

Oracle® Solaris Cluster Software Installation Guide

ORACLE®

Part No: E39580
September 2014, E39580-02

Copyright © 2000, 2014, Oracle and/or its affiliates. All rights reserved.

This software and related documentation are provided under a license agreement containing restrictions on use and disclosure and are protected by intellectual property laws. Except as expressly permitted in your license agreement or allowed by law, you may not use, copy, reproduce, translate, broadcast, modify, license, transmit, distribute, exhibit, perform, publish, or display any part, in any form, or by any means. Reverse engineering, disassembly, or decompilation of this software, unless required by law for interoperability, is prohibited.

The information contained herein is subject to change without notice and is not warranted to be error-free. If you find any errors, please report them to us in writing.

If this is software or related documentation that is delivered to the U.S. Government or anyone licensing it on behalf of the U.S. Government, the following notice is applicable:

U.S. GOVERNMENT END USERS. Oracle programs, including any operating system, integrated software, any programs installed on the hardware, and/or documentation, delivered to U.S. Government end users are "commercial computer software" pursuant to the applicable Federal Acquisition Regulation and agency-specific supplemental regulations. As such, use, duplication, disclosure, modification, and adaptation of the programs, including any operating system, integrated software, any programs installed on the hardware, and/or documentation, shall be subject to license terms and license restrictions applicable to the programs. No other rights are granted to the U.S. Government.

This software or hardware is developed for general use in a variety of information management applications. It is not developed or intended for use in any inherently dangerous applications, including applications that may create a risk of personal injury. If you use this software or hardware in dangerous applications, then you shall be responsible to take all appropriate fail-safe, backup, redundancy, and other measures to ensure its safe use. Oracle Corporation and its affiliates disclaim any liability for any damages caused by use of this software or hardware in dangerous applications.

Oracle and Java are registered trademarks of Oracle and/or its affiliates. Other names may be trademarks of their respective owners.

Intel and Intel Xeon are trademarks or registered trademarks of Intel Corporation. All SPARC trademarks are used under license and are trademarks or registered trademarks of SPARC International, Inc. AMD, Opteron, the AMD logo, and the AMD Opteron logo are trademarks or registered trademarks of Advanced Micro Devices. UNIX is a registered trademark of The Open Group.

This software or hardware and documentation may provide access to or information on content, products, and services from third parties. Oracle Corporation and its affiliates are not responsible for and expressly disclaim all warranties of any kind with respect to third-party content, products, and services. Oracle Corporation and its affiliates will not be responsible for any loss, costs, or damages incurred due to your access to or use of third-party content, products, or services.

Copyright © 2000, 2014, Oracle et/ou ses affiliés. Tous droits réservés.

Ce logiciel et la documentation qui l'accompagne sont protégés par les lois sur la propriété intellectuelle. Ils sont concédés sous licence et soumis à des restrictions d'utilisation et de divulgation. Sauf disposition de votre contrat de licence ou de la loi, vous ne pouvez pas copier, reproduire, traduire, diffuser, modifier, breveter, transmettre, distribuer, exposer, exécuter, publier ou afficher le logiciel, même partiellement, sous quelque forme et par quelque procédé que ce soit. Par ailleurs, il est interdit de procéder à toute ingénierie inverse du logiciel, de le désassembler ou de le décompiler, excepté à des fins d'interopérabilité avec des logiciels tiers ou tel que prescrit par la loi.

Les informations fournies dans ce document sont susceptibles de modification sans préavis. Par ailleurs, Oracle Corporation ne garantit pas qu'elles soient exemptes d'erreurs et vous invite, le cas échéant, à lui en faire part par écrit.

Si ce logiciel, ou la documentation qui l'accompagne, est concédé sous licence au Gouvernement des Etats-Unis, ou à toute entité qui délivre la licence de ce logiciel ou l'utilise pour le compte du Gouvernement des Etats-Unis, la notice suivante s'applique:

U.S. GOVERNMENT END USERS. Oracle programs, including any operating system, integrated software, any programs installed on the hardware, and/or documentation, delivered to U.S. Government end users are "commercial computer software" pursuant to the applicable Federal Acquisition Regulation and agency-specific supplemental regulations. As such, use, duplication, disclosure, modification, and adaptation of the programs, including any operating system, integrated software, any programs installed on the hardware, and/or documentation, shall be subject to license terms and license restrictions applicable to the programs. No other rights are granted to the U.S. Government.

Ce logiciel ou matériel a été développé pour un usage général dans le cadre d'applications de gestion des informations. Ce logiciel ou matériel n'est pas conçu ni n'est destiné à être utilisé dans des applications à risque, notamment dans des applications pouvant causer des dommages corporels. Si vous utilisez ce logiciel ou matériel dans le cadre d'applications dangereuses, il est de votre responsabilité de prendre toutes les mesures de secours, de sauvegarde, de redondance et autres mesures nécessaires à son utilisation dans des conditions optimales de sécurité. Oracle Corporation et ses affiliés déclinent toute responsabilité quant aux dommages causés par l'utilisation de ce logiciel ou matériel pour ce type d'applications.

Oracle et Java sont des marques déposées d'Oracle Corporation et/ou de ses affiliés. Tout autre nom mentionné peut correspondre à des marques appartenant à d'autres propriétaires qu'Oracle.

Intel et Intel Xeon sont des marques ou des marques déposées d'Intel Corporation. Toutes les marques SPARC sont utilisées sous licence et sont des marques ou des marques déposées de SPARC International, Inc. AMD, Opteron, le logo AMD et le logo AMD Opteron sont des marques ou des marques déposées d'Advanced Micro Devices. UNIX est une marque déposée d'The Open Group.

Ce logiciel ou matériel et la documentation qui l'accompagne peuvent fournir des informations ou des liens donnant accès à des contenus, des produits et des services émanant de tiers. Oracle Corporation et ses affiliés déclinent toute responsabilité ou garantie expresse quant aux contenus, produits ou services émanant de tiers. En aucun cas, Oracle Corporation et ses affiliés ne sauraient être tenus pour responsables des pertes subies, des coûts occasionnés ou des dommages causés par l'accès à des contenus, produits ou services tiers, ou à leur utilisation.

Contents

Using This Documentation	9
1 Planning the Oracle Solaris Cluster Configuration	11
Finding Oracle Solaris Cluster Installation Tasks	11
Planning the Oracle Solaris OS	12
Guidelines for Selecting Your Oracle Solaris Installation Method	12
Oracle Solaris OS Feature Requirements and Restrictions	13
System Disk Partitions	14
Volume Management Software	15
Planning the Oracle Solaris Cluster Environment	15
Oracle Solaris Cluster Software Version	16
Memory Requirements	16
Licensing	16
Software Updates	17
Geographic Edition	17
Public-Network IP Addresses	17
Console-Access Devices	18
Public Network Configuration	18
Quorum Server Configuration	20
NFS Guidelines	21
Service Restrictions	21
Network Time Protocol (NTP)	22
Oracle Solaris Cluster Configurable Components	23
SPARC: Guidelines for Oracle VM Server for SPARC in a Cluster	31
Zone Clusters	32
Planning Global Devices, Device Groups, and Cluster File Systems	36
Planning Global Devices	36
Planning Device Groups	37
Planning Cluster File Systems	37
Choosing Mount Options for UFS Cluster File Systems	38
Mount Information for Cluster File Systems	40

Planning Volume Management	41
Guidelines for Volume Manager Software	41
Guidelines for Solaris Volume Manager Software	42
UFS Cluster File System Logging	43
Mirroring Guidelines	43
2 Installing Software on Global-Cluster Nodes	45
Overview of Installing the Software	45
Installation Considerations	46
Installing the Software	46
▼ How to Prepare for Cluster Software Installation	47
▼ How to Install Oracle Solaris Software	48
▼ How to Install pconsole Software on an Administrative Console	52
▼ How to Install and Configure Oracle Solaris Cluster Quorum Server Software	55
▼ How to Configure Internal Disk Mirroring	57
▼ SPARC: How to Install Oracle VM Server for SPARC Software and Create Domains	58
▼ How to Install Oracle Solaris Cluster Software Packages	59
▼ How to Install the Availability Suite Feature of the Oracle Solaris Software	64
▼ How to Install Sun QFS Software	65
▼ How to Set Up the Root Environment	65
▼ How to Configure IP Filter	66
3 Establishing the Global Cluster	69
Overview of Establishing a New Cluster or Cluster Node	69
Establishing a New Global Cluster or New Global-Cluster Node	71
Configuring Oracle Solaris Cluster Software on All Nodes (scinstall)	71
▼ How to Configure Oracle Solaris Cluster Software on All Nodes (XML)	81
Establishing a New Oracle Solaris Cluster With the Automated Installer	89
▼ How to Prepare the Cluster for Additional Global-Cluster Nodes	114
▼ How to Change the Private Network Configuration When Adding Nodes or Private Networks	116
Configuring Oracle Solaris Cluster Software on Additional Global-Cluster Nodes (scinstall)	122
▼ How to Configure Oracle Solaris Cluster Software on Additional Global- Cluster Nodes (XML File)	129
▼ How to Update Quorum Devices After Adding a Node to a Global Cluster	134

▼ How to Configure Quorum Devices	137
▼ How to Verify the Quorum Configuration and Installation Mode	142
▼ How to Change Private Hostnames	143
Configuring Network Time Protocol (NTP)	144
▼ How to Validate the Cluster	148
▼ How to Record Diagnostic Data of the Cluster Configuration	148
4 Configuring Solaris Volume Manager Software	151
Configuring Solaris Volume Manager Software	151
▼ How to Install Solaris Volume Manager	151
▼ How to Create State Database Replicas	152
Creating Disk Sets in a Cluster	153
▼ How to Create a Disk Set	153
Adding Drives to a Disk Set	155
▼ How to Repartition Drives in a Disk Set	157
▼ How to Create an md.tab File	158
▼ How to Activate Volumes	160
Configuring Dual-String Mediators	161
Requirements for Dual-String Mediators	162
▼ How to Add Mediator Hosts	162
▼ How to Check For and Fix Bad Mediator Data	163
5 Creating a Cluster File System	165
Creating Cluster File Systems	165
▼ How to Create Cluster File Systems	165
Creating Oracle ACFS File Systems	168
Sample Configurations of an Oracle ACFS File System	169
▼ How to Register and Configure the Framework Resource Groups	173
▼ How to Create an Oracle ACFS File System	176
▼ How to Register and Configure the Scalable Device-Group Resource Group	178
▼ How to Register and Configure the Oracle ASM Resource Group	179
▼ How to Register and Configure the Oracle ACFS Proxy Resource Group	182
▼ How to Create an Oracle Grid Infrastructure Resource for Interoperation With Oracle Solaris Cluster	183
6 Creating Zone Clusters	187
Overview of Creating and Configuring a Zone Cluster	187

Creating and Configuring a Zone Cluster	188
Creating a Zone Cluster	188
Adding File Systems to a Zone Cluster	203
Adding Local File Systems to a Specific Zone-Cluster Node	217
Adding Storage Devices to a Zone Cluster	220
7 Uninstalling Software From the Cluster	227
Uninstalling the Software	227
▼ How to Unconfigure Oracle Solaris Cluster Software to Correct Installation Problems	227
▼ How to Uninstall Oracle Solaris Cluster Quorum Server Software	230
▼ How to Unconfigure a Zone Cluster	231
▼ How to Uninstall the Availability Suite Feature of the Oracle Solaris Software	233
Index	235

Using This Documentation

- **Overview** – Contains guidelines and procedures to install and configure Oracle Solaris Cluster software.
- **Audience** – Experienced system administrators with extensive knowledge of Oracle software and hardware.
- **Required knowledge** – Knowledge of the Oracle Solaris operating system, of Oracle Solaris Cluster software, and expertise with the volume manager software that is used with Oracle Solaris Cluster software.

This document is not to be used as a planning or presales guide.

Product Documentation Library

Late-breaking information and known issues for this product are included in the documentation library at <http://www.oracle.com/pls/topic/lookup?ctx=E39579>.

Access to Oracle Support

Oracle customers have access to electronic support through My Oracle Support. For information, visit <http://www.oracle.com/pls/topic/lookup?ctx=acc&id=info> or visit <http://www.oracle.com/pls/topic/lookup?ctx=acc&id=trs> if you are hearing impaired.

