
phys-schost# ./installer -no -saveState sc32state
```

```
phys-schost# ./installer -nodisplay -noconsole -state sc32state -altroot /sc32
```

```
phys-schost# cd /cdrom/cdrom0/Solaris_sparc/sun_cluster/Sol_9/Tools
```

```
phys-schost# ./scinstall -u update -R /sc32
```

```
phys-schost# /sc32/usr/cluster/bin/scinstall -u update -s all -d /cdrom/cdrom0 -R /sc32
```

```
phys-schost# cd /
```

```
phys-schost# eject cdrom
```

```
phys-schost# luumount sc32
phys-schost# luactivate sc32
Activation of boot environment sc32 successful.
 Upgrade all other nodes
```

```
 Boot all nodes
phys-schost# shutdown -y -g0 -i6
ok boot
```

At this point, you might upgrade data-service applications that cannot use the live upgrade method, before you reboot into cluster mode.

**Troubleshooting** **DID device name errors** - During the creation of the inactive BE, if you receive an error that a file system that you specified with its DID device name, `/dev/dsk/did/dNsX`, does not exist, but the device name does exist, you must specify the device by its physical device name. Then change the `vfstab` entry on the alternate BE to use the DID device name instead. Perform the following steps:

- 1) For all unrecognized DID devices, specify the corresponding physical device names as arguments to the `-m` or `-M` option in the `lucreate` command. For example, if `/global/.devices/node@nodeid` is mounted on a DID device, use `lucreate -m /global/.devices/node@nodeid:/dev/dsk/cNtXdYsZ:ufs [-m...] -n BE-name` to create the BE.
- 2) Mount the inactive BE by using the `luumount -n BE-name -m BE-mount-point` command.
- 3) Edit the `/BE-name/etc/vfstab` file to convert the physical device name, `/dev/dsk/cNtXdYsZ`, to its DID device name, `/dev/dsk/did/dNsX`.

**Mount point errors** - During creation of the inactive boot environment, if you receive an error that the mount point that you supplied is not mounted, mount the mount point and rerun the `lucreate` command.

**New BE boot errors** - If you experience problems when you boot the newly upgraded environment, you can revert to your original BE. For specific information, see “Failure Recovery: Falling Back to the Original Boot Environment (Command-Line Interface)” in *Solaris 9 9/04 Installation Guide* or Chapter 10, “Failure Recovery: Falling Back to the Original Boot Environment (Tasks),” in *Solaris 10 Installation Guide: Solaris Live Upgrade and Upgrade Planning*.

**Next Steps** Go to “[How to Verify Upgrade of Sun Cluster 3.2 Software](#)” on page 249.

**See Also** You can choose to keep your original, and now inactive, boot environment for as long as you need to. When you are satisfied that your upgrade is acceptable, you can then choose to remove the old environment or to keep and maintain it.



- If you used an unmirrored volume for your inactive BE, delete the old BE files. For specific information, see “Deleting an Inactive Boot Environment” in *Solaris 9 9/04 Installation Guide* or “Deleting an Inactive Boot Environment” in *Solaris 10 Installation Guide: Solaris Live Upgrade and Upgrade Planning*.
- If you detached a plex to use as the inactive BE, reattach the plex and synchronize the mirrors. For more information about working with a plex, see “Example of Detaching and Upgrading One Side of a RAID 1 Volume (Mirror) (Command-Line Interface)” in *Solaris 9 9/04 Installation Guide* or “Example of Detaching and Upgrading One Side of a RAID-1 Volume (Mirror) (Command-Line Interface)” in *Solaris 10 Installation Guide: Solaris Live Upgrade and Upgrade Planning*.

You can also maintain the inactive BE. For information about how to maintain the environment, see Chapter 37, “Maintaining Solaris Live Upgrade Boot Environments (Tasks),” in *Solaris 9 9/04 Installation Guide* or Chapter 11, “Maintaining Solaris Live Upgrade Boot Environments (Tasks),” in *Solaris 10 Installation Guide: Solaris Live Upgrade and Upgrade Planning*.

## Completing the Upgrade

This section provides the following information to complete all Sun Cluster 3.2 software upgrade methods:

- [“How to Verify Upgrade of Sun Cluster 3.2 Software” on page 249](#)
- [“How to Finish Upgrade to Sun Cluster 3.2 Software” on page 251](#)

### ▼ How to Verify Upgrade of Sun Cluster 3.2 Software

Perform this procedure to verify that the cluster is successfully upgraded to Sun Cluster 3.2 software. On the Solaris 10 OS, perform all steps from the global zone only.

---

**Note** – This procedure provides the long forms of the Sun Cluster commands. Most commands also have short forms. Except for the forms of the command names, the commands are identical. For a list of the commands and their short forms, see Appendix A, “Sun Cluster Object-Oriented Commands,” in *Sun Cluster System Administration Guide for Solaris OS*.

---

**Before You Begin** Ensure that all upgrade procedures are completed for all cluster nodes that you are upgrading.

- 1 **On each node, become superuser.**
- 2 **On each upgraded node, view the installed levels of Sun Cluster software.**

```
phys-schost# clnode show-rev -v
```

The first line of output states which version of Sun Cluster software the node is running. This version should match the version that you just upgraded to.

**3 From any node, verify that all upgraded cluster nodes are running in cluster mode (Online).**

```
phys-schost# clnode status
```

See the `clnode(1CL)` man page for more information about displaying cluster status.

**4 SPARC: If you upgraded from Solaris 8 to Solaris 9 software, verify the consistency of the storage configuration.**

**a. On each node, run the following command to verify the consistency of the storage configuration.**

```
phys-schost# cldevice check
```



**Caution** – Do not proceed to [Step b](#) until your configuration passes this consistency check. Failure to pass this check might result in errors in device identification and cause data corruption.

The following table lists the possible output from the `cldevice check` command and the action you must take, if any.

| Example Message                                                                                                    | Action                                                                                                                  |
|--------------------------------------------------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------|
| device id for 'phys-schost-1:/dev/rdisk/clt3d0' does not match physical device's id, device may have been replaced | Go to <a href="#">“Recovering From an Incomplete Upgrade” on page 257</a> and perform the appropriate repair procedure. |
| device id for 'phys-schost-1:/dev/rdisk/c0t0d0' needs to be updated, run <code>cldevice repair</code> to update    | None. You update this device ID in <a href="#">Step b</a> .                                                             |
| No output message                                                                                                  | None.                                                                                                                   |

See the `cldevice(1CL)` man page for more information.

**b. On each node, migrate the Sun Cluster storage database to Solaris 9 device IDs.**

```
phys-schost# cldevice repair
```

**c. On each node, run the following command to verify that storage database migration to Solaris 9 device IDs is successful.**

```
phys-schost# cldevice check
```

- If the `cldevice` command displays a message, return to [Step a](#) to make further corrections to the storage configuration or the storage database.
- If the `cldevice` command displays no messages, the device-ID migration is successful. When device-ID migration is verified on all cluster nodes, proceed to [“How to Finish Upgrade to Sun Cluster 3.2 Software” on page 251](#).

**Example 8-2** Verifying Upgrade to Sun Cluster 3.2 Software

The following example shows the commands used to verify upgrade of a two-node cluster to Sun Cluster 3.2 software. The cluster node names are `phys-schost-1` and `phys-schost-2`.

```
phys-schost# clnode show-rev -v
3.2
...
phys-schost# clnode status
=== Cluster Nodes ===

--- Node Status ---

Node Name Status

phys-schost-1 Online
phys-schost-2 Online
```

**Next Steps** Go to [“How to Finish Upgrade to Sun Cluster 3.2 Software”](#) on page 251.

## ▼ How to Finish Upgrade to Sun Cluster 3.2 Software

Perform this procedure to finish Sun Cluster upgrade. On the Solaris 10 OS, perform all steps from the global zone only. First, reregister all resource types that received a new version from the upgrade. Second, modify eligible resources to use the new version of the resource type that the resource uses. Third, re-enable resources. Finally, bring resource groups back online.

**Before You Begin** Ensure that all steps in [“How to Verify Upgrade of Sun Cluster 3.2 Software”](#) on page 249 are completed.

### 1 Copy the security files for the common agent container to all cluster nodes.

This step ensures that security files for the common agent container are identical on all cluster nodes and that the copied files retain the correct file permissions.

#### a. On each node, stop the Sun Java Web Console agent.

```
phys-schost# /usr/sbin/smcwebserver stop
```

#### b. On each node, stop the security file agent.

```
phys-schost# /usr/sbin/cacaoadm stop
```

#### c. On one node, change to the `/etc/cacao/instances/default/` directory.

```
phys-schost-1# cd /etc/cacao/instances/default/
```

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

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

- e. Copy the `/tmp/SECURITY.tar` file to each of the other cluster nodes.**

- f. On each node to which you copied the `/tmp/SECURITY.tar` file, extract the security files.**

Any security files that already exist in the `/etc/cacao/instances/default/` directory are overwritten.

```
phys-schost-2# cd /etc/cacao/instances/default/
phys-schost-2# tar xf /tmp/SECURITY.tar
```

- g. Delete the `/tmp/SECURITY.tar` file from each node in the cluster.**

You must delete each copy of the tar file to avoid security risks.

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

- h. On each node, start the security file agent.**

```
phys-schost# /usr/sbin/cacaoadm start
```

- i. On each node, start the Sun Java Web Console agent.**

```
phys-schost# /usr/sbin/smcwebserver start
```

- 2 If you upgraded any data services that are not supplied on the product media, register the new resource types for those data services.**

Follow the documentation that accompanies the data services.

- 3 If you upgraded Sun Cluster HA for SAP liveCache from the Sun Cluster 3.0 or 3.1 version to the Sun Cluster 3.2 version, modify the `/opt/SUNWsclc/livecache/bin/lccluster` configuration file.**

- a. Become superuser on a node that will host the liveCache resource.**

- b. Copy the new `/opt/SUNWsclc/livecache/bin/lccluster` file to the `/sapdb/LC_NAME/db/sap/` directory.**

Overwrite the `lccluster` file that already exists from the previous configuration of the data service.

- c. Configure this `/sapdb/LC_NAME/db/sap/lccluster` file as documented in “How to Register and Configure Sun Cluster HA for SAP liveCache” in *Sun Cluster Data Service for SAP liveCache Guide for Solaris OS*.**

**4 If you upgraded the Solaris OS and your configuration uses dual-string mediators for Solaris Volume Manager software, restore the mediator configurations.**

- a. Determine which node has ownership of a disk set to which you will add the mediator hosts.**

```
phys-schost# metaset -s setname
```

*-s setname* Specifies the disk set name.

- b. On the node that masters or will master the disk set, become superuser.**

- c. If no node has ownership, take ownership of the disk set.**

```
phys-schost# cldevicegroup switch -n node devicegroup
```

*node* Specifies the name of the node to become primary of the disk set.

*devicegroup* Specifies the name of the disk set.

- d. Re-create the mediators.**

```
phys-schost# metaset -s setname -a -m mediator-host-list
```

*-a* Adds to the disk set.

*-m mediator-host-list* Specifies the names of the nodes to add as mediator hosts for the disk set.

- e. Repeat these steps for each disk set in the cluster that uses mediators.**

**5 If you upgraded VxVM, upgrade all disk groups.**

- a. Bring online and take ownership of a disk group to upgrade.**

```
phys-schost# cldevicegroup switch -n node devicegroup
```

- b. Run the following command to upgrade a disk group to the highest version supported by the VxVM release you installed.**

```
phys-schost# vxdg upgrade dgname
```

See your VxVM administration documentation for more information about upgrading disk groups.

- c. Repeat for each remaining VxVM disk group in the cluster.**

**6 Migrate resources to new resource type versions.**

You must migrate all resources to the Sun Cluster 3.2 resource-type version.

---

**Note** – For Sun Cluster HA for SAP Web Application Server, if you are using a J2EE engine resource or a web application server component resource or both, you must delete the resource and recreate it with the new web application server component resource. Changes in the new web application server component resource includes integration of the J2EE functionality. For more information, see *Sun Cluster Data Service for SAP Web Application Server Guide for Solaris OS*.

---

See “Upgrading a Resource Type” in *Sun Cluster Data Services Planning and Administration Guide for Solaris OS*, which contains procedures which use the command line. Alternatively, you can perform the same tasks by using the Resource Group menu of the `clsetup` utility. The process involves performing the following tasks:

- Registering the new resource type.
- Migrating the eligible resource to the new version of its resource type.
- Modifying the extension properties of the resource type as specified in *Sun Cluster 3.2 Release Notes for Solaris OS*.

---

**Note** – The Sun Cluster 3.2 release introduces new default values for some extension properties, such as the `Retry_interval` property. These changes affect the behavior of any existing resource that uses the default values of such properties. If you require the previous default value for a resource, modify the migrated resource to set the property to the previous default value.

---

- 7 If your cluster runs the Sun Cluster HA for Sun Java System Application Server EE (HADB) data service and you shut down the HADB database before you began a dual-partition upgrade, re-enable the resource and start the database.**

```
phys-schost# clresource enable hadb-resource
phys-schost# hadbm start database-name
```

For more information, see the `hadbm(1m)` man page.

- 8 If you upgraded to the Solaris 10 OS and the Apache `httpd.conf` file is located on a cluster file system, ensure that the HTTPD entry in the Apache control script still points to that location.**

- a. View the HTTPD entry in the `/usr/apache/bin/apchectl` file.**

The following example shows the `httpd.conf` file located on the `/global` cluster file system.

```
phys-schost# cat /usr/apache/bin/apchectl | grep HTTPD=/usr
HTTPD="/usr/apache/bin/httpd -f /global/web/conf/httpd.conf"
```

- b. If the file does not show the correct HTTPD entry, update the file.**

```
phys-schost# vi /usr/apache/bin/apchectl
#HTTPD=/usr/apache/bin/httpd
HTTPD="/usr/apache/bin/httpd -f /global/web/conf/httpd.conf"
```

**9 From any node, start the `clsetup` utility.**

```
phys-schost# clsetup
```

The `clsetup` Main Menu is displayed.

**10 Re-enable all disabled resources.**

a. Type the number that corresponds to the option for Resource groups and press the Return key.  
The Resource Group Menu is displayed.

b. Type the number that corresponds to the option for Enable/Disable a resource and press the Return key.

c. Choose a resource to enable and follow the prompts.

d. Repeat [Step c](#) for each disabled resource.

e. When all resources are re-enabled, type `q` to return to the Resource Group Menu.

**11 Bring each resource group back online.**

This step includes the bringing online of resource groups in non-global zones.

a. Type the number that corresponds to the option for Online/Offline or Switchover a resource group and press the Return key.

b. Follow the prompts to put each resource group into the managed state and then bring the resource group online.

**12 When all resource groups are back online, exit the `clsetup` utility.**

Type `q` to back out of each submenu, or press `Ctrl-C`.

**13 If, before upgrade, you enabled automatic node reboot if all monitored disk paths fail, ensure that the feature is still enabled.**

Also perform this task if you want to configure automatic reboot for the first time.

a. Determine whether the automatic reboot feature is enabled or disabled.

```
phys-schost# clnode show
```

▪ If the `reboot_on_path_failure` property is set to enabled, no further action is necessary.

▪ If `reboot_on_path_failure` property is set to disabled, proceed to the next step to re-enable the property.

**b. Enable the automatic reboot feature.**

```
phys-schost# clnode set -p reboot_on_path_failure=enabled
```

-p Specifies the property to set

reboot\_on\_path\_failure=enable Specifies that the node will reboot if all monitored disk paths fail, provided that at least one of the disks is accessible from a different node in the cluster.

**c. Verify that automatic reboot on disk-path failure is enabled.**

```
phys-schost# clnode show
=== Cluster Nodes ===

Node Name: node
...
reboot_on_path_failure: enabled
...
```

**14 (Optional) Capture the disk partitioning information for future reference.**

```
phys-schost# prtvtoc /dev/rdisk/cNtXdYsZ > filename
```

Store the file in a location outside the cluster. If you make any disk configuration changes, run this command again to capture the changed configuration. If a disk fails and needs replacement, you can use this information to restore the disk partition configuration. For more information, see the prtvtoc(1M) man page.

**15 (Optional) Make a backup of your cluster configuration.**

An archived backup of your cluster configuration facilitates easier recovery of the your cluster configuration,

For more information, see “How to Back Up the Cluster Configuration” in *Sun Cluster System Administration Guide for Solaris OS*.

**Troubleshooting** **Resource-type migration failure** - Normally, you migrate resources to a new resource type while the resource is offline. However, some resources need to be online for a resource-type migration to succeed. If resource-type migration fails for this reason, error messages similar to the following are displayed:

```
phys-schost - Resource depends on a SUNW.HASStoragePlus type resource that is not
online anywhere.
(C189917) VALIDATE on resource nfsrs, resource group rg, exited with
non-zero exit status.
(C720144) Validation of resource nfsrs in resource group rg on node
phys-schost failed.
```



If resource-type migration fails because the resource is offline, use the `clsetup` utility to re-enable the resource and then bring its related resource group online. Then repeat migration procedures for the resource.

**Java binaries location change** - If the location of the Java binaries changed during the upgrade of shared components, you might see error messages similar to the following when you attempt to run the `cacaoadm start` or `smcwebserver start` commands:

```
/opt/SUNWcacao/bin/cacaoadm start
No suitable Java runtime found. Java 1.4.2_03 or higher is required.
Jan 3 17:10:26 ppups3 cacao: No suitable Java runtime found. Java 1.4.2_03 or higher
is required.
Cannot locate all the dependencies
```

```
smcwebserver start
/usr/sbin/smcwebserver: /usr/jdk/jdk1.5.0_04/bin/java: not found
```

These errors are generated because the start commands cannot locate the current location of the Java binaries. The `JAVA_HOME` property still points to the directory where the previous version of Java was located, but that previous version was removed during upgrade.

To correct this problem, change the setting of `JAVA_HOME` in the following configuration files to use the current Java directory:

```
/etc/webconsole/console/config.properties/etc/opt/SUNWcacao/cacao.properties
```

**Next Steps** If you have a SPARC based system and use Sun Management Center to monitor the cluster, go to [“SPARC: How to Upgrade Sun Cluster Module Software for Sun Management Center” on page 189](#).

To install or complete upgrade of Sun Cluster Geographic Edition 3.2 software, see *Sun Cluster Geographic Edition Installation Guide*.

Otherwise, the cluster upgrade is complete.

## Recovering From an Incomplete Upgrade

This section provides the following information to recover from certain kinds of incomplete upgrades:

- [“How to Recover from a Failed Dual-Partition Upgrade” on page 258](#)
- [“Recovering From Storage Configuration Changes During Upgrade” on page 260](#)

## ▼ How to Recover from a Failed Dual-Partition Upgrade

If you experience an unrecoverable error during upgrade, perform this procedure to back out of the upgrade.

---

**Note** – You cannot restart a dual-partition upgrade after the upgrade has experienced an unrecoverable error.

---

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

**2 Boot each node into noncluster mode.**

- **On SPARC based systems, perform the following command:**

```
ok boot -x
```

- **On x86 based systems, perform the following commands:**

- a. **In the GRUB menu, use the arrow keys to select the appropriate Solaris entry and type e to edit its commands.**

The GRUB menu appears similar to the following:

```
GNU GRUB version 0.95 (631K lower / 2095488K upper memory)
```

```
+-----+
| Solaris 10 /sol_10_x86 |
| Solaris failsafe |
| |
+-----+
```

Use the ^ and v keys to select which entry is highlighted.  
Press enter to boot the selected OS, 'e' to edit the  
commands before booting, or 'c' for a command-line.

For more information about GRUB based booting, see Chapter 11, “GRUB Based Booting (Tasks),” in *System Administration Guide: Basic Administration*.

- b. **In the boot parameters screen, use the arrow keys to select the kernel entry and type e to edit the entry.**

The GRUB boot parameters screen appears similar to the following:

```
GNU GRUB version 0.95 (615K lower / 2095552K upper memory)
```

```
+-----+
| root (hd0,0,a) |
| kernel /platform/i86pc/multiboot |
| module /platform/i86pc/boot_archive |
+-----+
```

Use the ^ and v keys to select which entry is highlighted.  
Press 'b' to boot, 'e' to edit the selected command in the

boot sequence, 'c' for a command-line, 'o' to open a new line after ('O' for before) the selected line, 'd' to remove the selected line, or escape to go back to the main menu.

**c. Add -x to the command to specify that the system boot into noncluster mode.**

[ Minimal BASH-like line editing is supported. For the first word, TAB lists possible command completions. Anywhere else TAB lists the possible completions of a device/filename. ESC at any time exits. ]

```
grub edit> kernel /platform/i86pc/multiboot -x
```

**d. Press Enter to accept the change and return to the boot parameters screen.**

The screen displays the edited command.