Documentation Accessibility

For information about Oracle's commitment to accessibility, visit the Oracle Accessibility Program web site at <http://www.oracle.com/pls/topic/lookup?ctx=acc&id=docacc>.

Feedback

Provide feedback about this documentation at <http://www.oracle.com/goto/docfeedback>.

Planning the Oracle Solaris Cluster Configuration

This chapter provides planning information and guidelines specific to an Oracle Solaris Cluster 4.2 configuration.

This chapter contains the following overview information:

- [“Finding Oracle Solaris Cluster Installation Tasks” on page 11](#)
- [“Planning the Oracle Solaris OS” on page 12](#)
- [“Planning the Oracle Solaris Cluster Environment” on page 15](#)
- [“Planning Global Devices, Device Groups, and Cluster File Systems” on page 36](#)
- [“Planning Volume Management” on page 41](#)

Finding Oracle Solaris Cluster Installation Tasks

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

TABLE 1-1 Oracle Solaris Cluster Software Installation Task Information

Task	Instructions
Set up cluster hardware.	“Oracle Solaris Cluster 4.2 Hardware Administration Manual ” Documentation that shipped with your server and storage devices
Plan global-cluster software installation.	Chapter 1, “Planning the Oracle Solaris Cluster Configuration”
Install software packages. Optionally, install and configure Sun QFS software.	“Installing the Software” on page 46 Sun QFS documentation
Establish a new global cluster or a new global-cluster node.	“Establishing a New Global Cluster or New Global-Cluster Node” on page 71
Configure Solaris Volume Manager software.	“Configuring Solaris Volume Manager Software” on page 151

Task	Instructions
	“Solaris Volume Manager Administration Guide”
Configure cluster file systems, if used.	“How to Create Cluster File Systems” on page 165
<i>(Optional)</i> Create zone clusters.	“Creating and Configuring a Zone Cluster” on page 188
Plan, install, and configure resource groups and data services. Create highly available local file systems, if used.	“Oracle Solaris Cluster Data Services Planning and Administration Guide”
Develop custom data services.	“Oracle Solaris Cluster Data Services Developer’s Guide”

Planning the Oracle Solaris OS

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

- [“Guidelines for Selecting Your Oracle Solaris Installation Method” on page 12](#)
- [“Oracle Solaris OS Feature Requirements and Restrictions” on page 13](#)
- [“System Disk Partitions” on page 14](#)

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

Guidelines for Selecting Your Oracle Solaris Installation Method

You can install Oracle Solaris software from a local DVD-ROM or from a network installation server by using the Automated Installer (AI) installation method. In addition, Oracle Solaris Cluster software provides a custom method for installing both the Oracle Solaris OS and Oracle Solaris Cluster software by using the AI installation method. During AI installation of Oracle Solaris software, you choose to either install the OS with defaults accepted or run an interactive installation of the OS where you can customize the installation for components such as the boot disk and the ZFS root pool. If you are installing several cluster nodes, consider a network installation.

See [“How to Install and Configure Oracle Solaris and Oracle Solaris Cluster Software \(IPS Repositories\)” on page 96](#) for details about the `scinstall` AI installation method. See your Oracle Solaris installation documentation for details about standard Oracle Solaris installation methods and what configuration choices you must make during installation of the OS.

Oracle Solaris OS Feature Requirements and Restrictions

Consider the following points when you plan the use of the Oracle Solaris OS in an Oracle Solaris Cluster configuration:

- **Minimum Oracle Solaris Package** – Oracle Solaris Cluster software requires at least the `solaris-small-server` package of Oracle Solaris software.
- **Oracle Solaris Operating System** – Oracle Solaris Cluster 4.2 core software and Oracle Solaris Cluster 4.2 Quorum Server software require a minimum of Oracle Solaris 11.1 software. The Oracle Solaris 11.0 release is not supported for Oracle Solaris Cluster 4.2 software.
- **Oracle Solaris Zones** – Install Oracle Solaris Cluster framework software *only* in the global zone.
- **Loopback file system (LOFS)** – During cluster creation, LOFS capability is enabled by default. If the cluster meets both of the following conditions, you must disable LOFS to avoid switchover problems or other failures:
 - 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 HA for NFS.

- **Power-saving shutdown** – Automatic power-saving shutdown is not supported in Oracle Solaris Cluster configurations and should not be enabled. See the [poweradm\(1M\)](#) man page for more information.
- **Network Auto-Magic (NWAM)** – The Oracle Solaris Network Auto-Magic (NWAM) feature activates a single network interface and disables all others. For this reason, NWAM cannot coexist with the Oracle Solaris Cluster software and you must disable it before you configure or run your cluster.
- **IP Filter feature** – Oracle Solaris Cluster relies on IP network multipathing (IPMP) for public network monitoring. Any IP Filter configuration must be made in accordance with IPMP configuration guidelines and restrictions concerning IP Filter.
- **fssnap** – Oracle Solaris Cluster software does not support the `fssnap` command, which is a feature of UFS. However, you can use the `fssnap` command on local systems that are not controlled by Oracle Solaris Cluster software. The following restrictions apply to `fssnap` support:
 - The `fssnap` command is supported on local files systems that are not managed by Oracle Solaris Cluster software.
 - The `fssnap` command is not supported on cluster file systems.

- The `fsnap` command is not supported on local file systems under the control of `HAStoragePlus`.
- **Minimum Java Version** – Java 1.7 is the minimum version required to run the Oracle Solaris Cluster software.

System Disk Partitions

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

- **root (/)** – The primary space requirements for the root (/) file system are as follows:
 - The Oracle Solaris Cluster software itself occupies less than 40 Mbytes of space in the root (/) file system.
 - Each Oracle Solaris Cluster data service might use between 1 Mbyte and 5 Mbytes.
 - Solaris Volume Manager software requires less than 5 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.
 - 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 stand-alone server. Therefore, allow at least 100 Mbytes for log files.
 - The `lofi` device for the global-devices namespace requires 100 Mbytes of free space. Starting in Oracle Solaris Cluster 4.2, a dedicated partition is no longer used for the global-devices namespace.

Note - The `lofi` device that is created for the global-devices namespace is restricted to that use only. Do not use this device for any other purpose, and never unmount the device.

- **/var** – The Oracle Solaris 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 stand-alone server. Therefore, allow at least 100 Mbytes for the /var file system.
- **swap** – The combined amount of swap space that is allocated for Oracle Solaris and Oracle Solaris Cluster software must be no less than 750 Mbytes. For best results, add at least 512 Mbytes for Oracle Solaris Cluster software to the amount that is required by the Oracle Solaris OS. In addition, allocate any additional swap amount that is required by applications that are to run on the Oracle Solaris host.

Note - If you create an additional swap file, do not create the swap file on a global device. Use only a local disk as a swap device for the host.

- **Volume manager** – Create a 20-Mbyte partition on slice 6 for volume manager use.

To support Solaris Volume Manager, you can create this partition on one of the following locations:

- A local disk other than the ZFS root pool
- The ZFS root pool, if the ZFS root pool is on a partition rather than a disk

Set aside a slice for this purpose on each local disk. However, if you have only one local disk on an Oracle Solaris host, you might need to create three state database replicas in the same slice for Solaris Volume Manager software to function properly. See [“Solaris Volume Manager Administration Guide”](#) for more information.

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

Volume Management Software

The following table describes the volume managers that work with Oracle Solaris on the SPARC and x86 platforms.

Volume Manager	Cluster Feature
Solaris Volume Manager	Solaris Volume Manager for Sun Cluster
Oracle Solaris ZFS Volume Manager	not applicable

Planning the Oracle Solaris Cluster Environment

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

- [“Oracle Solaris Cluster Software Version” on page 16](#)
- [“Memory Requirements” on page 16](#)
- [“Licensing” on page 16](#)
- [“Software Updates” on page 17](#)

- [“Geographic Edition” on page 17](#)
- [“Public-Network IP Addresses” on page 17](#)
- [“Console-Access Devices” on page 18](#)
- [“Public Network Configuration” on page 18](#)
- [“Quorum Server Configuration” on page 20](#)
- [“NFS Guidelines” on page 21](#)
- [“Service Restrictions” on page 21](#)
- [“Network Time Protocol \(NTP\)” on page 22](#)
- [“Oracle Solaris Cluster Configurable Components” on page 23](#)
- [“SPARC: Guidelines for Oracle VM Server for SPARC in a Cluster” on page 31](#)
- [“Zone Clusters” on page 32](#)

For detailed information about Oracle Solaris Cluster components, see the [“Oracle Solaris Cluster Concepts Guide”](#).

Oracle Solaris Cluster Software Version

All nodes in a cluster must run the same version of Oracle Solaris Cluster software.

Memory Requirements

The Oracle Solaris Cluster 4.2 software requires the following memory requirements for every cluster node:

- Minimum of 1.5 Gbytes of physical RAM (2 Gbytes typical)
- Minimum of 6 Gbytes of available hard drive space

Actual physical memory and hard drive requirements are determined by the applications that are installed. Consult the application's documentation or contact the application vendor to calculate additional memory and hard drive requirements.

Licensing

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

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

Software Updates

After installing each software product, you must also install any required software updates. For proper cluster operation, ensure that all cluster nodes maintain the same update level.

For general guidelines and procedures for applying software updates, see [Chapter 11, “Updating Your Software,”](#) in [“Oracle Solaris Cluster System Administration Guide”](#).

Geographic Edition

If a zone cluster will be configured in an Oracle Solaris Cluster Geographic Edition (Geographic Edition) configuration, the zone cluster must meet the following requirements:

- Each zone-cluster node must have a public-network IP address that corresponds to the zone-cluster node's hostname.
- The zone-cluster node's public-network IP address must be accessible by all nodes in the Geographic Edition configuration's partner cluster.
- Each zone-cluster node must have a failover IP address that maps to the hostname that corresponds to the zone-cluster name.

If you plan to use the Oracle Solaris Cluster Manager GUI to administer Geographic Edition components, all cluster nodes must have the same root password. For more information about Oracle Solaris Cluster Manager, see [Chapter 13, “Using the Oracle Solaris Cluster GUI,”](#) in [“Oracle Solaris Cluster System Administration Guide”](#).

Public-Network IP Addresses

For information about the use of public networks by the cluster, see [“Public Network Adapters and IP Network Multipathing”](#) in [“Oracle Solaris Cluster Concepts Guide”](#).

You must set up a number of public-network IP addresses for various Oracle Solaris Cluster components. The number of addresses that you need depends on which components you include in your cluster configuration. Each Oracle Solaris host 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 global-cluster node, after you install Oracle Solaris software
- The local `/etc/inet/hosts` file on any exclusive-IP non-global zone

TABLE 1-2 Oracle Solaris Cluster Components That Use Public-Network IP Addresses

Component	Number of IP Addresses Needed
Administrative console	1 IP address per subnet
Global-cluster nodes	1 IP address per node, per subnet
Zone-cluster nodes	1 IP address per node, per subnet
Domain console network interface	1 IP address per domain
<i>(Optional)</i> 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 [“Planning for Network Deployment in Oracle Solaris 11.2”](#).

Console-Access Devices

You must have console access to all cluster nodes. A service processor (SP) is used to communicate between the administrative console and the global-cluster node consoles.

For more information about console access, see the [“Oracle Solaris Cluster Concepts Guide”](#).

You can use the Oracle Solaris `pconsole` utility to connect to the cluster nodes. The utility also provides a master console window from which you can propagate your input to all connections that you opened. For more information, see the `pconsole(1)` man page that is available when you install the Oracle Solaris `terminal/pconsole` package.

Public Network Configuration

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.

Alternatively, create virtual NICs on the same physical interface and assign different virtual NICs to the private and public networks.

- **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.
- **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 additional information about data services and resources, also see the [“Oracle Solaris Cluster Concepts Guide”](#).
- **IPv4** – Oracle Solaris Cluster software supports IPv4 addresses on the public network.
- **IPv6** – Oracle Solaris Cluster software supports IPv6 addresses on the public network 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.

Unless there are one or more non-link-local IPv6 public network interfaces in the public network configuration, 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. These groups are link-based with transitive probes.

If the configuration contains any non-link-local IPv6 public network interfaces, you must manually configure in IPMP groups all interfaces that will be used for data-service traffic. You can configure the IPMP groups either before or after the cluster is established.