```
GNU GRUB version 0.95 (615K lower / 2095552K upper memory)
```

```
+-----+
| root (hd0,0,a) |
| kernel /platform/i86pc/multiboot -x |
| module /platform/i86pc/boot_archive |
+-----+
```

Use the ^ and v keys to select which entry is highlighted. Press 'b' to boot, 'e' to edit the selected command in the boot sequence, 'c' for a command-line, 'o' to open a new line after ('O' for before) the selected line, 'd' to remove the selected line, or escape to go back to the main menu.-

**e. Type b to boot the node into noncluster mode.**

---

**Note** – This change to the kernel boot parameter command does not persist over the system boot. The next time you reboot the node, it will boot into cluster mode. To boot into noncluster mode instead, perform these steps to again to add the -x option to the kernel boot parameter command.

---

**3 On each node, run the upgrade recovery script from the installation media.**

If the node successfully upgraded to Sun Cluster 3.2 software, you can alternatively run the `scinstall` command from the `/usr/cluster/bin` directory.

```
phys-schost# cd /cdrom/cdrom0/Solaris_arch/Product/sun_cluster/Solaris_ver/Tools
phys-schost# ./scinstall -u recover
```

-u Specifies upgrade.

recover Restores the `/etc/vfstab` file and the Cluster Configuration Repository (CCR) database to their original state before the start of the dual-partition upgrade.

The recovery process leaves the cluster nodes in noncluster mode. Do **not** attempt to reboot the nodes into cluster mode.

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

#### 4 Perform either of the following tasks.

- **Restore the old software from backup to return the cluster to its original state.**
- **Continue to upgrade software on the cluster by using the standard upgrade method.**

This method requires that all cluster nodes remain in noncluster mode during the upgrade. See the task map for standard upgrade, [Table 8–1](#). You can resume the upgrade at the last task or step in the standard upgrade procedures that you successfully completed before the dual-partition upgrade failed.

## Recovering From Storage Configuration Changes During Upgrade

This section provides the following repair procedures to follow if changes were inadvertently made to the storage configuration during upgrade:

- [“How to Handle Storage Reconfiguration During an Upgrade” on page 260](#)
- [“How to Resolve Mistaken Storage Changes During an Upgrade” on page 261](#)

### ▼ How to Handle Storage Reconfiguration During an Upgrade

Any changes to the storage topology, including running Sun Cluster commands, should be completed before you upgrade the cluster to Solaris 9 or Solaris 10 software. If, however, changes were made to the storage topology during the upgrade, perform the following procedure. This procedure ensures that the new storage configuration is correct and that existing storage that was not reconfigured is not mistakenly altered.

---

**Note** – This procedure provides the long forms of the Sun Cluster commands. Most commands also have short forms. Except for the forms of the command names, the commands are identical. For a list of the commands and their short forms, see Appendix A, “Sun Cluster Object-Oriented Commands,” in *Sun Cluster System Administration Guide for Solaris OS*.

---

**Before You Begin** Ensure that the storage topology is correct. Check whether the devices that were flagged as possibly being replaced map to devices that actually were replaced. If the devices were not replaced, check for and correct possible accidental configuration changes, such as incorrect cabling.

**1 On a node that is attached to the unverified device, become superuser.**

**2 Manually update the unverified device.**

```
phys-schost# cldevice repair device
```

See the `cldevice(1CL)` man page for more information.

### 3 Update the DID driver.

```
phys-schost# scdidadm -ui
phys-schost# scdidadm -r
```

- u Loads the device-ID configuration table into the kernel.
- i Initializes the DID driver.
- r Reconfigures the database.

### 4 Repeat [Step 2](#) through [Step 3](#) on all other nodes that are attached to the unverified device.

**Next Steps** Return to the remaining upgrade tasks. Go to [Step 4](#) in “[How to Upgrade Sun Cluster 3.2 Software \(Standard\)](#)” on page 207.

## ▼ How to Resolve Mistaken Storage Changes During an Upgrade

If accidental changes are made to the storage cabling during the upgrade, perform the following procedure to return the storage configuration to the correct state.

---

**Note** – This procedure assumes that no physical storage was actually changed. If physical or logical storage devices were changed or replaced, instead follow the procedures in “[How to Handle Storage Reconfiguration During an Upgrade](#)” on page 260.

---

**Before You Begin** Return the storage topology to its original configuration. Check the configuration of the devices that were flagged as possibly being replaced, including the cabling.

### 1 On each node of the cluster, become superuser.

### 2 Update the DID driver on each node of the cluster.

```
phys-schost# scdidadm -ui
phys-schost# scdidadm -r
```

- u Loads the device-ID configuration table into the kernel.
- i Initializes the DID driver.
- r Reconfigures the database.

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

### 3 If the `scdidadm` command returned any error messages in [Step 2](#), make further modifications as needed to correct the storage configuration, then repeat [Step 2](#).

**Next Steps** Return to the remaining upgrade tasks. Go to [Step 4](#) in “[How to Upgrade Sun Cluster 3.2 Software \(Standard\)](#)” on page 207.

# Uninstalling Software From the Cluster

---

This chapter provides procedures for uninstalling or removing certain software from a Sun Cluster configuration. The following procedures are in this chapter:

- “How to Unconfigure Sun Cluster Software to Correct Installation Problems” on page 263
- “How to Remove Sun Cluster Information From a JumpStart Install Server” on page 267
- “How to Uninstall SunPlex Manager Software” on page 268
- “How to Uninstall the SUNWscrdt Package” on page 269
- “How to Unload the RSMRDT Driver Manually” on page 270

## Uninstalling the Software

This section provides procedures to uninstall or remove certain software products from a cluster.

### ▼ How to Unconfigure Sun Cluster Software to Correct Installation Problems

Perform this procedure if the installed node cannot join the cluster or if you need to correct configuration information. For example, perform this procedure on all nodes to reconfigure the transport adapters or the private-network address.

---

**Note** – If the node has already joined the cluster and is no longer in installation mode, as described in Step 2 of “How to Verify the Quorum Configuration and Installation Mode” on page 121, do not perform this procedure. Instead, go to “How to Uninstall Sun Cluster Software From a Cluster Node” in *Sun Cluster System Administration Guide for Solaris OS*.

---

**Before You Begin** Attempt to rerun cluster configuration of the node. You can correct certain cluster node configuration failures by repeating Sun Cluster software configuration on the node.

**1 Add to the cluster's node-authentication list each node that you intend to unconfigure.**

If you are unconfiguring a single-node cluster, skip to [Step 2](#).

**a. On an active cluster member other than the node that you are unconfiguring, become superuser.****b. Specify the name of the node to add to the authentication list.**

```
phys-schost# /usr/cluster/bin/claccess allow -h nodename
```

-h *nodename* Specifies the name of the node to add to the authentication list.

You can also use the `clsetup` utility to perform this task. See “How to Add a Node to the Authorized Node List” in *Sun Cluster System Administration Guide for Solaris OS* for procedures.

**2 On a node that you intend to unconfigure, become superuser.****3 Shut down the node.**

```
phys-schost# shutdown -g0 -y -i0
```

**4 Reboot the node into noncluster mode.****■ On SPARC based systems, perform the following command:**

```
ok boot -x
```

**■ On x86 based systems, perform the following commands:****a. In the GRUB menu, use the arrow keys to select the appropriate Solaris entry and type `e` to edit its commands.**

The GRUB menu appears similar to the following:

```
GNU GRUB version 0.95 (631K lower / 2095488K upper memory)
```

```
+-----+
| Solaris 10 /sol_10_x86 |
| Solaris failsafe |
| |
+-----+
```

Use the `^` and `v` keys to select which entry is highlighted.

Press `enter` to boot the selected OS, `'e'` to edit the commands before booting, or `'c'` for a command-line.

For more information about GRUB based booting, see Chapter 11, “GRUB Based Booting (Tasks),” in *System Administration Guide: Basic Administration*.



- b. In the boot parameters screen, use the arrow keys to select the kernel entry and type e to edit the entry.**

The GRUB boot parameters screen appears similar to the following:

```
GNU GRUB version 0.95 (615K lower / 2095552K upper memory)
```

```
+-----+
| root (hd0,0,a) |
| kernel /platform/i86pc/multiboot |
| module /platform/i86pc/boot_archive |
+-----+
```

Use the ^ and v keys to select which entry is highlighted. Press 'b' to boot, 'e' to edit the selected command in the boot sequence, 'c' for a command-line, 'o' to open a new line after ('O' for before) the selected line, 'd' to remove the selected line, or escape to go back to the main menu.

- c. Add -x to the command to specify that the system boot into noncluster mode.**

[ Minimal BASH-like line editing is supported. For the first word, TAB lists possible command completions. Anywhere else TAB lists the possible completions of a device/filename. ESC at any time exits. ]

```
grub edit> kernel /platform/i86pc/multiboot -x
```

- d. Press Enter to accept the change and return to the boot parameters screen.**

The screen displays the edited command.

```
GNU GRUB version 0.95 (615K lower / 2095552K upper memory)
```

```
+-----+
| root (hd0,0,a) |
| kernel /platform/i86pc/multiboot -x |
| module /platform/i86pc/boot_archive |
+-----+
```

Use the ^ and v keys to select which entry is highlighted. Press 'b' to boot, 'e' to edit the selected command in the boot sequence, 'c' for a command-line, 'o' to open a new line after ('O' for before) the selected line, 'd' to remove the selected line, or escape to go back to the main menu.-

- e. Type b to boot the node into noncluster mode.**

---

**Note** – This change to the kernel boot parameter command does not persist over the system boot. The next time you reboot the node, it will boot into cluster mode. To boot into noncluster mode instead, perform these steps to again to add the -x option to the kernel boot parameter command.

---

- 5 Change to a directory, such as the root (/) directory, that does not contain any files that are delivered by the Sun Cluster packages.**

```
phys-schost# cd /
```

- 6 Remove the node from the cluster configuration.**

```
phys-schost# /usr/cluster/bin/clnode remove
```

The node is removed from the cluster configuration but Sun Cluster software is not removed from the node.

See the `clnode(1CL)` man page for more information.

- 7 Repeat Step 2 through Step 6 for each additional node to unconfigure.**

- 8 (Optional) Uninstall Sun Cluster framework and data service software packages.**

---

**Note** – If you do not need to remove or reinstall Sun Cluster software, you can skip this step.

---

This step also removes the Sun Cluster entry from the Sun Java Enterprise System (Java ES) product registry. If the Java ES product registry contains a record that Sun Cluster software is installed, the Java ES installer will show the Sun Cluster component grayed out and will not permit reinstallation.

- a. Start the `uninstall` program.**

Run the following command, where *ver* is the version of the Java ES distribution from which you installed Sun Cluster software.

```
phys-schost# /var/sadm/prod/SUNWentsysver/uninstall
```

- b. Follow the onscreen instructions to select the Sun Cluster components to uninstall.**

---

**Note** – If Sun Cluster Geographic Edition software is installed, you must uninstall it as well.

---

For more information about using the `uninstall` program, see Chapter 8, “Uninstalling Product Components,” in *Sun Java Enterprise System 5 Installation Guide for UNIX*.

**Next Steps** To reinstall or reconfigure Sun Cluster software on the node, refer to [Table 2–1](#). This table lists all installation tasks and the order in which to perform the tasks.

To physically remove the node from the cluster, see “How to Remove an Interconnect Component” in *Sun Cluster 3.1 - 3.2 Hardware Administration Manual for Solaris OS* and the removal procedure in the Sun Cluster Hardware Administration Collection manual for your storage array.

## ▼ How to Remove Sun Cluster Information From a JumpStart Install Server

Perform this procedure to remove Sun Cluster information from a JumpStart install server that was used to install and configure a cluster. You can remove information for one or more individual nodes or for one or more entire clusters. For more information about the JumpStart feature, see *Solaris 10 Installation Guide: Custom JumpStart and Advanced Installations*.

**1 Become superuser on the JumpStart install server.**

**2 Change to the JumpStart directory that you used to install Sun Cluster software.**

```
installserver# cd jumpstart-dir
```

**3 Remove from the `rules` file any entries that were created by the `scinstall` command that contain the name of a node that you want to remove.**

The Sun Cluster entries reference `autostinstall.class` or `autoscinstall.finish` or both. Entries appear similar to the following, where *release* is the version of Sun Cluster software:

```
hostname phys-schost-1 - autoscinstall.d/release/autoscinstall.class \
autoscinstall.d/release/autoscinstall.finish
```

**4 Regenerate the `rules.ok` file.**

Run the `check` command that is located in the `jumpstart-dir/` directory to regenerate the `rules.ok` file.

```
installserver# ./check
```

**5 Remove from the appropriate `clusters/clustername/` directory the symbolic link for each node that you want to remove.**

- To remove the symbolic link for one or more nodes in a cluster, remove the link that is named for each node to remove.

```
installserver# rm -f autoscinstall.d/clusters/clustername/nodename
```

- To remove the symbolic links for an entire cluster, recursively remove the directory that is named for the cluster to remove.

```
installserver# rm -rf autoscinstall.d/clusters/clustername
```

- To remove the symbolic links for all clusters, recursively remove the `clusters/` directory.

```
installserver# rm -rf autoscinstall.d/clusters
```

- 6 Remove from the `autosinstall.d/` directory the node configuration directory that is named for each node that you want to remove.**

If you are removing information for an entire cluster, remove the directory for each node in the cluster.

- To remove information for one or more nodes in a cluster, recursively remove the directory for each node.

```
installserver# rm -rf autosinstall.d/nodes/nodename
```

- To remove all entries for all clusters, recursively remove the `autosinstall.d` directory.

```
installserver# rm -rf autosinstall.d
```

- 7 Remove the `.autosinstall.log.3` file.**

```
installserver# rm .autosinstall.log.3
```

- 8 (Optional) If you used a flash archive to JumpStart install the cluster, remove the flash archive if you no longer need the file.**

```
installserver# rm filename.flar
```

**Next Steps** If you intend to use custom JumpStart to reinstall a cluster from which you removed information for one or more nodes that were removed from that cluster, you must rerun interactive `scinstall` to update the cluster node list. See [“How to Install Solaris and Sun Cluster Software \(JumpStart\)”](#) on page 79.

## ▼ How to Uninstall SunPlex Manager Software

Use this procedure to uninstall SunPlex Manager software that was installed by the Java ES `installer` utility up to and including the Sun Java Enterprise System 2005Q4 distribution or by any other installation method.

To remove Sun Cluster Manager software that was installed with the Sun Java Enterprise System 2006Q4 or compatible distribution of the Java ES `installer` utility, instead use the Java ES `uninstall` utility to remove these packages. For more information, see Chapter 8, “Uninstalling Product Components,” in *Sun Java Enterprise System 5 Installation Guide for UNIX*.

---

**Note** – Do not uninstall SunPlex Manager or Sun Cluster Manager or its shared components if you intend to use the graphical user interface (GUI) to Sun Cluster Geographic Edition software or to Sun Cluster features that are available through the GUI. These features include the data-service configuration wizards or system resource monitoring.

However, if you use the command-line interface to administer these features, it is safe to uninstall SunPlex Manager or Sun Cluster Manager software.

---

Perform this procedure on each node in the cluster to uninstall SunPlex Manager software and related Sun Java Enterprise System (Java ES) shared components.

---

**Note** – SunPlex Manager software must be installed on all cluster nodes or on none.

---

- 1 Become superuser on a cluster node.**
- 2 Remove the SunPlex Manager software packages.**  

```
phys-schost# pkgrm SUNWscspm SUNWscspmu SUNWscspmr
```
- 3 (Optional) Remove Sun Java Web Console software packages, if you have no other need for them.**  

```
phys-schost# pkgrm SUNWmctag SUNWmconr SUNWmcon SUNWmcos SUNWmcosx
```
- 4 (Optional) If you removed Sun Java Web Console packages, remove Apache Tomcat and Java Studio Enterprise Web Application Framework (JATO) software packages, if you have no other need for them.**

Remove the packages listed below for each additional product that you want to uninstall, in the order that the set of packages is listed.

```
phys-schost# pkgrm packages
```

---

| Product       | Package Names                    |
|---------------|----------------------------------|
| Apache Tomcat | SUNWtcatu                        |
| JATO          | SUNWjato SUNWjatodmo SUNWjatodoc |

---

## ▼ How to Uninstall the SUNWscrdt Package

Perform this procedure on each node in the cluster.

**Before You Begin** Verify that no applications are using the RSMRDT driver before performing this procedure.

- 1 Become superuser on the node to which you want to uninstall the SUNWscrdt package.**

**2 Uninstall the SUNWscrdt package.**

```
phys-schost# pkgrm SUNWscrdt
```

**▼ How to Unload the RSMRDT Driver Manually**

If the driver remains loaded in memory after completing “[How to Uninstall the SUNWscrdt Package](#)” on page 269, perform this procedure to unload the driver manually.

**1 Start the adb utility.**

```
phys-schost# adb -kw
```

**2 Set the kernel variable `clifrsrmdt_modunload_ok` to 1.**

```
physmem NNNN
clifrsrmdt_modunload_ok/W 1
```

**3 Exit the adb utility by pressing Control-D.****4 Find the `clif_rsmrdt` and `rsmrdt` module IDs.**

```
phys-schost# modinfo | grep rdt
```

**5 Unload the `clif_rsmrdt` module.**

You must unload the `clif_rsmrdt` module before you unload the `rsmrdt` module.

```
phys-schost# modunload -i clif_rsmrdt_id
```

*clif\_rsmrdt\_id* Specifies the numeric ID for the module that is being unloaded.

**6 Unload the `rsmrdt` module.**

```
phys-schost# modunload -i rsmrdt_id
```

*rsmrdt\_id* Specifies the numeric ID for the module that is being unloaded.

**7 Verify that the module was successfully unloaded.**

```
phys-schost# modinfo | grep rdt
```

**Example 9-1 Unloading the RSMRDT Driver**

The following example shows the console output after the RSMRDT driver is manually unloaded.

```
phys-schost# adb -kw
physmem fc54
clifrsrmdt_modunload_ok/W 1
clifrsrmdt_modunload_ok: 0x0 = 0x1
^D
phys-schost# modinfo | grep rsm
```