The `scinstall` utility ignores adapters that are already configured in an IPMP group. You can use probe-based IPMP groups or link-based IPMP groups in a cluster. Probe-based IPMP groups, which test the target IP address, provide the most protection by recognizing more conditions that might compromise availability.

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 on IPMP groups, see [Chapter 2, “About IPMP Administration,”](#) in [“Administering TCP/IP Networks, IPMP, and IP Tunnels in Oracle Solaris 11.2”](#).

To modify IPMP groups after cluster installation, follow the guidelines in [“How to Administer IP Network Multipathing Groups in a Cluster”](#) in [“Oracle Solaris Cluster System Administration Guide”](#) and procedures in [Chapter 3, “Administering IPMP,”](#) in [“Administering TCP/IP Networks, IPMP, and IP Tunnels in Oracle Solaris 11.2”](#).

- **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. Oracle Solaris Cluster software does not support a `local-mac-address?` value of `false` for Ethernet adapters.

For more information about public-network interfaces, see [“Oracle Solaris Cluster Concepts Guide”](#).

Quorum Server Configuration

You can use Oracle Solaris Cluster Quorum Server software to configure a machine as a quorum server and then configure the quorum server as your cluster's quorum device. You can use a quorum server instead of or in addition to shared disks and NAS filers.

Consider the following points when you plan the use of a quorum server in an Oracle Solaris Cluster configuration.

- **Network connection** – The quorum-server computer must connect to your cluster through the public network on the same subnet that is used by the cluster nodes it serves. Otherwise, the quorum server might be unavailable to the cluster during a node reboot and prevent the cluster from forming.
- **Supported hardware** – The supported hardware platforms for a quorum server are the same as for a global-cluster node.
- **Operating system** – Oracle Solaris software requirements for Oracle Solaris Cluster software apply as well to Quorum Server software.
- **Restriction for non-global zones** – In the Oracle Solaris Cluster 4.2 release, a quorum server cannot be installed and configured in a non-global zone.
- **Service to multiple clusters** – You can configure a quorum server as a quorum device to more than one cluster.
- **Mixed hardware and software** – You do not have to configure a quorum server on the same hardware and software platform as the cluster or clusters for which it provides quorum. For example, a SPARC based machine that runs the Oracle Solaris 10 OS can be configured as a quorum server for an x86 based cluster that runs the Oracle Solaris 11 OS.
- **Spanning tree algorithm** – You must disable the spanning tree algorithm on the Ethernet switches for the ports that are connected to the cluster public network where the quorum server will run.
- **Using a cluster node as a quorum server** – You can configure a quorum server on a cluster node to provide quorum for clusters other than the cluster that the node belongs to. However, a quorum server that is configured on a cluster node is not highly available.

NFS Guidelines

Consider the following points when you plan the use of Network File System (NFS) in an Oracle Solaris Cluster configuration:

- **NFS client** – No Oracle Solaris Cluster node can be an NFS client of an HA for NFS exported file system that is being mastered on a node in the same cluster. Such cross-mounting of HA for NFS is prohibited. Use the cluster file system to share files among global-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 HA for NFS data service on the same cluster node. If you do, certain HA for NFS data-service activities might cause the NFS daemons to stop and restart, interrupting NFS services. However, you can safely run the 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` or `fcntl`) might interfere with the ability to restart the lock manager (`lockd`). During restart, a blocked local process might be granted a lock which might be intended to be reclaimed by a remote client. This situation would cause unpredictable behavior.
- **NFS security features** – Oracle Solaris Cluster software does not support the following options of the `share_nfs(1M)` command:
 - `secure`
 - `sec=dh`

However, Oracle Solaris 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 nfssrv:nfs_portmon=1` to the `/etc/system` file on cluster nodes.
- The use of Kerberos with NFS.
- **Fencing** – Zone clusters support fencing for all supported NAS devices, shared disks, and storage arrays.

Service Restrictions

Observe the following service restrictions for Oracle Solaris Cluster configurations:

- **Routers** – Do not configure cluster nodes as routers (gateways) due to the following reasons:
 - Routing protocols might inadvertently broadcast the cluster interconnect as a publicly reachable network to other routers, despite the setting of the `IFF_PRIVATE` flag on the interconnect interfaces.

- Routing protocols might interfere with the failover of IP addresses across cluster nodes that impact client accessibility.
- Routing protocols might compromise proper functionality of scalable services by accepting client network packets and dropping them, instead of forwarding the packets to other cluster nodes.
- **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.
- **Install servers** – Do not use an Oracle Solaris Cluster configuration to provide a highly available installation service on client systems.
- **RARP** – Do not use an Oracle Solaris Cluster configuration to provide an rarpd service.
- **Remote procedure call (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 Oracle Solaris 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** – Oracle Solaris 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

Oracle Solaris 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 Oracle Solaris Cluster kernel threads from acquiring needed CPU cycles.

Network Time Protocol (NTP)

Observe the following guidelines for NTP:

- **Synchronization** – 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.
- **Accuracy** – 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 [“Oracle Solaris Cluster Concepts Guide”](#) for further information about cluster time. For more information about NTP, see the `ntpd(1M)` man page that is delivered in the Oracle Solaris `service/network/ntp` package.

Oracle Solaris Cluster Configurable Components

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

- [“Global-Cluster Name” on page 23](#)
- [“Global-Cluster Node Names and Node IDs” on page 23](#)
- [“Private Network Configuration” on page 24](#)
- [“Private Hostnames” on page 26](#)
- [“Cluster Interconnect” on page 26](#)
- [“Global Fencing” on page 28](#)
- [“Quorum Devices” on page 29](#)

Global-Cluster Name

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

For information about naming a zone cluster, see [“Zone Clusters” on page 32](#).

Global-Cluster Node Names and Node IDs

The name of a node in a global cluster is the same name that you assign to the physical or virtual host when you install it with the Oracle Solaris OS. See the [`hosts\(4\)`](#) man page for information about naming requirements.

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

During Oracle Solaris Cluster configuration, you specify the names of all nodes that you are installing in the global cluster. The node name must be same as the output of the `uname -n` command..

A node ID number is assigned to each cluster node for intracluster use, beginning with the number 1. Node ID numbers are assigned to each cluster node in the order that the node becomes a cluster member. If you configure all cluster nodes in one operation, the node from which you run the `scinstall` utility is the last node assigned a node ID number. You cannot change a node ID number after it is assigned to a cluster node.

A node that becomes a cluster member is assigned the lowest available node ID number. If a node is removed from the cluster, its node ID becomes available for assignment to a new node. For example, if in a four-node cluster the node that is assigned node ID 3 is removed and a new node is added, the new node is assigned node ID 3, not node ID 5.

If you want the assigned node ID numbers to correspond to certain cluster nodes, configure the cluster nodes one node at a time in the order that you want the node ID numbers to be assigned. For example, to have the cluster software assign node ID 1 to `phys-schost-1`, configure that node as the sponsoring node of the cluster. If you next add `phys-schost-2` to the cluster established by `phys-schost-1`, `phys-schost-2` is assigned node ID 2.

For information about node names in a zone cluster, see [“Zone Clusters” on page 32](#).

Private Network Configuration

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

Oracle Solaris Cluster software uses the private network for internal communication among nodes and among non-global zones that are managed by Oracle Solaris Cluster software. An Oracle Solaris Cluster configuration requires at least two connections to the cluster interconnect on the private network. When you configure Oracle Solaris 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 default netmask (255.255.240.0). This IP address range supports a combined maximum of 64 nodes and non-global zones, a maximum of 12 zone clusters, 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 or software configuration can currently 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 utility also prompts you for the number of zone clusters that you want to support. The number of global-cluster nodes that you specify should also include the expected number of unclustered 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, zone clusters, and private networks that you specified. The calculated netmask might support more than the supplied number of nodes, including non-global zones, zone clusters, and private networks. The `scinstall` utility also calculates a second netmask that would be the minimum to support twice the number of nodes, zone clusters, 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 - Changing the cluster private IP address range might be necessary to support the addition of nodes, non-global zones, zone clusters, or private networks.

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 [“Oracle Solaris Cluster System Administration Guide”](#). You must bring down the cluster to make these changes.

However, the cluster can remain in cluster mode if you use the `cluster set-netprops` command to change only the netmask. For any zone cluster that is already configured in the cluster, the private IP subnets and the corresponding private IP addresses that are allocated for that zone cluster will also be updated.

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`. However, 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 provided that the clusters are on different private networks. Private IP network addresses are not accessible from outside the physical cluster.
- **Oracle VM Server for SPARC** - When guest domains are created on the same physical machine and are connected to the same virtual switch, the private network is shared by such guest domains and is visible to all these domains. Proceed with caution before you specify a private-network IP address range to the `scinstall` utility for use by a cluster of guest domains. Ensure that the address range is not already in use by another guest domain that exists on the same physical machine and shares its virtual switch.
- **VLANs shared by multiple clusters** – Oracle Solaris Cluster configurations support the sharing of the same private-interconnect VLAN among multiple clusters. You do not have

to configure a separate VLAN for each cluster. However, for the highest level of fault isolation and interconnect resilience, limit the use of a VLAN to a single cluster.

- **IPv6** – Oracle Solaris 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. However, internode communication on the private network does not use these IPv6 addresses.

See “[Planning for Network Deployment in Oracle Solaris 11.2](#)” 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 Oracle Solaris Cluster configuration of a global cluster or a zone cluster. These private hostnames follow the naming convention `clusternode $node-id$ -priv`, where *node-id* is the numeral of the internal node ID. During Oracle Solaris Cluster configuration, the node ID number is automatically assigned to each node when the node becomes a cluster member. A node of the global cluster and a node of a zone cluster can both have the same private hostname, but each hostname resolves to a different private-network IP address.

After a global cluster is configured, you can rename its private hostnames by using the `clsetup(1CL)` utility. Currently, you cannot rename the private hostname of a zone-cluster node.

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

For more information about the purpose and function of the cluster interconnect, see “[Cluster Interconnect](#)” in “[Oracle Solaris Cluster Concepts Guide](#)”.

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

During Oracle Solaris Cluster configuration, you specify configuration information for one or two cluster interconnects.

- If the number of available adapter ports is limited, you can use tagged VLANs to share the same adapter with both the private and public network. For more information, see the guidelines for tagged VLAN adapters in [“Transport Adapters” on page 27](#).
- You can set up from one to six cluster interconnects in a cluster. While a single cluster interconnect reduces the number of adapter ports that are used for the private interconnect, it provides no redundancy and less availability. If a single interconnect fails, the cluster is at a higher risk of having to perform automatic recovery. Whenever possible, install two or more cluster interconnects to provide redundancy and scalability, and therefore higher availability, by avoiding a single point of failure.

You can configure additional cluster interconnects, up to six interconnects total, after the cluster is established by using the `clsetup` utility.

For guidelines about cluster interconnect hardware, see [“Interconnect Requirements and Restrictions”](#) in [“Oracle Solaris Cluster 4.2 Hardware Administration Manual”](#). For general information about the cluster interconnect, see [“Cluster Interconnect”](#) in [“Oracle Solaris Cluster Concepts Guide”](#).

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-host cluster, you also specify whether your interconnect is a point-to-point connection (adapter to adapter) or uses a transport switch.

Consider the following guidelines and restrictions:

- **IPv6** – Oracle Solaris 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 for scalable data services, are derived from the local MAC addresses.
- **Tagged VLAN adapters** – Oracle Solaris Cluster software supports tagged Virtual Local Area Networks (VLANs) to share an adapter between the private cluster interconnect and the public network. You must use the `dladm create-vlan` command to configure the adapter as a tagged VLAN adapter before you configure it with the cluster.

To configure a tagged VLAN adapter for the cluster interconnect, 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 `net2`, the VLAN instance number would be calculated as $(1000*73)+2$. You would therefore specify the adapter name as `net73002` 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 “[Oracle Solaris Cluster 4.2 Hardware Administration Manual](#)”.

For information about creating and administering VLANs, see the `dladm(1M)` man page and Chapter 3, “[Configuring Virtual Networks by Using Virtual Local Area Networks](#),” in “[Managing Network Datalinks in Oracle Solaris 11.2](#)”.