```
88 f064a5cb 974 - 1 rsmops (RSMOPS module 1.1)
93 f08e07d4 b95 - 1 clif_rsmrdt (CLUSTER-RSMRDT Interface module)
94 f0d3d000 13db0 194 1 rsmrdt (Reliable Datagram Transport dri)
phys-schost# modunload -i 93
phys-schost# modunload -i 94
phys-schost# modinfo | grep rsm
88 f064a5cb 974 - 1 rsmops (RSMOPS module 1.1)
```

**Troubleshooting** If the `modunload` command fails, applications are probably still using the driver. Terminate the applications before you run `modunload` again.





# Sun Cluster Installation and Configuration Worksheets

---

This appendix provides worksheets to plan various components of your cluster configuration and examples of completed worksheets for your reference. See “Installation and Configuration Worksheets” in *Sun Cluster Data Services Planning and Administration Guide for Solaris OS* for configuration worksheets for resources, resource types, and resource groups.

# Installation and Configuration Worksheets

If necessary, make additional copies of a worksheet to accommodate all the components in your cluster configuration. Follow planning guidelines in [Chapter 1](#) to complete these worksheets. Then refer to your completed worksheets during cluster installation and configuration.

---

**Note** – The data used in the worksheet examples is intended as a guide only. The examples do not represent a complete configuration of a functional cluster.

---

The following table lists the planning worksheets and examples provided in this appendix, as well as the titles of sections in [Chapter 1](#) that contain related planning guidelines.

**TABLE A-1** Cluster Installation Worksheets and Related Planning Guidelines

| Worksheet                                                | Example                                                                                    | Section Titles of Related Planning Guidelines                                                                           |
|----------------------------------------------------------|--------------------------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------|
| “Local File System Layout Worksheet” on page 275         | “Example: Local File System Layout Worksheets, With and Without Mirrored Root” on page 276 | “System Disk Partitions” on page 16<br>“Guidelines for Mirroring the Root Disk” on page 42                              |
| “Local Devices Worksheets” on page 277                   | “Example: Local Devices Worksheets” on page 278                                            | ---                                                                                                                     |
| “Device Group Configurations Worksheet” on page 279      | “Example: Device Group Configurations Worksheet” on page 280                               | “Device Groups” on page 31<br>“Planning Volume Management” on page 36                                                   |
| “Volume-Manager Configurations Worksheet” on page 281    | “Example: Volume-Manager Configurations Worksheet” on page 282                             | “Planning Volume Management” on page 36<br>Your volume manager documentation                                            |
| “Volumes Worksheet (Solaris Volume Manager)” on page 283 | “Example: Volumes Worksheet (Solaris Volume Manager)” on page 284                          | “Planning Volume Management” on page 36<br><i>Solaris Volume Manager Administration Guide</i> (Solaris 9 or Solaris 10) |

# Local File System Layout Worksheet

Node name: \_\_\_\_\_

**TABLE A-2** Local File Systems With Mirrored Root Worksheet

| Volume Name | Component | Component | File System    | Size |
|-------------|-----------|-----------|----------------|------|
|             |           |           | /              |      |
|             |           |           | swap           |      |
|             |           |           | /globaldevices |      |
|             |           |           |                |      |
|             |           |           |                |      |
|             |           |           |                |      |
|             |           |           |                |      |

**TABLE A-3** Local File Systems With Nonmirrored Root Worksheet

| Device Name | File System    | Size |
|-------------|----------------|------|
|             | /              |      |
|             | swap           |      |
|             | /globaldevices |      |
|             |                |      |
|             |                |      |
|             |                |      |
|             |                |      |

## Example: Local File System Layout Worksheets, With and Without Mirrored Root

Node name: **phys-schost-1**

TABLE A-4 Example: Local File Systems With Mirrored Root Worksheet

| Volume Name | Component       | Component       | File System        | Size           |
|-------------|-----------------|-----------------|--------------------|----------------|
| <b>d1</b>   | <b>c0t0d0s0</b> | <b>c1t0d0s0</b> | /                  | <b>6.75 GB</b> |
| <b>d2</b>   | <b>c0t0d0s1</b> | <b>c1t0d0s1</b> | swap               | <b>750 MB</b>  |
| <b>d3</b>   | <b>c0t0d0s3</b> | <b>c1t0d0s3</b> | /globaldevices     | <b>512 MB</b>  |
|             |                 |                 |                    |                |
|             |                 |                 |                    |                |
|             |                 |                 |                    |                |
| <b>d7</b>   | <b>c0t0d0s7</b> | <b>c1t0d0s7</b> | <b>SVM replica</b> | <b>20 MB</b>   |

TABLE A-5 Example: Local File Systems With Nonmirrored Root Worksheet

| Device Name     | File System        | Size           |
|-----------------|--------------------|----------------|
| <b>c0t0d0s0</b> | /                  | <b>6.75 GB</b> |
| <b>c0t0d0s1</b> | swap               | <b>750 MB</b>  |
| <b>c0t0d0s3</b> | /globaldevices     | <b>512 MB</b>  |
|                 |                    |                |
|                 |                    |                |
|                 |                    |                |
| <b>c0t0d0s7</b> | <b>SVM replica</b> | <b>20 MB</b>   |

# Local Devices Worksheets

Node name: \_\_\_\_\_

TABLE A-6 Local Disks Worksheet

| Local Disk Name | Size |
|-----------------|------|
|                 |      |
|                 |      |
|                 |      |
|                 |      |
|                 |      |
|                 |      |
|                 |      |
|                 |      |

TABLE A-7 Other Local Devices Worksheet

| Device Type | Name |
|-------------|------|
|             |      |
|             |      |
|             |      |
|             |      |

## Example: Local Devices Worksheets

Node name: **phys-schost-1**

TABLE A-8 Example: Local Disks Worksheet

| Local Disk Name | Size      |
|-----------------|-----------|
| <b>c0t0d0</b>   | <b>2G</b> |
| <b>c0t1d0</b>   | <b>2G</b> |
| <b>c1t0d0</b>   | <b>2G</b> |
| <b>c1t1d0</b>   | <b>2G</b> |
|                 |           |
|                 |           |
|                 |           |
|                 |           |

TABLE A-9 Example: Other Local Devices Worksheet

| Device Type | Name              |
|-------------|-------------------|
| tape        | <b>/dev/rmt/0</b> |
|             |                   |
|             |                   |
|             |                   |

## Device Group Configurations Worksheet

Volume manager (circle one):

Solaris Volume Manager | VxVM

**TABLE A-10** Device Groups Worksheet

| Disk Group/<br>Disk Set Name | Node Names<br>(indicate priority if ordered list) | Ordered priority?<br>(circle one) | Failback?<br>(circle one) |
|------------------------------|---------------------------------------------------|-----------------------------------|---------------------------|
|                              |                                                   | Yes   No                          | Yes   No                  |
|                              |                                                   | Yes   No                          | Yes   No                  |
|                              |                                                   | Yes   No                          | Yes   No                  |
|                              |                                                   | Yes   No                          | Yes   No                  |
|                              |                                                   | Yes   No                          | Yes   No                  |
|                              |                                                   | Yes   No                          | Yes   No                  |
|                              |                                                   | Yes   No                          | Yes   No                  |
|                              |                                                   | Yes   No                          | Yes   No                  |
|                              |                                                   | Yes   No                          | Yes   No                  |
|                              |                                                   | Yes   No                          | Yes   No                  |
|                              |                                                   | Yes   No                          | Yes   No                  |
|                              |                                                   | Yes   No                          | Yes   No                  |
|                              |                                                   | Yes   No                          | Yes   No                  |
|                              |                                                   | Yes   No                          | Yes   No                  |

## Example: Device Group Configurations Worksheet

Volume manager (circle one):

Solaris Volume Manager

TABLE A-11 Example: Device Groups Configurations Worksheet

| Disk Group/ Disk Set Name | Node Names<br>(indicate priority if ordered list) | Ordered priority?<br>(circle one) | Failback?<br>(circle one) |
|---------------------------|---------------------------------------------------|-----------------------------------|---------------------------|
| <b>dg-schost-1</b>        | <b>1) phys-schost-1,<br/>2) phys-schost-2</b>     | <b>Yes</b>                        | <b>Yes</b>                |
|                           |                                                   | Yes   No                          | Yes   No                  |
|                           |                                                   | Yes   No                          | Yes   No                  |
|                           |                                                   | Yes   No                          | Yes   No                  |
|                           |                                                   | Yes   No                          | Yes   No                  |
|                           |                                                   | Yes   No                          | Yes   No                  |
|                           |                                                   | Yes   No                          | Yes   No                  |
|                           |                                                   | Yes   No                          | Yes   No                  |
|                           |                                                   | Yes   No                          | Yes   No                  |
|                           |                                                   | Yes   No                          | Yes   No                  |
|                           |                                                   | Yes   No                          | Yes   No                  |
|                           |                                                   | Yes   No                          | Yes   No                  |