- **SPARC: Oracle VM Server for SPARC guest domains** – Specify adapter names by their virtual names, `vnetN`, such as `vnet0` and `vnet1`. Virtual adapter names are recorded in the `/etc/path_to_inst` file.
- **Logical network interfaces** – Logical network interfaces are reserved for use by Oracle Solaris Cluster software.

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 Oracle Solaris host that hosts the adapter end of the cable. However, you cannot use the default port name for certain adapter types.

Clusters with three or more nodes *must* use transport switches. Direct connection between cluster nodes is supported only for two-host clusters. If your two-host 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.

Global Fencing

Fencing is a mechanism that is used by the cluster to protect the data integrity of a shared disk during split-brain situations. By default, the `scinstall` utility in Typical Mode leaves global fencing enabled, and each shared disk in the configuration uses the default global fencing setting of `prefer3`. With the `prefer3` setting, the SCSI-3 protocol is used.

If any device is unable to use the SCSI-3 protocol, the `pathcount` setting should be used instead, where the fencing protocol for the shared disk is chosen based on the number of DID paths that are attached to the disk. Non-SCSI-3 capable devices are limited to two DID device paths within the cluster. Fencing can be turned off for devices which do not support either SCSI-3 or SCSI-2 fencing. However, data integrity for such devices cannot be guaranteed during split-brain situations.

In Custom Mode, the `scinstall` utility prompts you whether to disable global fencing. For most situations, respond **No** to keep global fencing enabled. However, you can disable global fencing in certain situations.



Caution - If you disable fencing under other situations than the ones described, your data might be vulnerable to corruption during application failover. Examine this data corruption possibility carefully when you consider turning off fencing.

The situations in which you can disable global fencing are as follows:

- The shared storage does not support SCSI reservations.
If you turn off fencing for a shared disk that you then configure as a quorum device, the device uses the software quorum protocol. This is true regardless of whether the disk supports SCSI-2 or SCSI-3 protocols. Software quorum is a protocol in Oracle Solaris Cluster software that emulates a form of SCSI Persistent Group Reservations (PGR).
- You want to enable systems that are outside the cluster to gain access to storage that is attached to the cluster.

If you disable global fencing during cluster configuration, fencing is turned off for all shared disks in the cluster. After the cluster is configured, you can change the global fencing protocol or override the fencing protocol of individual shared disks. However, to change the fencing protocol of a quorum device, you must first unconfigure the quorum device. Then set the new fencing protocol of the disk and reconfigure it as a quorum device.

For more information about fencing behavior, see [“Failfast Mechanism”](#) in [“Oracle Solaris Cluster Concepts Guide”](#). For more information about setting the fencing protocol of individual shared disks, see the `cldevice(1CL)` man page. For more information about the global fencing setting, see the `cluster(1CL)` man page.

Quorum Devices

Oracle Solaris 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 [“Oracle Solaris Cluster Concepts Guide”](#).

During Oracle Solaris Cluster installation of a two-host cluster, you can choose to have the `scinstall` utility automatically configure an available shared disk in the configuration as a quorum device. The `scinstall` utility assumes that all available shared disks are supported as quorum devices.

If you want to use a quorum server or an Oracle ZFS Storage 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` utility.

Note - You do not need to configure quorum devices for a single-host 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-host cluster must have at least one quorum device, which can be a shared disk, a quorum server, or a NAS device. For other topologies, quorum devices are optional.
- **Odd-number rule** – If more than one quorum device is configured in a two-host cluster or in a pair of hosts 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 shared-disk quorum device is configured, its fencing protocol is automatically set to SCSI-2 in a two-host cluster or SCSI-3 in a cluster with three or more nodes.
- **Changing the fencing protocol of quorum devices** – For SCSI disks that are configured as a quorum device, you must unconfigure the quorum device before you can enable or disable its SCSI fencing protocol.
- **Software quorum protocol** – You can configure supported shared disks that do not support SCSI protocol, such as SATA disks, as quorum devices. You must disable fencing for such disks. The disks would then use the software quorum protocol, which emulates SCSI PGR. The software quorum protocol would also be used by SCSI-shared disks if fencing is disabled for such disks.
- **Replicated devices** – Oracle Solaris Cluster software does not support replicated devices as quorum devices.
- **ZFS storage pools** – Do not add a configured quorum device to a 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.

After 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 [“Oracle Solaris Cluster Concepts Guide ”](#).

SPARC: Guidelines for Oracle VM Server for SPARC in a Cluster

Consider the following points when you create an Oracle VM Server for SPARC I/O domain or guest domain on a physically clustered machine that is SPARC hypervisor capable:

- **SR-IOV devices** – An SR-IOV device is supported with a logical domain that is configured to run as a cluster node. See the [Oracle Solaris Cluster 4 Compatibility Guide \(http://www.oracle.com/technetwork/server-storage/solaris-cluster/overview/solariscluster4-compatibilityguide-1429037.pdf\)](http://www.oracle.com/technetwork/server-storage/solaris-cluster/overview/solariscluster4-compatibilityguide-1429037.pdf) for information about supported SR-IOV devices.
- **SCSI LUN requirement** – The virtual shared storage device, or virtual disk back end, of an Oracle VM Server for SPARC guest domain must be a full SCSI LUN in the I/O domain. You cannot use an arbitrary virtual device.
- **Fencing** – Do not export a storage LUN to more than one guest domain on the same physical machine unless you also disable fencing for that device. Otherwise, if two different guest domains on the same machine both are visible to a device, the device will be fenced whenever one of the guest domains halts. The fencing of the device will panic any other guest domain that subsequently tries to access the device.
- **Network isolation** – Guest domains that are located on the same physical machine but are configured in different clusters must be network isolated from each other. Use one of the following methods:
 - Configure the clusters to use different network interfaces in the I/O domain for the private network.
 - Use different network addresses for each of the clusters when you perform initial configuration of the clusters.
- **Networking in guest domains** – Network packets to and from guest domains must traverse service domains to reach the network drivers through virtual switches. Virtual switches use kernel threads that run at system priority. The virtual-switch threads must be able to acquire needed CPU resources to perform critical cluster operations, including heartbeats, membership, checkpoints, and so forth. Configuring virtual switches with the `mode=sc` setting enables expedited handling of cluster heartbeat packets. However, the reliability of other critical cluster operations can be enhanced by adding more CPU resources to the service domain under the following workloads:
 - High-interrupt load, for example, due to network or disk I/O. Under extreme load, virtual switches can preclude system threads from running for a long time, including virtual-switch threads.
 - Real-time threads that are overly aggressive in retaining CPU resources. Real-time threads run at a higher priority than virtual-switch threads, which can restrict CPU resources for virtual-switch threads for an extended time.
- **Non-shared storage** – For non-shared storage, such as for Oracle VM Server for SPARC guest-domain OS images, you can use any type of virtual device. You can back such

virtual devices by any implement in the I/O domain, such as files or volumes. However, do not copy files or clone volumes in the I/O domain for the purpose of mapping them into different guest domains of the same cluster. Such copying or cloning would lead to problems because the resulting virtual devices would have the same device identity in different guest domains. Always create a new file or device in the I/O domain, which would be assigned a unique device identity, then map the new file or device into a different guest domain.

- **Exporting storage from I/O domains** – If you configure a cluster that is composed of Oracle VM Server for SPARC I/O domains, do not export its storage devices to other guest domains that also run Oracle Solaris Cluster software.
- **Oracle Solaris I/O multipathing** – Do not run Oracle Solaris I/O multipathing software (MPxIO) from guest domains. Instead, run Oracle Solaris I/O multipathing software in the I/O domain and export it to the guest domains.
- **Virtual disk multipathing** - Do not configure the virtual disk multipathing feature of Oracle VM Server for SPARC on a logical domain that is configured as a cluster node.
- **Live migration restriction** - Live migration is not supported for logical domains that are configured to run as cluster nodes. However, logical domains that are configured to be managed by the HA for Oracle VM Server for SPARC data service can use live migration.

For more information about Oracle VM Server for SPARC, see the [“Oracle VM Server for SPARC 3.1 Administration Guide”](#).

Zone Clusters

A zone cluster is a cluster of Oracle Solaris non-global zones. You can use the `clsetup` utility to create a zone cluster and add a network address, file system, ZFS storage pool, or storage device. You can also use a command line interface (the `clzonecluster` utility) to create a zone cluster, make configuration changes, and remove a zone cluster. For more information about using the `clzonecluster` utility, see the `clzonecluster(1CL)` man page.

Supported brands for zone clusters are `solaris`, `solaris10`, and `labeled`. The `labeled` brand is used exclusively in a Trusted Extensions environment. To use the Trusted Extensions feature of Oracle Solaris, you must configure the Trusted Extensions feature for use in a zone cluster. No other use of Trusted Extensions is supported in an Oracle Solaris Cluster configuration.

You can also specify a shared-IP zone cluster or an exclusive-IP zone cluster when you run the `clsetup` utility.

- Shared-IP zone clusters work with `solaris` or `solaris10` brand zones. A shared-IP zone cluster shares a single IP stack between all the zones on the node, and each zone is allocated an IP address.
- Exclusive-IP zone clusters work with `solaris` and `solaris10` brand zones. An exclusive-IP zone cluster supports a separate IP instance stack.

Consider the following points when you plan the creation of a zone cluster:

- [“Global-Cluster Requirements and Guidelines” on page 33](#)
- [“Zone-Cluster Requirements and Guidelines” on page 33](#)
- [“Guidelines for Trusted Extensions in a Zone Cluster” on page 35](#)

Global-Cluster Requirements and Guidelines

- **Global cluster** – The zone cluster must be configured on a global Oracle Solaris Cluster configuration. A zone cluster cannot be configured without an underlying global cluster.
- **Cluster mode** – The global-cluster node from which you create or modify a zone cluster must be in cluster mode. If any other nodes are in noncluster mode when you administer a zone cluster, the changes that you make are propagated to those nodes when they return to cluster mode.
- **Adequate private-IP addresses** – The private IP-address range of the global cluster must have enough free IP-address subnets for use by the new zone cluster. If the number of available subnets is insufficient, the creation of the zone cluster fails.
- **Changes to the private IP-address range** – The private IP subnets and the corresponding private IP-addresses that are available for zone clusters are automatically updated if the global cluster's private IP-address range is changed. If a zone cluster is deleted, the cluster infrastructure frees the private IP-addresses that were used by that zone cluster, making the addresses available for other use within the global cluster and by any other zone clusters that depend on the global cluster.
- **Supported devices** – Devices that are supported with Oracle Solaris zones can be exported to a zone cluster. Such devices include the following:
 - Oracle Solaris disk devices (*cNtXdYsZ*)
 - DID devices (*/dev/did/*dsk/dN*)
 - Solaris Volume Manager and Solaris Volume Manager for Sun Cluster multi-owner disk sets (*/dev/md/setname/*dsk/dN*)

Zone-Cluster Requirements and Guidelines

- **Distribution of nodes** – You cannot host multiple nodes of the same zone cluster on the same host machine. A host can support multiple zone-cluster nodes as long as each zone-cluster node on that host is a member of a different zone cluster.
- **Node creation** – You must create at least one zone-cluster node at the time that you create the zone cluster. You can use the `clsetup` utility or the `clzonecluster` command to create the zone cluster. The name of the zone-cluster node must be unique within the zone cluster. The infrastructure automatically creates an underlying non-global zone on each host that supports the zone cluster. Each non-global zone is given the same zone name, which is derived from, and identical to, the name that you assign to the zone cluster when you create

the cluster. For example, if you create a zone cluster that is named `zc1`, the corresponding non-global zone name on each host that supports the zone cluster is also `zc1`.

- **Cluster name** – Each zone-cluster name must be unique throughout the cluster of machines that host the global cluster. The zone-cluster name cannot also be used by a non-global zone elsewhere in the cluster of machines, nor can the zone-cluster name be the same as that of a global-cluster node. You cannot use “all” or “global” as a zone-cluster name, because these are reserved names.
- **Public-network IP addresses** – You can optionally assign a specific public-network IP address to each zone-cluster node.

Note - If you do not configure an IP address for each zone cluster node, two things will occur:

- That specific zone cluster will not be able to configure NAS devices for use in the zone cluster. The cluster uses the IP address of the zone cluster node when communicating with the NAS device, so not having an IP address prevents cluster support for fencing NAS devices.
- The cluster software will activate any Logical Host IP address on any NIC.

-
- **Private hostnames** – During creation of the zone cluster, a private hostname is automatically created for each node of the zone cluster, in the same way that hostnames are created in global clusters. Currently, you cannot rename the private hostname of a zone-cluster node. For more information about private hostnames, see [“Private Hostnames” on page 26](#).
 - **Oracle Solaris Zones brands** – All nodes of a zone cluster are configured as non-global zones of the `solaris`, `solaris10`, or `labeled` brand that is set with the `cluster` attribute. No other brand types are permitted in a zone cluster.
For Trusted Extensions, you must use only the `labeled` brand.
 - **IP type** - You can create a zone cluster that is either the shared IP type or the exclusive IP type. If the IP type is not specified, a shared-IP zone cluster is created by default.
 - **Global_zone=TRUE resource-type property** – To register a resource type that uses the `Global_zone=TRUE` resource-type property, the resource-type file must reside in the `/usr/cluster/global/rgm/rtreg/` directory of the zone cluster. If that resource-type file resides in any other location, the command to register the resource type is rejected.
 - **Conversion to a zone-cluster node** – You cannot add to a zone cluster a non-global zone that resides outside that zone cluster. You must use only the `clzonecluster` command to add new nodes to a zone cluster.
 - **File systems** – You can use the `clsetup` utility or the `clzonecluster` command to add the following types of file systems for use by the zone cluster. A file system is exported to a zone cluster by using either a direct mount or a loopback mount. Adding a file system with the `clsetup` utility is done in cluster scope, which affects the entire zone cluster.
 - By direct mount:
 - UFS local file system

- Sun QFS stand-alone file system
- Sun QFS shared file system, only when used to support Oracle RAC
- Oracle Solaris ZFS (exported as a data set)
- NFS from supported NAS devices
- By loopback mount:
 - UFS local file system
 - Sun QFS stand-alone file system
 - Sun QFS shared file system, only when used to support Oracle RAC
 - UFS cluster file system

You configure an `HAStoragePlus` or `ScalMountPoint` resource to manage the mounting of the file system.

Guidelines for Trusted Extensions in a Zone Cluster

Consider the following points when you use the Trusted Extensions feature of Oracle Solaris in a zone cluster:

- **Only zone-cluster support** – In an Oracle Solaris Cluster configuration with Trusted Extensions enabled, applications must run only in a zone cluster. No other non-global zones can be used on the cluster. You must use only the `clzonecluster` command to create a zone cluster. Do not use the `txzonemgr` command to create a non-global zone on a cluster that has Trusted Extensions enabled.
- **Trusted Extensions scope** – You can either enable or disable Trusted Extensions for the entire cluster configuration. When Trusted Extensions is enabled, all non-global zones in the cluster configuration must belong to one of the zone clusters. You cannot configure any other kind of non-global zone without compromising security.
- **IP addresses** – Each zone cluster that uses Trusted Extensions must use its own IP addresses. The special networking feature in Trusted Extensions that enables an IP address to be shared between multiple non-global zones is not supported with Oracle Solaris Cluster software.
- **Loopback mounts** – You cannot use loopback mounts that have write permissions in a zone cluster that uses Trusted Extensions. Use only direct mounts of file systems that permit write access, or use loopback mounts that have only read permissions.
- **File systems** – Do not configure in the zone cluster the global device that underlies a file system. Configure only the file system itself in the zone cluster.
- **Storage device name** – Do not add an individual slice of a storage device to a zone cluster. You must add the entire device to a single zone cluster. The use of slices of the same storage device in different zone clusters compromises the security of those zone clusters.
- **Application installation** – Install applications only in the zone cluster or in the global cluster and then export to the zone cluster by using read-only loopback mounts.

- **Zone cluster isolation** – When Trusted Extensions is used, the name of a zone cluster is a security label. In some cases, the security label itself might be information that cannot be disclosed, and the name of a resource or resource group might be a sensitive piece of information that cannot be disclosed. When an inter-cluster resource dependency or inter-cluster resource-group affinity is configured, the name of the other cluster becomes visible as well as the name of any affected resource or resource group. Therefore, before you establish any inter-cluster relationships, evaluate whether this information can be made visible according to the your requirements.

Planning Global Devices, Device Groups, and Cluster File Systems

This section provides the following information:

- [“Planning Global Devices” on page 36](#)
- [“Planning Device Groups” on page 37](#)
- [“Planning Cluster File Systems” on page 37](#)
- [“Choosing Mount Options for UFS Cluster File Systems” on page 38](#)
- [“Mount Information for Cluster File Systems” on page 40](#)

Planning Global Devices

For information about the purpose and function of global devices, see [“Global Devices” in “Oracle Solaris Cluster Concepts Guide”](#).

Oracle Solaris 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 data from a cluster file system is accessible from a non-global zone.

Planning Device Groups

For information about the purpose and function of device groups, see [“Device Groups” in “Oracle Solaris Cluster Concepts Guide”](#).

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 43](#) for additional guidelines. See [“Configuring Solaris Volume Manager Software” on page 151](#) 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.

Planning Cluster File Systems

For information about the purpose and function of cluster file systems, see [“Cluster File Systems” in “Oracle Solaris Cluster Concepts Guide”](#).

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 “Oracle Solaris Cluster Data Services Planning and Administration Guide”](#).

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.
- **Zone clusters** – You cannot configure cluster file systems that use UFS for use in a zone cluster. Use highly available local file systems instead. You can use a Sun QFS shared file system in a zone cluster only to support Oracle RAC.
- **Loopback file system (LOFS)** – During cluster creation, LOFS is enabled by default. You must manually disable LOFS on each cluster node if the cluster meets both of the following conditions:

- HA for NFS (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 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 Oracle Solaris software by which one would put a communication endpoint in the file system namespace. Therefore, do not attempt to use the `fattach` command from any node other than the local node.
 - 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.
- **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.

Choosing Mount Options for *UFS* Cluster File Systems

This section describes requirements and restrictions for mount options of the following types of *UFS* cluster file systems:

Note - You can alternatively configure *this* and other types of file systems as highly available local file systems. For more information, see [“Enabling Highly Available Local File Systems”](#) in [“Oracle Solaris Cluster Data Services Planning and Administration Guide”](#).

Follow the guidelines in the following list of mount options *in the following list of mount options* to determine what mount options to use when you create your *UFS* cluster file systems.

`global`

Required. This option makes the file system globally visible to all nodes in the cluster.

`logging`

Required. This option enables logging.

`forcedirectio`

Conditional. This option is required only for cluster file systems that will host Oracle RAC RDBMS data files, log files, and control files.

`onerror=panic`

Required. You do not have to explicitly specify the `onerror=panic` mount option in the `/etc/vfstab` file. This mount option is already the default value if no other `onerror` mount option is specified.

Note - Only the `onerror=panic` mount option is supported by Oracle Solaris Cluster software. Do not use the `onerror=umount` or `onerror=lock` mount options. These mount options are not supported on cluster file systems for the following reasons:

- Use of the `onerror=umount` or `onerror=lock` 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 `onerror=umount` or `onerror=lock` 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.

A node might require rebooting to recover from these states.

`syncdir`

Optional. If you specify `syncdir`, you are guaranteed POSIX-compliant file system behavior for the `write()` system call. If a `write()` succeeds, then this mount option ensures that sufficient space is on the disk.

If you do not specify `syncdir`, the same behavior occurs that is seen with UFS file systems. When you do not specify `syncdir`, performance of writes that allocate disk blocks, such as when appending data to a file, can significantly improve. However, in some cases, without `syncdir` you would not discover an out-of-space condition (ENOSPC) until you close a file.

You see ENOSPC on close only during a very short time after a failover. With `syncdir`, 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.

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

See the [mount_ufs\(1M\)](#) man page for more information about UFS mount options.

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.

- **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`. Ignoring 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 for cluster file systems on UFS. You can nest the mount points if the devices for the two file systems have the same physical host connectivity, for example, different slices on the same disk.

Note - This restriction still applies to Sun QFS shared file systems even if the two file system devices have the same physical host connectivity.

- `forcedirectio` – Oracle Solaris 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

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

- [“Guidelines for Volume Manager Software” on page 41](#)
- [“Guidelines for Solaris Volume Manager Software” on page 42](#)
- [“UFS Cluster File System Logging” on page 43](#)
- [“Mirroring Guidelines” on page 43](#)

Oracle Solaris Cluster software uses volume manager software to group disks into device groups that can then be administered as one unit. You must install Solaris Volume Manager software on all nodes of the cluster.

See your volume manager documentation and [“Configuring Solaris Volume Manager Software” on page 151](#) 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 “Oracle Solaris Cluster Concepts Guide ”](#) and [“Device Groups” in “Oracle Solaris Cluster Concepts Guide ”](#).

Guidelines for Volume Manager Software

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

- **Software RAID** – Oracle Solaris 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 43 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 ZFS root pool ensures high availability, but such mirroring is not required. See “[Mirroring Guidelines](#)” on page 43 for guidelines to help determine whether to mirror the ZFS root pool.
- **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 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 “[Oracle Solaris Cluster Data Services Planning and Administration Guide](#)” 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.
- **Hot-spare disks** – You can use hot-spare disks to increase availability, but hot spare disks are not required.

See your volume manager software 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:

- **Unique naming** – Disk set names must be unique throughout the cluster.
- **Disk set reserved names** – Do not name a disk set `admin` or `shared`.
- **Dual-string mediators** – A disk string consists of a disk enclosure, its physical disks, cables from the enclosure to the host or hosts, and the interface adapter cards. Each disk set configured with exactly two disk strings and mastered by exactly two or three Oracle Solaris hosts is called a dual-string disk set. This type of disk set must have Solaris Volume Manager dual-string mediators configured. Observe the following rules when you configure dual-string mediators:
 - You must configure each disk set with two or three hosts that act as mediator hosts.
 - You must use the hosts that can master a disk set as mediators for that disk set. If you have a campus cluster, you can also configure a third node or a non-clustered host on the cluster network as a third mediator host to improve availability.
 - 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.

UFS Cluster File System Logging

Logging is required for UFS cluster file systems. Oracle Solaris Cluster software supports *Oracle Solaris UFS logging*. See the [mount_ufs\(1M\)](#) man page for more information.

Mirroring Guidelines

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

- [“Guidelines for Mirroring Multihost Disks”](#) on page 43
- [“Guidelines for Mirroring the ZFS Root Pool”](#) on page 43

Guidelines for Mirroring Multihost Disks

Mirroring all multihost disks in an Oracle Solaris Cluster configuration enables the configuration to tolerate single-device failures. Oracle Solaris 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 supports three-way mirroring. However, Oracle Solaris 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 Devices”](#) in [“Oracle Solaris Cluster Concepts Guide”](#).

Guidelines for Mirroring the ZFS Root Pool

Oracle Solaris ZFS is the default root file system in the Oracle Solaris release. See [“How to Configure a Mirrored Root Pool \(SPARC or x86/VTOC\)”](#) in [“Managing ZFS File Systems in Oracle Solaris 11.2”](#) for instructions about how to mirror the ZFS root pool. Also see [Chapter](#)

[4, “Managing ZFS Root Pool Components,”](#) in [“Managing ZFS File Systems in Oracle Solaris 11.2”](#) for information about how to manage the different root pool components.

For maximum availability, mirror root (/), /usr, /var, /opt, and swap on the local disks. However, Oracle Solaris Cluster software does not require that you mirror the ZFS root pool.

Consider the following points when you decide whether to mirror the ZFS root pool:

- **Boot disk** – You can set up the mirror to be a bootable root pool. You can then boot from the mirror if the primary boot disk fails.
- **Backups** – Regardless of whether you mirror the root pool, 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 pool.
- **Separate controllers** – Highest availability includes mirroring the root pool on a separate controller.

Installing Software on Global-Cluster Nodes

This chapter provides the following information to install Oracle Solaris Cluster software on global-cluster nodes.