|                           |                                                   | Yes   No                          | Yes   No                  |
|                           |                                                   | Yes   No                          | Yes   No                  |











# Index

---

## A

### adapters

#### local MAC address

changes during upgrade, 210, 231

NIC support, 23

required setting, 23

SBus SCI restriction, 29

#### SCI-PCI

installing Solaris packages, 52, 87

installing Sun Cluster packages, 88

installing Sun Cluster packages with JumpStart, 87

package requirements, 16

#### tagged VLAN

cluster interconnect guidelines, 28

public network guidelines, 22

### adding

*See also* configuring

*See also* installing

cluster file system mount points to new nodes, 52

drives to a disk set, 150-151

mediator hosts, 156-157

nodes to the Sun Cluster module to Sun Management Center, 186-187

### administrative console

installing CCP software, 47-50

IP addresses, 21

MANPATH, 49

PATH, 49

alternate boot environment, disk space requirements, 196

alternate boot path, displaying, 133

### Apache

disabling run control scripts, 206

modifying scripts during upgrade, 202, 224

applications, installing on a cluster file system, 33

authentication, *See* authorized-node list

### authorized-node list

adding nodes, 264

removing nodes, 107

autoboot property, 180

automatic power-saving shutdown, restriction, 15

automatic reboot on disk-path failures

enabling, 70

verifying, 70

autoscinstall.class file, 86

Availability Suite software

preparing for cluster upgrade, 200, 216, 238

## B

boot devices, alternate boot path, 133

### boot environment

creating for live upgrade, 241

disk space requirements, 196

### booting

into noncluster mode, 264

into single-user noncluster mode, 204, 225

## C

cconsole command, 49

installing the software, 47-50

using, 50, 88

ccp command, 49

ce\_taskq\_disable variable, 53

### changing

private hostnames, 122-123

- changing (*Continued*)
  - private IP address range, 98-103
- claccess command
  - adding nodes to the authorized-node list, 264
  - removing nodes from the authorized-node list, 107
- class file, modifying, 86
- cldevice command
  - determining device-ID names, 119
  - error messages, 250
  - migrating device IDs after upgrade, 250, 260
  - updating the global-devices namespace, 148
  - verifying command processing, 148
- cldevicegroup command
  - enabling the localonly property, 133, 137, 141, 144
  - removing nodes from raw-disk device groups, 133, 136, 140, 144, 165
  - verifying disk-group configurations, 170-172
- clnode command
  - moving resource groups and device groups, 164
  - verifying cluster mode, 250
  - verifying device-ID migration, 250
  - viewing private hostnames, 122
- clsetup command
  - adding cluster interconnects, 96
  - changing private hostnames, 122
  - postinstallation setup, 119
  - registering device groups, 168
- cluster command
  - adding nodes, 110-114
  - creating new clusters, 72-79
  - verifying installation mode, 121
- Cluster Control Panel (CCP) software
  - installing, 47-50
  - starting, 49
- cluster file systems
  - See also* shared file systems
  - adding to a new node, 52
  - caution notice, 175
  - configuring, 175-179
  - installing applications, 33
  - LOFS restriction, 32
  - mount options, 177
  - planning, 31-36
  - restrictions
    - communication end-points, 33
    - fatattach command, 33
- cluster file systems, restrictions (*Continued*)
  - forcedirectio, 36
  - LOFS, 15
  - quotas, 32
  - VxFS, 36
  - verifying the configuration, 177
- cluster interconnects
  - changing the private IP address range, 98-103
  - configuring on a single-node cluster, 96
  - planning, 27-29
- cluster mode, verifying, 250
- cluster name, 25
- cluster nodes
  - adding nodes
    - by using an XML file, 110-114
    - by using JumpStart, 79-95
    - by using scinstall, 103-110
    - updating quorum devices with new nodes, 114-117
  - adding to the Sun Cluster module to Sun Management Center, 186-187
  - creating the cluster
    - by using an XML file, 72-79
    - by using JumpStart, 79-95
    - by using scinstall, 64-72
  - determining the node-ID number, 172
  - planning, 25
  - verifying
    - cluster mode, 250
    - installation mode, 121
    - status, 69, 75, 93, 107
- clusters file, administrative console, 48
- clvxvm command, installing VxVM, 160-162
- common agent container, enabling the daemon, 84
- communication end-points, restriction on cluster file systems, 33
- configuring
  - additional nodes
    - by using an XML file, 110-114
    - by using JumpStart, 79-95
    - by using scinstall, 103-110
  - cluster file systems, 175-179
  - cluster interconnects on a single-node cluster, 96
  - disk sets, 147-150
  - IPMP groups, 53
  - local disk groups, 167-168
  - md.tab file, 152-154

configuring (*Continued*)  
 multipathing software, 55-57  
 Network Time Protocol (NTP), 123-125  
 new clusters  
   by using an XML file, 72-79  
   by using JumpStart, 79-95  
   by using `scinstall`, 64-72  
 non-global zones, 179-182  
 quorum devices, 117-120  
 Solaris Volume Manager, 127-146  
 state database replicas, 129-130  
 user work environment, 61  
 VERITAS Volume Manager (VxVM), 159-166  
 console-access devices  
   IP addresses, 21  
   planning, 21-22  
   serial-port numbers, 48  
 CVM, *See* VERITAS Volume Manager (VxVM) cluster feature

## D

data services  
 installing, 58-60  
 upgrading  
   Sun Cluster HA for SAP liveCache, 252  
   Sun Cluster HA for SAP Web Application Server, 210, 232, 254  
 departing device groups, 171  
 device groups  
   *See also* disk groups  
   *See also* raw-disk device groups  
 departing, 171  
 importing, 171  
 moving, 164  
 planning, 31-32  
 registering changes to, 169  
 registering disk groups as, 168  
 reminding, 169-170  
 replicated disks, 32  
 setting the replication property, 149  
 status, 171  
 verifying registration, 169

device-ID names  
 determining, 119

device-ID names (*Continued*)  
 migrating after upgrade, 260  
 DID driver, updating, 261  
 Dirty Region Logging (DRL), planning, 40  
 disabling  
   Apache, 206  
   installation mode, 120  
   LOFS, 70, 75, 92, 108  
   LOFS after upgrade, 211, 233  
   NTP daemon, 124  
   resources, 198, 199  
   spanning tree algorithm, 118  
 disk device groups, *See* device groups  
 disk drives, *See* drives  
 disk groups  
   *See also* device groups  
   configuring, 167-168  
   description of types, 166  
   registering as device groups, 168  
   verifying the configuration, 170-172  
 disk-path failures  
   enabling automatic reboot, 70  
   verifying enabled automatic reboot, 70  
 disk sets  
   adding drives, 150-151  
   configuring, 147-150  
   planning the maximum number, 39  
   repartitioning drives, 152  
   setting the maximum number, 128-129  
 disk strings, dual-string mediator requirements, 156  
 disks, *See* drives  
 disksets, *See* disk sets  
 domain console network interfaces, IP addresses, 21  
 drives  
   adding to disk sets, 150-151  
   mirroring differing device sizes, 42  
   repartitioning, 152  
 DRL, planning, 40  
 dual-partition upgrade, 214-236  
   preparing the cluster, 215-224  
   recovering from unrecoverable error, 258-260  
   requirements, 195  
   Solaris software, 224  
   Sun Cluster software, 229-236  
   troubleshooting, 236

- dual-string mediators
  - adding hosts, 156-157
  - overview, 155-158
  - planning, 38
  - repairing data, 157-158
  - restoring after upgrade, 253
  - status, 157
  - unconfiguring during upgrade, 199
- Dynamic Multipathing (DMP), 40
- E**
- enabling
  - automatic reboot on disk-path failures, 70
  - common agent container daemon, 84
  - kernel cage, 53
  - LOFS requirements, 15
  - multipathing, 56
  - NTP, 125
  - remote access, 68
- encapsulated root disks
  - configuring, 162-163
  - mirroring, 164-166
  - planning, 40
  - unconfiguring, 172-174
- Enclosure-Based Naming, planning, 40
- error messages
  - cluster, 12
  - metainit command, 134
  - NTP, 95
  - scdidadm command, 250
- establishing a new cluster, by using an XML file, 72-79
- /etc/clusters file, 48
- /etc/inet/hosts file
  - configuring, 53, 84
  - planning, 21
- /etc/inet/ntp.conf.cluster file
  - configuring, 123-125
  - starting NTP, 125
  - stopping NTP, 124
- /etc/inet/ntp.conf file
  - changes during upgrade, 210, 231
  - configuring, 123-125
  - starting NTP, 125
  - stopping NTP, 124
- /etc/init.d/xntpd.cluster command, starting NTP, 125
- /etc/init.d/xntpd command
  - starting NTP, 125
  - stopping NTP, 124
- /etc/lvm/md.tab file, 152-154
- /etc/name\_to\_major file
  - non-VxVM nodes, 52, 161
  - VxVM-installed nodes, 161
- /etc/nsswitch.conf file, non-global zone modifications, 181
- /etc/serialports file, 48
- /etc/system file
  - ce adapter setting, 53
  - kernel\_cage\_enable variable, 53
  - LOFS setting, 70, 75, 92, 108, 211, 233
  - stack-size setting, 57
  - thread stack-size setting, 169
- /etc/vfstab file
  - adding mount points, 177
  - modifying during upgrade
    - dual-partition upgrade, 225
    - standard upgrade, 203
  - verifying the configuration, 177
- examples
  - activating volumes in the md.tab file, 155
  - adding drives to a disk set, 151
  - adding mediator hosts, 157
  - assigning a minor number to a device group, 170
  - configuring Sun Cluster software on additional nodes
    - by using scinstall, 108-110
  - configuring Sun Cluster software on all nodes
    - by using an XML file, 77
    - by using scinstall, 71
  - creating a cluster file system, 178-179
  - creating a disk set, 150
  - creating state database replicas, 130
  - live upgrade, 247-248
  - live upgrade verification, 251
  - md.tab file, 153-154
  - mirroring file systems that can be unmounted, 145
  - mirroring file systems that cannot be unmounted, 141-142
  - mirroring the encapsulated root disk, 166
  - mirroring the global devices namespace, 137-138
  - mirroring the root (/) file system, 133-134



examples (*Continued*)

- unloading the RSMRDT driver, 270-271
- updating SCSI quorum devices after adding a node, 116-117

**F**

- failover file system, *See* highly available local file system
- fat attach command, restriction on cluster file systems, 33
- file-system logging, planning, 41
- finish script, JumpStart, 87
- forcedirectio command, restriction, 36

**G**

- global devices
  - caution notice, 173
  - /global/.devices/ directory
    - mirroring, 134-138
    - node@nodeid file system, 38
  - /globaldevices partition
    - creating, 51
    - planning, 17
  - node@nodeid file system, 194
  - planning, 31-36
  - updating the namespace, 148
- /global directory, 35
- global file systems
  - See* cluster file systems
- global zones, *See* Solaris zones

**H**

- HADB
  - disabling before upgrade, 195
  - disabling for dual-partition upgrade, 218
  - re-enabling after dual-partition upgrade, 254
- hardware RAID, internal disk mirroring, 54-55
- help, 12
- high-priority processes, restriction, 24
- hosts file
  - configuring, 53, 84
  - planning, 21

- hot spare disks, planning, 38

**I**

- importing device groups, 171
- inactive boot environment, 241
- initialization files, 61
- installation mode
  - disabling, 120
  - verifying, 121
- installing
  - See also* adding
  - See also* configuring
  - Cluster Control Panel (CCP), 47-50
  - data services, 58-60
  - in zones, 15
  - man pages, 48
  - multipathing software, 55-57
  - RSMAPI
    - Solaris packages, 52, 87
    - Sun Cluster packages, 59, 87
    - Sun Cluster packages with JumpStart, 88
  - RSMRDT drivers
    - Solaris packages, 52, 87
    - Sun Cluster packages, 59
    - Sun Cluster packages with JumpStart, 87
  - SCI-PCI adapters
    - Solaris packages, 52
    - Sun Cluster packages, 59
    - Sun Cluster packages with JumpStart, 87
  - Solaris software
    - alone, 50-54
    - with Sun Cluster software, 79-95
  - Sun Cluster Geographic Edition, 59
  - Sun Cluster Manager, 59
  - Sun Cluster software
    - packages, 58-60
    - verifying, 121-122
  - Sun Management Center
    - requirements, 184
    - Sun Cluster module, 185
  - Sun StorEdge QFS software, 60
  - Sun StorEdge Traffic Manager, 55-57
  - VERITAS File System (VxFS), 57
  - VERITAS Volume Manager (VxVM), 159-166

- integrated mirroring, 54-55
- internal hardware disk mirroring, 54-55
- IP addresses
  - changing the private IP address range, 98-103
  - planning the private network, 26-27
  - planning the public network, 21
- IP Filter, restriction, 16
- IP network multipathing (IPMP) groups
  - automatic creation during installation, 23
  - converting from NAFO groups, 194
  - upgrading from NAFO groups, 210, 231
- IPMP groups
  - See* IP network multipathing (IPMP) groups
  - configuring, 53
- IPv6 addresses
  - private network restriction, 27, 28
  - public network use, 22
  
- J**
- Java ES, removing the Sun Cluster entry, 266
- JumpStart
  - class file, 86
  - finish script, 87
  - installing Solaris and Sun Cluster software, 79-95
  - troubleshooting installation, 95
- junctions, *See* switches
  
- K**
- kernel\_cage\_enable variable, 53
- /kernel/drv/md.conf file, 38
  - caution notice, 39, 129
  - configuring, 128-129
- /kernel/drv/scsi\_vhci.conf file, 56
  
- L**
- licenses, planning, 20
- live upgrade, 236-249
  - preparing the cluster, 237-239
  - requirements, 195
  - Solaris software, 239-249
- live upgrade (*Continued*)
  - Sun Cluster software, 239-249
  - troubleshooting, 248
- loading the Sun Cluster module to Sun Management Center, 187-188
- local disk groups
  - configuring, 167-168
  - description, 166
- local MAC address
  - changes during upgrade, 210, 231
  - NIC support, 23
  - required setting, 23
- local volumes
  - unique name requirement, 38
- LOFS
  - disabling, 70, 75, 92, 108
  - disabling after upgrade, 211, 233
  - restriction, 15, 32
- log files, Sun Cluster installation, 69
- logging for cluster file systems, planning, 41
- logical addresses, planning, 22
- logical network interfaces, restriction, 29
- loopback file system (LOFS)
  - disabling, 70, 75, 92, 108
  - disabling after upgrade, 211, 233
  - restriction, 15, 32
  
- M**
- man pages, installing, 48
- MANPATH
  - administrative console, 49
  - cluster nodes, 61
- md.conf file
  - caution notice, 129
  - configuring, 128-129
  - planning, 38
- md\_nsets field
  - configuring, 128-129
  - planning, 39
- md.tab file, configuring, 152-154
- mediators, *See* dual-string mediators
- messages files, cluster, 12
- minor-number conflicts, repairing, 169-170

- mirroring
    - differing device sizes, 42
    - global devices namespace, 134-138
    - internal disks, 54-55
    - multihost disks, 42
    - planning, 42-43
    - root (/) file system, 131-134
    - root disks, 130-131
      - caution notice, 165
      - planning, 42-43
    - troubleshooting, 138
  - mount options for cluster file systems
    - QFS, 35
    - requirements, 177
    - UFS, 34
    - VxFS, 35,36
  - mount points
    - adding to new nodes, 52
    - cluster file systems, 35-36
    - modifying the /etc/vfstab file, 177
    - nested, 36
  - mpxio-disable parameter, 56
  - multi-user services
    - verifying, 69,75,91,107,180
  - multihost disks
    - mirroring, 42
    - planning, 38
  - multipathing software
    - enabling, 56
    - installing, 55-57
    - troubleshooting installation, 56
  - multiported disks, *See* multihost disks
- N**
- NAFO groups
    - See also* IPMP groups
    - converting to IPMP groups, 194
    - upgrading to IPMP groups, 210,231
  - name\_to\_major file
    - non-VxVM nodes, 52,161
    - VxVM-installed nodes, 161
  - naming conventions
    - cluster, 25
    - cluster nodes, 25
    - naming conventions (*Continued*)
      - local volumes, 38
      - private hostnames, 27
      - raw-disk devices, 176
      - tagged VLAN adapters, 28
      - zones, 25
    - Network Appliance NAS devices, configuring as quorum devices, 117-120
    - Network File System (NFS)
      - See also* Sun Cluster HA for NFS
      - guidelines for cluster nodes, 23-24
    - Network Time Protocol (NTP)
      - configuring, 123-125
      - error messages, 95
      - starting, 125
      - stopping, 124
    - NFS, *See* Network File System (NFS)
    - NIS servers, restriction for cluster nodes, 24
    - nmd field, configuring, 128-129
    - node lists
      - device groups, 38
      - raw-disk device groups
        - removing nodes from, 165
        - viewing, 165
    - nodes, *See* cluster nodes
    - non-global zones
      - See* Solaris zones
      - nsswitch.conf file modifications, 181
    - noncluster mode
      - booting into, 264
      - booting into single-user, 204,225
    - nsswitch.conf file, non-global zone modifications, 181
    - NTP
      - configuring, 123-125
      - error messages, 95
      - starting, 125
      - stopping, 124
    - ntp.conf.cluster file
      - configuring, 123-125
      - starting NTP, 125
      - stopping NTP, 124
    - ntp.conf file
      - changes during upgrade, 210,231
      - configuring, 123-125
      - starting NTP, 125

ntp.conf file (*Continued*)  
stopping NTP, 124

## O

online help, Sun Cluster module to Sun Management Center, 188  
/opt/SUNWcluster/bin/ directory, 49  
/opt/SUNWcluster/bin/cconsole command, 49  
installing the software, 47-50  
using, 50, 88  
/opt/SUNWcluster/bin/ccp command, 49  
/opt/SUNWcluster/man/ directory, 49  
Oracle Parallel Server, *See* Oracle RAC

## P

package installation  
Cluster Control Panel (CCP) software, 47-50  
data services, 58-60  
RSMAPI, 87  
RSMRDT drivers, 52, 88  
SCI-PCI adapters, 52  
Sun Cluster man pages, 48  
Sun Cluster software, 58-60  
partitions  
/global/.devices/node@nodeid file system, 194  
/globaldevices, 17, 51  
repartitioning drives, 152  
root (/) file system, 17-18  
swap, 16  
volume manager, 17  
patches, planning, 20-21  
PATH  
administrative console, 49  
cluster nodes, 61  
PCI adapters, *See* SCI-PCI adapters  
ports, *See* serial ports  
private hostnames  
changing, 122-123  
planning, 27  
verifying, 122  
private network  
changing the IP address range, 98-103

private network (*Continued*)  
IPv6 address restriction, 28  
planning, 26-27  
profile, JumpStart, 86  
public network  
IPv6 support, 22  
planning, 22-23

## Q

QFS, *See* Sun StorEdge QFS software  
quorum devices  
caution notice, 165  
initial configuration, 117-120  
and mirroring, 43  
NAS devices, 118  
planning, 29-30  
quorum servers, 117  
restriction for replicated devices, 30  
SCSI protocol setting, 30  
troubleshooting configuration, 120  
updating after adding a node, 114-117  
verifying, 121-122  
ZFS restriction, 30  
quorum server, requirements as a quorum device, 117  
quorum servers  
configuring as quorum devices, 117-120  
disabling the spanning tree algorithm, 118  
quotas, restriction on cluster file systems, 32

## R

RAID, restriction, 37  
raidctl command, 54-55  
rarpd service, restriction for cluster nodes, 24  
raw-disk device groups  
*See also* device groups  
removing nodes from node lists, 165  
viewing node lists, 165  
raw-disk devices, naming conventions, 176  
rebooting  
into noncluster mode, 264  
into single-user noncluster mode, 204, 225

- recovering
    - See also* repairing
    - See also* troubleshooting
    - dual-partition upgrade unrecoverable error, 258-260
    - storage configuration changes during
      - upgrade, 260-262
    - unsuccessful cluster node creation, 72
  - registering VxVM device groups, 168
  - Remote Shared Memory Application Programming Interface (RSMAPI)
    - package requirements, 16
    - Solaris packages
      - installing by using JumpStart, 87
      - installing by using pkgadd, 52
    - Sun Cluster packages
      - installing by using JumpStart, 87, 88
      - installing by using pkgadd, 59
  - removing
    - See also* uninstalling
    - Java ES product registry, 266
    - Sun Cluster software, 263-266
  - repairing
    - See also* recovering
    - See also* troubleshooting
    - mediator data, 157-158
    - minor-number conflicts, 169-170
    - storage reconfiguration during upgrade, 260-262
  - replicated devices
    - disk requirement, 32
    - restriction as quorum devices, 30
    - setting the replication property, 149
  - resource groups
    - making unmanaged, 199
    - moving, 164
  - resource types, registering after upgrade, 252
  - resources
    - disabling, 198, 199
    - listing, 199
  - root (/) file system, mirroring, 131-134
  - root disk groups
    - configuring
      - on encapsulated root disks, 162-163
      - on nonroot disks, 163-164
    - planning, 40
    - simple, 40
    - unconfiguring encapsulated root disks, 172-174
  - root disks
    - encapsulating, 162-163
    - internal disk mirroring, 54-55
    - mirroring, 130-131
      - caution notice, 165
      - planning, 42-43
    - unencapsulating, 172-174
  - root environment, configuring, 61
  - rootdg, *See* root disk groups
  - routers, restriction for cluster nodes, 24
  - RPC service, restricted program numbers, 24
  - rpcmod settings, 57
  - RSMAPI, *See* Remote Shared Memory Application Programming Interface (RSMAPI)
  - RSMRDT drivers
    - installing
      - Solaris packages, 52
      - Sun Cluster packages, 59
      - Sun Cluster packages with JumpStart, 87
    - package installation, 88
    - requirements, 87
    - Solaris packages, 87
    - troubleshooting unloading, 271
    - uninstalling, 269
    - unloading, 270
- ## S
- SBus SCI adapters, restriction, 29
  - sccheck command, vfstab file check, 177
  - SCI-PCI adapters
    - package requirements, 16
    - Solaris packages
      - installing by using JumpStart, 87
      - installing by using pkgadd, 52
    - Sun Cluster packages
      - installing by using JumpStart, 87, 88
      - installing by using pkgadd, 59
  - scinstall command
    - adding nodes, 103-110
    - adding nodes by using JumpStart, 79-95
    - creating the cluster, 64-72
    - creating the cluster by using JumpStart, 79-95
    - unconfiguring Sun Cluster software, 263-266
    - upgrading Sun Cluster software, 209, 231

- scinstall command (*Continued*)
  - verifying Sun Cluster software, 249
- scrgadm command, listing resources, 199
- SCSI devices
  - correcting reservations after adding a third node, 114-117
  - installing quorum devices, 117-120
  - protocol setting in quorum devices, 30
- SCSI fencing protocol, quorum devices, 30
- scswitch command
  - disabling resources, 199
  - making resource groups unmanaged, 199
  - taking resource groups offline, 199
- secondary root disks, 43
- security files, distributing upgraded files, 251
- serial ports
  - configuring on the administrative console, 48
  - Simple Network Management Protocol (SNMP), 184
- serialports file, 48
- Service Management Facility (SMF)
  - verifying online services, 69, 75, 91, 107, 180
- shared disk groups, description, 166
- shared file systems
  - See also* cluster file systems
  - required mount parameters for QFS, 35
- Simple Network Management Protocol (SNMP), port for
  - Sun Management Center, 184
- single-user noncluster mode
  - booting into, 204, 225
- SMF
  - verifying online services, 69, 75, 91, 107, 180
- SNMP, port for Sun Management Center, 184
- software RAID, restriction, 37
- Solaris software
  - installing
    - alone, 50-54
    - with Sun Cluster software, 79-95
  - planning, 14-20
    - /globaldevices file system, 18
    - partitions, 16-19
    - root (/) file system, 17-18
    - software groups, 16
    - volume managers, 18
  - restrictions
    - automatic power-saving shutdown, 15
    - interface groups, 15
- Solaris software, restrictions (*Continued*)
  - IP Filter, 16
  - zones, 15
- upgrading
  - dual-partition upgrade, 224
  - standard upgrade, 202
  - verifying device-ID migration, 250
- Solaris Volume Manager
  - coexistence with VxVM, 161
  - configuring, 127-146
  - disk sets
    - adding drives, 150-151
    - configuring, 147-150
    - repartitioning drives, 152
    - setting the maximum number, 128-129
  - dual-string mediators
    - adding hosts, 156-157
    - overview, 155-158
    - repairing bad data, 157-158
    - status, 157
  - error messages, 134