- [“Overview of Installing the Software” on page 45](#)
- [“Installing the Software” on page 46](#)

Overview of Installing the Software

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

TABLE 2-1 Task Map: Installing the Software

Task	Instructions
Plan the layout of your cluster configuration and prepare to install software.	“How to Prepare for Cluster Software Installation” on page 47
Install the Oracle Solaris OS on all nodes and optionally on an administrative console and a quorum server. Optionally, enable Oracle Solaris I/O multipathing on the nodes.	“How to Install Oracle Solaris Software” on page 48
(Optional) Install pconsole software on an administrative console.	“How to Install pconsole Software on an Administrative Console” on page 52
(Optional) Install and configure a quorum server.	“How to Install and Configure Oracle Solaris Cluster Quorum Server Software” on page 55
(Optional) Configure internal disk mirroring.	“How to Configure Internal Disk Mirroring” on page 57
(Optional) Install Oracle VM Server for SPARC software and create domains.	“How to Install Oracle VM Server for SPARC Software and Create Domains” on page 58
Install Oracle Solaris Cluster software and any data services that you will use.	“How to Install Oracle Solaris Cluster Software Packages” on page 59
(Optional) Install and configure the Availability Suite feature of Oracle Solaris software	“How to Install the Availability Suite Feature of the Oracle Solaris Software” on page 64
(Optional) Install Sun QFS software.	“How to Install Sun QFS Software” on page 65
Set up directory paths.	“How to Set Up the Root Environment” on page 65
(Optional) Configure the IP Filter feature of Oracle Solaris.	“How to Configure IP Filter” on page 66

Installation Considerations

The following table describes the file systems that work with the Oracle Solaris Cluster 4.2 software on SPARC and x86 platforms.

File System	Additional Information
Oracle Solaris UFS	
Sun QFS - Standalone file system	Supported data services: All failover data services External volume management: Solaris Volume Manager only
Sun QFS - Shared QFS file system	Supported data services: Oracle RAC External volume management: Solaris Volume Manager for Sun Cluster
Sun QFS - Shared QFS clients outside the cluster (SC-COTC)	Supported data services: None; only a shared file system is supported External volume management: No external volume manager is supported
Sun QFS - HA-SAM Failover	Supported data services: None; only a shared file system is supported External volume management: No external volume manager is supported
Oracle Solaris ZFS	
Network File System (NFS)	
Oracle Solaris Cluster proxy file system (PxFS)	Only a UFS file system can be configured as a cluster file system

For information about the supported versions of the file systems for this release, see [Oracle Solaris Cluster 4 Compatibility Guide \(http://www.oracle.com/technetwork/server-storage/solaris-cluster/overview/solariscluster4-compatibilityguide-1429037.pdf\)](http://www.oracle.com/technetwork/server-storage/solaris-cluster/overview/solariscluster4-compatibilityguide-1429037.pdf).

Installing the Software

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

- [“How to Prepare for Cluster Software Installation” on page 47](#)
- [“How to Install Oracle Solaris Software” on page 48](#)
- [“How to Install pconsole Software on an Administrative Console” on page 52](#)

- “How to Install and Configure Oracle Solaris Cluster Quorum Server Software” on page 55
- “How to Configure Internal Disk Mirroring” on page 57
- “How to Install Oracle VM Server for SPARC Software and Create Domains” on page 58
- “How to Install Oracle Solaris Cluster Software Packages” on page 59
- “How to Install the Availability Suite Feature of the Oracle Solaris Software” on page 64
- “How to Install Sun QFS Software” on page 65
- “How to Set Up the Root Environment” on page 65
- “How to Configure IP Filter” on page 66

▼ How to Prepare for Cluster Software Installation

1. **Ensure that the combination of hardware and software that you choose for your cluster is currently a supported Oracle Solaris Cluster configuration.**
 - See “Cluster Nodes” in “Oracle Solaris Cluster Concepts Guide ” for information about physical and virtual machines that are supported as cluster nodes.
 - See *Oracle Solaris Cluster 4 Compatibility Guide* (<http://www.oracle.com/technetwork/server-storage/solaris-cluster/overview/solariscluster4-compatibilityguide-1429037.pdf>) for information about supported software and hardware for this release.
 - Contact your Oracle 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.**
 - “Oracle Solaris Cluster 4.2 Release Notes ” – Restrictions, bug workarounds, and other late-breaking information.
 - “Oracle Solaris Cluster Concepts Guide ” - Overview of the Oracle Solaris Cluster product.
 - *Oracle® Solaris Cluster Software Installation Guide* (this manual) – Planning guidelines and procedures for installing and configuring Oracle Solaris, Oracle Solaris Cluster, and volume manager software.
 - “Oracle Solaris Cluster Data Services Planning and Administration Guide ” – 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:

 - Oracle Solaris OS

- Solaris Volume Manager software
- Sun QFS software
- Third-party applications

4. Plan your cluster configuration.

Use the planning guidelines in [Chapter 1, “Planning the Oracle Solaris Cluster Configuration”](#) and in the [“Oracle Solaris Cluster Data Services Planning and Administration Guide”](#) to determine how to install and configure your cluster.



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

5. Obtain all necessary updates for your cluster configuration.

For instructions on updating your software, see [Chapter 11, “Updating Your Software,”](#) in [“Oracle Solaris Cluster System Administration Guide”](#).

Next Steps

- If you want to install a machine as a quorum server to use as the quorum device in your cluster, go next to [“How to Install and Configure Oracle Solaris Cluster Quorum Server Software”](#) on page 55.
- Otherwise, if you want to use an administrative console to communicate with the cluster nodes, go to [“How to Install pconsole Software on an Administrative Console”](#) on page 52.
- Otherwise, choose the Oracle Solaris installation procedure to use.
 - To configure Oracle Solaris Cluster software by using the `scinstall(1M)` utility, go to [“How to Install Oracle Solaris Software”](#) on page 48 to first install Oracle Solaris software.
 - To install and configure both Oracle Solaris and Oracle Solaris Cluster software in the same operation (Automated Installer method), go to [“How to Install and Configure Oracle Solaris and Oracle Solaris Cluster Software \(IPS Repositories\)”](#) on page 96.

▼ How to Install Oracle Solaris Software

Use this procedure to install the Oracle Solaris OS on the following systems, as applicable to your cluster configuration:

1. (Optional) An administrative console that you will install with `pconsole` software. For more information, see [“How to Install pconsole Software on an Administrative Console”](#) on page 52.

2. (Optional) A quorum server. For more information, see [“How to Install and Configure Oracle Solaris Cluster Quorum Server Software”](#) on page 55.

3. Each node in the global cluster, if you will not use the `scinstall` custom Automated Installer method to install software. For more information about Automated Installer installation of a cluster, see [“How to Install and Configure Oracle Solaris and Oracle Solaris Cluster Software \(IPS Repositories\)”](#) on page 96.

If your nodes are already installed with the Oracle Solaris OS but do not meet Oracle Solaris Cluster installation requirements, you might need to reinstall the Oracle Solaris software. Follow the steps in this procedure to ensure subsequent successful installation of the Oracle Solaris Cluster software. See [“Planning the Oracle Solaris OS”](#) on page 12 for information about required root-disk partitioning and other Oracle Solaris 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 Oracle Solaris software. See the [“Oracle Solaris Cluster 4.2 Hardware Administration Manual”](#) 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 47 for requirements and guidelines.
- 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 17 for planning guidelines. See your Oracle Solaris system administrator documentation for information about using Oracle Solaris naming services.

1. Connect to the consoles of each node.

2. Install the Oracle Solaris OS.

Follow installation instructions in [“Installing Oracle Solaris 11.2 Systems”](#).

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

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

- a. **(Cluster nodes) Choose Manual Layout to set up the file systems.**
 - **Specify a slice that is at least 20 Mbytes in size.**
 - **Create any other file system partitions that you need, as described in [“System Disk Partitions”](#) on page 14.**

- b. (Cluster nodes) For ease of administration, set the same root password on each node.**

Note - This step is required if you plan to use the Oracle Solaris Cluster Manager GUI to administer Geographic Edition components. For more information about Oracle Solaris Cluster Manager, see [Chapter 13, “Using the Oracle Solaris Cluster GUI,”](#) in “Oracle Solaris Cluster System Administration Guide”.

- 3. Ensure that the solaris publisher is valid.**

```
# pkg publisher
PUBLISHER          TYPE    STATUS  URI
solaris            origin  online  solaris-repository
```

For information about setting the solaris publisher, see [“Adding and Updating Software in Oracle Solaris 11.2”](#).

- 4. (Cluster nodes) If you will use role-based access control (RBAC) instead of the root role to access the cluster nodes, set up an RBAC role that provides authorization for all Oracle Solaris Cluster commands.**

This series of installation procedures requires the following Oracle Solaris Cluster RBAC authorizations if the user is not the root role:

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

See [“User Rights Management”](#) in [“Securing Users and Processes in Oracle Solaris 11.2”](#) for more information about using RBAC roles. See the Oracle Solaris Cluster man pages for the RBAC authorization that each Oracle Solaris Cluster subcommand requires.

- 5. (Cluster nodes) 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.

6. Install any required Oracle Solaris OS software updates and hardware-related firmware and updates.

Include those updates for storage array support. Also download any needed firmware that is contained in the hardware updates.

For instructions on updating your software, see [Chapter 11, “Updating Your Software,”](#) in [“Oracle Solaris Cluster System Administration Guide”](#).

7. (x86 only) (Cluster nodes) Set the default boot file.

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

```
grub edit> kernel /platform/i86pc/kernel/amd64/unix -B $ZFS-BOOTFS -k
```

For more information, see [“How to Boot a System With the Kernel Debugger \(kldb\) Enabled”](#) in [“Booting and Shutting Down Oracle Solaris 11.2 Systems”](#).

8. (Cluster nodes) Update the /etc/inet/hosts 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.

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.

9. (Optional) (Cluster nodes) 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 stand-alone system. See [Chapter 3, “Administering IPMP,”](#) in [“Administering TCP/IP Networks, IPMP, and IP Tunnels in Oracle Solaris 11.2”](#) 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.

10. (Optional) (Cluster nodes) If the Oracle Solaris Cluster software is not already installed and you want to use Oracle Solaris I/O multipathing, enable multipathing on each node.



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

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

-e Enables Oracle Solaris I/O multipathing.

See [“How to Enable Multipathing”](#) in [“Managing SAN Devices and Multipathing in Oracle Solaris 11.2”](#) and the [stmsboot\(1M\)](#) man page for more information.

Next Steps If you want to use the pconsole utility, go to [“How to Install pconsole Software on an Administrative Console”](#) on page 52.

If you want to use a quorum server, go to [“How to Install and Configure Oracle Solaris Cluster Quorum Server Software”](#) on page 55.

If your cluster nodes support the mirroring of internal hard drives and you want to configure internal disk mirroring, go to [“How to Configure Internal Disk Mirroring”](#) on page 57.

SPARC: If you want to install Oracle VM Server for SPARC, go to [“How to Install Oracle VM Server for SPARC Software and Create Domains”](#) on page 58.

Otherwise, install the Oracle Solaris Cluster software on the cluster nodes.

- If you already installed the Oracle Solaris OS on the cluster nodes, go to [“How to Install Oracle Solaris Cluster Software Packages”](#) on page 59.
- If you want to use the scinstall custom Automated Installer (AI) method to install both Oracle Solaris OS and Oracle Solaris Cluster software on the cluster nodes, go to [“How to Install and Configure Oracle Solaris and Oracle Solaris Cluster Software \(IPS Repositories\)”](#) on page 96.

See Also See the [“Oracle Solaris Cluster System Administration Guide”](#) for procedures to perform dynamic reconfiguration tasks in an Oracle Solaris Cluster configuration.

▼ How to Install pconsole 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.

You cannot use this software to connect to Oracle VM Server for SPARC guest domains.

This procedure describes how to install the Parallel Console Access (pconsole) software on an administrative console. The pconsole utility is part of the Oracle Solaris terminal/pconsole package.

The `pconsole` utility creates a host terminal window for each remote host that you specify on the command line. The utility also opens a central, or master, console window that you can use to send input to all nodes at one time. For additional information, see the [pconsole\(1\)](#) man page that is installed with the `terminal/pconsole` package.

You can use any desktop machine that runs a version of the Oracle Solaris OS that is supported by Oracle Solaris Cluster 4.2 software as an administrative console.

Before You Begin Ensure that a supported version of the Oracle Solaris OS and any Oracle Solaris software updates are installed on the administrative console.