  - md.tab file, 152-154
  - mediators
    - See* dual-string mediators
  - mirroring
    - global devices namespace, 134-138
    - root (/) file system, 131-134
    - root disks, 130-131
  - planning, 38-39
  - raw-disk device names, 176
  - state database replicas, 129-130
  - transactional-volume logging, 41
  - volumes
    - activating, 154-155
    - planning the maximum number, 39
    - setting the maximum number, 128-129
- Solaris zones
  - autoboot property, 180
  - configuring, 179-182
  - guidelines, 19-20
  - installation requirements, 15
  - LOFS coexistence with Sun Cluster HA for NFS, 19-20
  - LOFS requirement, 15
  - naming conventions, 25
  - spanning tree algorithm, 118
  - SSP, *See* console-access devices

- stack-size setting, 57, 169
  - standard upgrade, 196-214
    - preparing the cluster, 197-202
    - Solaris software, 202
    - Sun Cluster software, 207-214
    - volume managers, 202
  - starting
    - Cluster Control Panel (CCP), 49
    - Sun Management Center, 186
  - state database replicas, configuring, 129-130
  - status
    - device groups, 171
    - dual-string mediators, 157
    - verifying, 121-122
  - Sun Cluster Geographic Edition
    - installing, 59
    - uninstalling, 202, 224, 238
    - upgrading, 257
  - Sun Cluster HA for NFS
    - restriction with LOFS, 15, 32
  - Sun Cluster HA for SAP liveCache, upgrading, 252
  - Sun Cluster HA for SAP Web Application Server
    - upgrading, 210, 232, 254
  - Sun Cluster HA for Sun Java System Application Server EE (HADB)
    - disabling before upgrade, 195
    - disabling for dual-partition upgrade, 218
    - re-enabling after dual-partition upgrade, 254
  - Sun Cluster Manager
    - See also* SunPlex Manager
    - installing, 59
    - required packages, 87
  - Sun Cluster module to Sun Management Center, 183-188
    - adding nodes, 186-187
    - installing, 185
    - loading, 187-188
    - online help, 188
    - requirements, 184
    - upgrade, 189-190
  - Sun Enterprise 10000 servers
    - dynamic reconfiguration support, 53
    - kernel\_cage\_enable variable, 53
    - serialports file, 49
  - Sun Fire 15000 servers
    - IP addresses, 21
    - serial-port numbers, 49
  - Sun Management Center
    - installation requirements, 184
    - starting, 186
    - stopping, 190
    - Sun Cluster module, 183-188
      - adding nodes, 186-187
      - installing, 185
      - loading, 187-188
      - online help, 188
      - upgrading, 189-190
    - upgrading, 190-192
  - Sun StorageTek Availability Suite software
    - preparing for cluster upgrade, 200, 216, 238
  - Sun StorEdge Availability Suite software
    - preparing for cluster upgrade, 200, 216, 238
  - Sun StorEdge QFS software
    - installing, 60
    - mounting shared file systems, 35
  - Sun StorEdge Traffic Manager software
    - enabling, 56
    - installing, 55-57
    - troubleshooting installation, 56
  - SunPlex Manager
    - See also* Sun Cluster Manager
    - uninstalling, 268-269
  - swap, planning, 16
  - switches, *See* transport switches
  - system controllers (SC), *See* console-access devices
  - system file
    - kernel\_cage\_enable variable, 53
    - stack-size setting, 57
    - thread stack-size setting, 169
  - System Service Processor (SSP), *See* console-access devices
- ## T
- tagged VLAN adapters
    - cluster interconnect guidelines, 28
    - public network guidelines, 22
  - technical support, 12
  - telnet command, serial-port numbers, 49
  - terminal concentrators (TC), *See* console-access devices
  - test IP addresses, converting from NAFO groups, 194
  - thread stack-size setting, 169
  - three-way mirroring, 42

- Traffic Manager software
  - enabling, 56
  - installing, 55-57
  - troubleshooting installation, 56
- transport adapters, *See* adapters
- transport junctions, *See* transport switches
- transport switches, planning, 29
- troubleshooting
  - See also* recovering
  - See also* repairing
  - adding nodes to the Sun Cluster module to Sun Management Center, 187
  - configuring
    - additional nodes, 110
    - new clusters, 72
    - quorum devices, 120
  - dual-partition upgrade, 236
  - installing multipathing software, 56
  - JumpStart installation, 95
  - live upgrade, 248
  - local VxVM disk groups, 171
  - mirroring root file systems, 134, 138
  - recovering from unrecoverable error during
    - dual-partition upgrade, 258-260
  - storage configuration changes during
    - upgrade, 260-262
  - Sun StorEdge Traffic Manager installation, 56
  - unloading the RSMRDT driver, 271
- U**
- UFS logging, planning, 41
- unconfiguring Sun Cluster software, 263-266
- unencapsulating the root disk, 172-174
- uninstalling
  - See also* removing
  - RSMRDT package, 269
  - Sun Cluster Geographic Edition, 202, 224, 238
  - SunPlex Manager, 268-269
- upgrading, 193-262
  - choosing an upgrade method, 195-196
  - dual-partition upgrade, 214-236
    - preparing the cluster, 215-224
    - Solaris software, 224
    - Sun Cluster software, 229-236
  - upgrading, dual-partition upgrade (*Continued*)
    - troubleshooting, 236
  - live upgrade, 236-249
    - preparing the cluster, 237-239
    - Solaris, 239-249
    - Sun Cluster software, 239-249
    - troubleshooting, 248
    - volume managers, 239-249
  - recovering
    - storage configuration changes, 260-262
    - unrecoverable error during dual-partition
      - upgrade, 258-260
  - requirements, 193-194, 195
  - resource types, 252
  - restoring mediators, 253
  - standard upgrade, 196-214
    - preparing the cluster, 197-202
    - Solaris software, 202
    - Sun Cluster software, 207-214
    - volume managers, 202
  - Sun Cluster Geographic Edition, 257
  - Sun Cluster HA for SAP liveCache, 252
  - Sun Cluster HA for SAP Web Application Server, 210, 232, 254
  - Sun Cluster HA for Sun Java System Application Server EE (HADB), 218, 254
  - Sun Cluster module to Sun Management Center, 189-190
  - Sun Management Center, 190-192
  - Sun StorageTek Availability Suite configuration
    - device, 200, 216, 238
  - Sun StorEdge Availability Suite configuration
    - device, 200, 216, 238
  - troubleshooting
    - dual-partition upgrade, 236
    - live upgrade, 248
    - recovering from unrecoverable error during
      - dual-partition upgrade, 258-260
      - storage configuration changes, 260-262
  - unconfiguring mediators, 199
  - verifying
    - device-ID conversion, 250
    - Sun Cluster, 249-251
  - volume manager software
    - dual-partition upgrade, 224
    - live upgrade, 239-249



upgrading, volume manager software (*Continued*)  
     standard upgrade, 202  
     VxVM disk groups, 253  
 user-initialization files, modifying, 61  
 /usr/cluster/bin/ directory, 61  
 /usr/cluster/bin/claccess command  
     adding nodes to the authorized-node list, 264  
     removing nodes from the authorized-node list, 107  
 /usr/cluster/bin/cldevice command  
     determining device-ID names, 119  
     error messages, 250  
     migrating device IDs after upgrade, 250, 260  
     updating the global-devices namespace, 148  
     verifying command processing, 148  
 /usr/cluster/bin/cldevicegroup command  
     enabling the localonly property, 133, 137, 141, 144  
     removing nodes from raw-disk device groups, 133,  
         136, 140, 144, 165  
     verifying disk-group configurations, 170-172  
 /usr/cluster/bin/clnode command  
     moving resource groups and device groups, 164  
     verifying cluster mode, 250  
     verifying device-ID migration, 250  
     viewing private hostnames, 122  
 /usr/cluster/bin/clsetup command  
     adding cluster interconnects, 96  
     changing private hostnames, 122  
     postinstallation setup, 119  
     registering device groups, 168  
 /usr/cluster/bin/cluster command  
     adding nodes, 110-114  
     creating new clusters, 72-79  
     verifying installation mode, 121  
 /usr/cluster/bin/clvxvm command, installing  
     VxVM, 160-162  
 /usr/cluster/bin/sccheck command, vfstab file  
     check, 177  
 /usr/cluster/bin/scinstall command  
     adding nodes, 103-110  
     adding nodes by using JumpStart, 79-95  
     creating the cluster, 64-72  
     creating the cluster by using JumpStart, 79-95  
     unconfiguring Sun Cluster software, 263-266  
     verifying Sun Cluster software, 249  
 /usr/cluster/bin/scrgadm command, listing  
     resources, 199

/usr/cluster/bin/scswitch command  
     disabling resources, 199  
     making resource groups unmanaged, 199  
     taking resource groups offline, 199  
 /usr/cluster/man/ directory, 61

## V

/var/adm/messages file, 12  
 verifying  
     automatic reboot on disk-path failures, 70  
     cldevice command processing, 148  
     cluster node status, 69, 75, 93, 107  
     device-ID migration, 250  
     installation mode, 121  
     private hostnames, 122  
     quorum configurations, 121-122  
     SMF services, 69, 75, 91, 107, 180  
     Sun Cluster software version, 249  
     vfstab configuration, 177  
     VxVM disk-group configurations, 170-172  
 VERITAS File System (VxFS)  
     administering, 178  
     dual-partition upgrade, 224  
     installing, 57  
     mounting cluster file systems, 36, 178  
     planning, 36, 41  
     restrictions, 36  
     upgrading  
         dual-partition upgrade, 224  
         live upgrade, 239-249  
         standard upgrade, 202  
 VERITAS Volume Manager (VxVM)  
     cluster feature  
         creating shared disk groups, 167  
         installation requirement, 37  
     configuring, 159-166  
         disk groups, 167-168  
         local disk groups, 167-168  
         non-VxVM nodes, 161  
         volumes, 167-168  
     device groups  
         importing and deporting, 171  
         reminoring, 169-170

VERITAS Volume Manager (VxVM) *(Continued)*

- disk groups
  - configuring, 167-168
  - description of types, 166
  - registering, 168
  - troubleshooting local disk groups, 171
  - troubleshooting registration, 169
  - verifying configurations, 170-172
- dual-partition upgrade, 224
- encapsulating the root disk, 162-163
- Enclosure-Based Naming, 40
- installing, 159-166
- local disk groups
  - configuring, 167-168
  - description, 166
  - troubleshooting, 171
- mirroring the encapsulated root disk, 164-166
- planning, 18, 39-40
- raw-disk device names, 176
- root disk groups
  - configuring on nonroot disks, 163-164
  - configuring on root disks, 162-163
  - planning, 40, 160
  - simple, 40
  - unconfiguring from root disks, 172-174
- root disks
  - caution when unencapsulating, 173
  - encapsulating, 162-163
  - unencapsulating, 172-174
- shared disk groups
  - description, 166
- Solaris Volume Manager coexistence, 37
- unencapsulating the root disk, 172-174
- upgrading
  - dual-partition upgrade, 224
  - live upgrade, 239-249
  - standard upgrade, 202
- upgrading disk groups, 253
- vfstab file
  - adding mount points, 177
  - modifying during upgrade
    - dual-partition upgrade, 225
    - standard upgrade, 203
  - verifying the configuration, 177
- VLAN adapters
  - cluster interconnect guidelines, 28

VLAN adapters *(Continued)*

- public network guidelines, 22
- volume managers, 224
  - See also* Solaris Volume Manager
  - See also* VERITAS Volume Manager (VxVM)
- partitions for, 17
- planning
  - general, 36-43
  - Solaris Volume Manager, 38-39
  - VERITAS Volume Manager, 39-40
- upgrading
  - dual-partition upgrade, 224
  - live upgrade, 239-249
  - standard upgrade, 202
- volumes
  - Solaris Volume Manager
    - activating, 154-155
    - planning the maximum number, 39
    - setting the maximum number, 128-129
  - VxVM
    - configuring, 167-168
    - verifying, 171
- VxFS, *See* VERITAS File System (VxFS)
- vxio driver major number
  - non-VxVM nodes, 161
  - VxVM-installed nodes, 161
- VxVM, *See* VERITAS Volume Manager (VxVM)

**X**

- xntpd.cluster command, starting NTP, 125
- xntpd command
  - starting NTP, 125
  - stopping NTP, 124

**Z**

- Zettabyte File System (ZFS) storage pools, restriction
  - against adding quorum disks, 30
- ZFS storage pools, restriction against adding quorum disks, 30
- zones, *See* Solaris zones