1. **Assume the root role on the administrative console.**
2. **Ensure that the data service package is available from the configured publisher and that the `solaris` and `ha-cluster` publishers are valid.**

```
# pkg list -a package-name
# pkg publisher
PUBLISHER                TYPE    STATUS  P  LOCATION
solaris                   origin  online  F  solaris-repository
ha-cluster                 origin  online  F  ha-cluster-repository
```

For information about setting the `solaris` publisher, see “[Adding, Modifying, or Removing Package Publishers](#)” in “[Adding and Updating Software in Oracle Solaris 11.2](#)”.

Tip - Use the `-nv` options whenever you install or update to see what changes will be made, such as which versions of which packages will be installed or updated and whether a new BE will be created.

If you do not get any error messages when you use the `-nv` options, run the command again without the `-n` option to actually perform the installation or update. If you do get error messages, run the command again with more `-v` options (for example, `-nvv`) or more of the package FMRI pattern to get more information to help you diagnose and fix the problem. For troubleshooting information, see [Appendix A, “Troubleshooting Package Installation and Update,”](#) in “[Adding and Updating Software in Oracle Solaris 11.2](#)”.

3. **Install the `terminal/pconsole` package.**
4. **(Optional) Install the Oracle Solaris Cluster man page packages.**

```
adminconsole# pkg install terminal/pconsole
```

```
adminconsole# pkg install pkgname ...
```

Package Name	Description
<code>ha-cluster/system/manual</code>	Oracle Solaris Cluster framework man pages

Package Name	Description
ha-cluster/system/manual/data-services	Oracle Solaris Cluster data service man pages
ha-cluster/service/quorum-server/manual	Oracle Solaris Cluster Quorum Server man pages
ha-cluster/geo/manual	Oracle Solaris Cluster Geographic Edition man pages

When you install the Oracle Solaris Cluster man page packages on the administrative console, you can view them from the administrative console before you install Oracle Solaris Cluster software on the cluster nodes or on a quorum server.

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

- a. **If you installed the `ha-cluster/system/manual/data-services` package, ensure that the `/opt/SUNWcluster/bin/` directory is in the PATH.**
- b. **If you installed any other man page package, ensure that the `/usr/cluster/bin/` directory is in the PATH.**

6. Start the pconsole utility.

Specify in the command each node that you want to connect to.

```
adminconsole# pconsole host[:port] [...] &
```

See the procedures “[Logging Into the Cluster Remotely](#)” in “[Oracle Solaris Cluster System Administration Guide](#)” and “[How to Connect Securely to Cluster Consoles](#)” in “[Oracle Solaris Cluster System Administration Guide](#)” for additional information about how to use the pconsole utility. Also see the [pconsole\(1\)](#) man page that is installed as part of the Oracle Solaris terminal/pconsole package.

Next Steps If you want to use a quorum server, go to “[How to Install and Configure Oracle Solaris Cluster Quorum Server Software](#)” on page 55.

If your cluster nodes support the mirroring of internal hard drives and you want to configure internal disk mirroring, go to “[How to Configure Internal Disk Mirroring](#)” on page 57.

SPARC: If you want to install Oracle VM Server for SPARC, go to “[How to Install Oracle VM Server for SPARC Software and Create Domains](#)” on page 58.

Otherwise, install the Oracle Solaris Cluster software on the cluster nodes.

- If you already installed the Oracle Solaris OS on the cluster nodes, go to “[How to Install Oracle Solaris Cluster Software Packages](#)” on page 59.
- If you want to use the `scinstall` custom Automated Installer (AI) method to install both Oracle Solaris OS and Oracle Solaris Cluster software on the cluster nodes, go to

[“How to Install and Configure Oracle Solaris and Oracle Solaris Cluster Software \(IPS Repositories\)”](#) on page 96

▼ How to Install and Configure Oracle Solaris Cluster Quorum Server Software

Perform this procedure to configure a host server as a quorum server.

Before You Begin Perform the following tasks:

- Ensure that the quorum server machine is connected to a public network that is accessible to the cluster nodes and on the same subnet that is used by the cluster nodes.
- Disable the spanning tree algorithm on the Ethernet switches for the ports that are connected to the cluster public network where the quorum server will run.

1. **Assume the `root` role on the machine on which you want to install the Oracle Solaris Cluster Quorum Server software.**
2. **Ensure that the data service package is available from the configured publisher and that the `solaris` and `ha-cluster` publishers are valid.**

```
# pkg list -a package-name
# pkg publisher
PUBLISHER                TYPE    STATUS  P  LOCATION
solaris                   origin  online  F  solaris-repository
ha-cluster                 origin  online  F  ha-cluster-repository
```

For information about setting the `solaris` publisher, see [“Adding, Modifying, or Removing Package Publishers”](#) in [“Adding and Updating Software in Oracle Solaris 11.2”](#).

Tip - Use the `-nv` options whenever you install or update to see what changes will be made, such as which versions of which packages will be installed or updated and whether a new BE will be created.

If you do not get any error messages when you use the `-nv` options, run the command again without the `-nv` option to actually perform the installation or update. If you do get error messages, run the command again with more `-v` options (for example, `-nvv`) or more of the package FMRI pattern to get more information to help you diagnose and fix the problem. For troubleshooting information, see [Appendix A, “Troubleshooting Package Installation and Update,”](#) in [“Adding and Updating Software in Oracle Solaris 11.2”](#).

3. **Install the Quorum Server group package.**

```
quorumserver# pkg install ha-cluster-quorum-server-full
```

4. **(Optional) Add the Oracle Solaris Cluster Quorum Server binary location to your PATH environment variable.**

```
quorumserver# PATH=$PATH:/usr/cluster/bin
```

5. **Configure the quorum server by adding the following entry to the `/etc/scqsd/scqsd.conf` file to specify configuration information about the quorum server.**

Identify the quorum server by specifying the port number and optionally the instance name.

- If you provide an instance name, that name must be unique among your quorum servers.
- If you do not provide an instance name, always refer to this quorum server by the port on which it listens.

The format for the entry is as follows:

```
/usr/cluster/lib/sc/scqsd [-d quorum-directory] [-i instance-name] -p port
```

`-d quorum-directory`

The path to the directory where the quorum server can store quorum data.

The quorum server process creates one file per cluster in this directory to store cluster-specific quorum information.

By default, the value of this option is `/var/scqsd`. This directory must be unique for each quorum server that you configure.

`-i instance-name`

A unique name that you choose for the quorum-server instance.

`-p port`

The port number on which the quorum server listens for requests from the cluster.

6. **(Optional) To serve more than one cluster but use a different port number or instance, configure an additional entry for each additional instance of the quorum server that you need.**

7. **Save and close the `/etc/scqsd/scqsd.conf` file.**

8. **Start the newly configured quorum server.**

```
quorumserver# /usr/cluster/bin/clquorumserver start quorum-server
```

`quorum-server`

Identifies the quorum server. You can use the port number on which the quorum server listens. If you provided an instance name in the configuration file, you can use that name instead.

- To start a single quorum server, provide either the instance name or the port number.

- To start all quorum servers when you have multiple quorum servers configured, use the + operand.

Troubleshooting Oracle Solaris Cluster Quorum Server software consists of the following packages:

- `ha-cluster/service/quorum-server`
- `ha-cluster/service/quorum-server/locale`
- `ha-cluster/service/quorum-server/manual`
- `ha-cluster/service/quorum-server/manual/locale`

These packages are contained in the `ha-cluster/group-package/ha-cluster-quorum-server-full` and `ha-cluster/group-package/ha-cluster-quorum-server-l10n` group packages.

The installation of these packages adds software to the `/usr/cluster` and `/etc/scqsd` directories. You cannot modify the location of the Oracle Solaris Cluster Quorum Server software.

If you receive an installation error message regarding the Oracle Solaris Cluster Quorum Server software, verify that the packages were properly installed.

Next Steps If your cluster nodes support the mirroring of internal hard drives and you want to configure internal disk mirroring, go to [“How to Configure Internal Disk Mirroring” on page 57](#).

SPARC: If you want to install Oracle VM Server for SPARC, go to [“How to Install Oracle VM Server for SPARC Software and Create Domains” on page 58](#).

Otherwise, install the Oracle Solaris Cluster software on the cluster nodes.

- If you already installed the Oracle Solaris OS on the cluster nodes, go to [“How to Install Oracle Solaris Cluster Software Packages” on page 59](#).
- If you want to use the `scinstall` custom Automated Installer (AI) method to install both Oracle Solaris OS and Oracle Solaris Cluster software on the cluster nodes, go to [“How to Install and Configure Oracle Solaris and Oracle Solaris Cluster Software \(IPS Repositories\)” on page 96](#).

▼ How to Configure Internal Disk Mirroring

Perform this procedure on each node of the global 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 “[Oracle Solaris Cluster 4.2 Hardware Administration Manual](#)”.

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

1. **Assume the root role.**
2. **Configure an internal mirror.**

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

```
-c clt0d0 clt1d0
```

Creates the mirror of primary disk to the mirror disk. Provide the name of your primary disk as the first argument and 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 SPARC: If you want to install Oracle VM Server for SPARC, go to “[How to Install Oracle VM Server for SPARC Software and Create Domains](#)” on page 58.

Otherwise, install the Oracle Solaris Cluster software on the cluster nodes.

- If you already installed the Oracle Solaris OS on the cluster nodes, go to “[How to Install Oracle Solaris Cluster Software Packages](#)” on page 59.
- If you want to use the `scinstall` custom Automated Installer (AI) method to install both Oracle Solaris OS and Oracle Solaris Cluster software on the cluster nodes, go to “[How to Install and Configure Oracle Solaris and Oracle Solaris Cluster Software \(IPS Repositories\)](#)” on page 96.

▼ SPARC: How to Install Oracle VM Server for SPARC Software and Create Domains

Perform this procedure to install Oracle VM Server for SPARC software on a physically clustered machine and to create I/O and guest domains.

Before You Begin Perform the following tasks:

- Ensure that the machine is SPARC hypervisor capable.
- Have available [“Oracle VM Server for SPARC 3.1 Administration Guide ”](#) and [“Oracle VM Server for SPARC 3.1.1.1, 3.1.1, and 3.1 Release Notes ”](#).
- Read the requirements and guidelines in [“SPARC: Guidelines for Oracle VM Server for SPARC in a Cluster”](#) on page 31.

1. **Assume the root role on the machine.**
2. **Install Oracle VM Server for SPARC software and configure domains by following the procedures in Chapter 2, “Installing and Enabling Software,” in “Oracle VM Server for SPARC 3.1 Administration Guide ”.**

Observe the following special instructions:

- If you create guest domains, adhere to the Oracle Solaris Cluster guidelines for creating guest domains in a cluster.
- Use the mode=sc option for all virtual switch devices that connect the virtual network devices that are used as the cluster interconnect.
- For shared storage, map only the full SCSI disks into the guest domains.

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 57.

Otherwise, install the Oracle Solaris Cluster software packages. Go to [“How to Install Oracle Solaris Cluster Software Packages”](#) on page 59.

▼ How to Install Oracle Solaris Cluster Software Packages

Follow this procedure to perform one or more of the following installation tasks:

- To install the Oracle Solaris Cluster framework and data service software packages on each node in the global cluster. These nodes can be physical machines or (SPARC only) Oracle VM Server for SPARC I/O domains or guest domains, or a combination of any of these types of nodes.

Note - If your physically clustered machines are configured with Oracle VM Server for SPARC, install Oracle Solaris Cluster software only in I/O domains or guest domains.

- To install data services.

Note - You cannot add or remove individual packages that are part of the `ha-cluster-minimal` group package except by complete reinstallation or uninstallation. See [“How to Unconfigure Oracle Solaris Cluster Software to Correct Installation Problems”](#) on page 227 and [“How to Uninstall Oracle Solaris Cluster Software From a Cluster Node”](#) in [“Oracle Solaris Cluster System Administration Guide”](#) for procedures to remove the cluster framework packages.

However, you can add or remove other, optional packages without removing the `ha-cluster-minimal` group package.

If you choose to install the `ha-cluster-full` package, the GUI will be installed automatically. If you choose to install another package, you can manually install the GUI later.

Before You Begin Perform the following tasks:

- Ensure that the Oracle Solaris OS is installed to support Oracle Solaris Cluster software.
If the Oracle Solaris software is already installed on the node, you must ensure that the Oracle Solaris installation meets the requirements for the Oracle Solaris Cluster software and any other software that you intend to install on the cluster. See [“How to Install Oracle Solaris Software”](#) on page 48 for more information about installing the Oracle Solaris software to meet Oracle Solaris Cluster software requirements.

- Choose which Oracle Solaris Cluster software packages to install.

The following table lists the primary group packages for the Oracle Solaris Cluster 4.2 software and the principal features that each group package contains. You must install at least the `ha-cluster-framework-minimal` group package.

Feature	<code>ha-cluster-full</code>	<code>ha-cluster-framework-full</code>	<code>ha-cluster-data-services-full</code>	<code>ha-cluster-geo-full</code>	<code>ha-cluster-minimal</code>	<code>ha-cluster-framework-minimal</code>
Framework	X	X	X	X	X	X
Agents	X		X			
Localization	X	X				
Framework man pages	X	X				
Data Service man pages	X		X			
Agent Builder	X	X				
Generic Data Service	X	X	X			

Feature	ha-cluster-full	ha-cluster-framework-full	ha-cluster-data-services-full	ha-cluster-geo-full	ha-cluster-minimal	ha-cluster-framework-minimal
Graphical User Interface	X					
Geographic Edition	X			X		

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

- **If pconsole software is installed and configured on your administrative console, use the pconsole utility to display the individual console screens.**

As the root role, use the following command to start the pconsole utility:

```
adminconsole# pconsole host[:port] [...] &
```

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

2. Restore external access to remote procedure call (RPC) communication.

During the installation of the Oracle Solaris OS, a restricted network profile is used that disables external access for certain network services. The restricted services include the RPC communication service, which is required for cluster communication.

Perform the following commands to restore external access to RPC communication.

```
# svccfg
svc:> select network/rpc/bind
svc:/network/rpc/bind> setprop config/local_only=false
svc:/network/rpc/bind> quit
# svcadm refresh network/rpc/bind:default
# svcprop network/rpc/bind:default | grep local_only
```

The output of the last command should show that the local_only property is now set to false.

3. Assume the root role on the cluster node to install.

Alternatively, if your user account is assigned the System Administrator profile, issue commands as non-root through a profile shell, or prefix the command with the pfexec command.

4. Disable Network Auto-Magic (NWAM).

NWAM activates a single network interface and disables all others. For this reason, NWAM cannot coexist with the Oracle Solaris Cluster software and you must disable it before you configure or run your cluster. To disable NWAM, you enable the defaultfixed profile.

```
# netadm enable -p ncp defaultfixed
# netadm list -p ncp defaultfixed
```

5. Set up the repository for the Oracle Solaris Cluster software packages.

- **If the cluster nodes have direct access or web proxy access to the Internet, perform the following steps.**

- a. Go to <https://pkg-register.oracle.com>.
- b. Choose Oracle Solaris Cluster software.
- c. Accept the license.
- d. Request a new certificate by choosing Oracle Solaris Cluster software and submitting a request.

The certification page is displayed with download buttons for the key and the certificate.

- e. Download the key and certificate files and install them as described in the returned certification page.
- f. Configure the ha-cluster publisher with the downloaded SSL keys and set the location of the Oracle Solaris Cluster 4.2 repository.

In the following example the repository name is `https://pkg.oracle.com/repository-location/`.

```
# pkg set-publisher \  
-k /var/pkg/ssl/Oracle_Solaris_Cluster_4.0.key.pem \  
-c /var/pkg/ssl/Oracle_Solaris_Cluster_4.0.certificate.pem \  
-O https://pkg.oracle.com/repository-location/ ha-cluster
```

```
-k /var/pkg/ssl/Oracle_Solaris_Cluster_4.0.key.pem  
    Specifies the full path to the downloaded SSL key file.
```

```
-c /var/pkg/ssl/Oracle_Solaris_Cluster_4.0.certificate.pem  
    Specifies the full path to the downloaded certificate file.
```

```
-O https://pkg.oracle.com/repository-location/  
    Specifies the URL to the Oracle Solaris Cluster 4.2 package repository.
```

For more information, see the [pkg\(1\)](#) man page.

■ **If you are using an ISO image of the software, perform the following steps.**

- a. **Download the Oracle Solaris Cluster 4.2 ISO image from Oracle Software Delivery Cloud at <https://edelivery.oracle.com/>.**

Note - A valid Oracle license is required to access Oracle Software Delivery Cloud.

Oracle Solaris Cluster software is part of the Oracle Solaris Product Pack. Follow online instructions to complete selection of the media pack and download the software.

- b. **Make the Oracle Solaris Cluster 4.2 ISO image available.**

```
# lofiadm -a path-to-iso-image
/dev/lofi/N
# mount -F hsfs /dev/lofi/N /mnt
```

-a *path-to-iso-image*

Specifies the full path and file name of the ISO image.

- c. **Set the location of the Oracle Solaris Cluster 4.2 package repository.**

```
# pkg set-publisher -g file:///mnt/repo ha-cluster
```

6. **Ensure that the data service package is available from the configured publisher and that the solaris and ha-cluster publishers are valid.**

```
# pkg list -a package-name
# pkg publisher
PUBLISHER                TYPE    STATUS  P  LOCATION
solaris                   origin  online  F  solaris-repository
ha-cluster                 origin  online  F  ha-cluster-repository
```

For information about setting the solaris publisher, see “[Adding, Modifying, or Removing Package Publishers](#)” in “[Adding and Updating Software in Oracle Solaris 11.2](#)”.

Tip - Use the -nv options whenever you install or update to see what changes will be made, such as which versions of which packages will be installed or updated and whether a new BE will be created.

If you do not get any error messages when you use the -nv options, run the command again without the -n option to actually perform the installation or update. If you do get error messages, run the command again with more -v options (for example, -nvv) or more of the

package FMRI pattern to get more information to help you diagnose and fix the problem. For troubleshooting information, see [Appendix A, “Troubleshooting Package Installation and Update,”](#) in [“Adding and Updating Software in Oracle Solaris 11.2”](#).

7. Install the Oracle Solaris Cluster 4.2 software.

```
# /usr/bin/pkg install ha-cluster-package
```

8. Verify that the package installed successfully.

```
$ pkg info -r ha-cluster-package
```

Package installation succeeded if the state is Installed.

9. Perform any necessary updates to the Oracle Solaris Cluster software.

For instructions on updating your software, see [Chapter 11, “Updating Your Software,”](#) in [“Oracle Solaris Cluster System Administration Guide”](#).

Next Steps If you want to install the Sun QFS file system software, follow the procedures for initial installation. See [“How to Install Sun QFS Software”](#) on page 65.

If you want to use the Availability Suite feature of Oracle Solaris software, install the Availability Suite software. Go to [“How to Install the Availability Suite Feature of the Oracle Solaris Software”](#) on page 64.

Otherwise, to set up the root user environment, go to [“How to Set Up the Root Environment”](#) on page 65.

▼ How to Install the Availability Suite Feature of the Oracle Solaris Software

1. Assume the root role.

2. Ensure that the solaris publishers is valid.

```
# pkg publisher
PUBLISHER          TYPE    STATUS  URI
solaris            origin  online  solaris-repository
```

For information about setting the solaris publisher, see [“Copying and Creating Package Repositories in Oracle Solaris 11.2”](#).

3. Install the IPS package for the Availability Suite feature of the Oracle Solaris software.

```
# /usr/bin/pkg install storage/avs
```


4. Configure the Availability Suite feature.

For details, see “[Initial Configuration Settings](http://docs.oracle.com/cd/E19359-01/819-6147-10/config_proc.html#pgfId-998170)” (http://docs.oracle.com/cd/E19359-01/819-6147-10/config_proc.html#pgfId-998170) in *Sun StorageTek Availability Suite 4.0 Software Installation and Configuration Guide*.

5. Verify that the packages installed successfully.

```
# pkg info group/features/storage-avs \
  storage/avs/avs-cache-management \
  storage/avs/avs-point-in-time-copy \
  storage/avs/avs-remote-mirror \
  driver/storage/sv
```

Next Steps If you want to install the Sun QFS file system software, follow the procedures for initial installation. See “[How to Install Sun QFS Software](#)” on page 65.

Otherwise, to set up the root user environment, go to “[How to Set Up the Root Environment](#)” on page 65.

▼ How to Install Sun QFS Software

Perform this procedure on each node in the global cluster.

1. Ensure that the Oracle Solaris Cluster software is installed.

See “[How to Install Oracle Solaris Cluster Software Packages](#)” on page 59.

2. Assume the root role on a cluster node.**3. Install Sun QFS file system software.**

Follow procedures for initial installation in your Sun QFS documentation.

Next Steps Set up the root user environment. Go to “[How to Set Up the Root Environment](#)” on page 65.

▼ How to Set Up the Root Environment

Note - In an Oracle Solaris 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 “[About the User’s Work Environment](#)” in “[Managing User Accounts and User Environments in Oracle Solaris 11.2](#)” for more information.

Perform this procedure on each node in the global cluster.

1. **Assume the root role on a cluster node.**
2. **Add `/usr/cluster/bin/` and `/usr/sbin/` to the PATH.**

Note - Always make `/usr/cluster/bin` the *first* entry in the PATH. This placement ensures that Oracle Solaris Cluster commands take precedence over any other binaries that have the same name, thus avoiding unexpected behavior.

See your Oracle 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 If you want to use the IP Filter feature of Oracle Solaris, go to [“How to Configure IP Filter” on page 66](#).

Otherwise, configure Oracle Solaris Cluster software on the cluster nodes. Go to [“Establishing a New Global Cluster or New Global-Cluster Node” on page 71](#).

▼ How to Configure IP Filter

Perform this procedure to configure the IP Filter feature of Oracle Solaris software on the global cluster.

Note - Only use IP Filter with failover data services. The use of IP Filter with scalable data services is not supported.

For more information about the IP Filter feature, see [Chapter 4, “About IP Filter in Oracle Solaris,”](#) in [“Securing the Network in Oracle Solaris 11.2”](#).

Before You Begin Read the guidelines and restrictions to follow when you configure IP Filter in a cluster. See the “IP Filter” bullet item in [“Oracle Solaris OS Feature Requirements and Restrictions” on page 13](#).

1. **Assume the root role.**
2. **Add filter rules to the `/etc/ipf/ipf.conf` file on all affected nodes.**

Observe the following guidelines and requirements when you add filter rules to Oracle Solaris Cluster nodes.

- In the `ipf.conf` file on each node, add rules to explicitly allow cluster interconnect traffic to pass unfiltered. Rules that are not interface specific are applied to all interfaces, including cluster interconnects. Ensure that traffic on these interfaces is not blocked mistakenly. If interconnect traffic is blocked, the IP Filter configuration interferes with cluster handshakes and infrastructure operations.

For example, suppose the following rules are currently used:

```
# Default block TCP/UDP unless some later rule overrides
block return-rst in proto tcp/udp from any to any

# Default block ping unless some later rule overrides
block return-rst in proto icmp all
```

To unblock cluster interconnect traffic, add the following rules. The subnets used are for example only. Derive the subnets to use by using the `ifconfig show-addr | grep interface` command.

```
# Unblock cluster traffic on 172.16.0.128/25 subnet (physical interconnect)
pass in quick proto tcp/udp from 172.16.0.128/25 to any
pass out quick proto tcp/udp from 172.16.0.128/25 to any

# Unblock cluster traffic on 172.16.1.0/25 subnet (physical interconnect)
pass in quick proto tcp/udp from 172.16.1.0/25 to any
pass out quick proto tcp/udp from 172.16.1.0/25 to any

# Unblock cluster traffic on 172.16.4.0/23 (clprivnet0 subnet)
pass in quick proto tcp/udp from 172.16.4.0/23 to any
pass out quick proto tcp/udp from 172.16.4.0/23 to any
```

- You can specify either the adapter name or the IP address for a cluster private network. For example, the following rule specifies a cluster private network by its adapter's name:

```
# Allow all traffic on cluster private networks.
pass in quick on net1 all
...
```

- Oracle Solaris Cluster software fails over network addresses from node to node. No special procedure or code is needed at the time of failover.
- All filtering rules that reference IP addresses of logical hostname and shared address resources must be identical on all cluster nodes.
- Rules on a standby node will reference a nonexistent IP address. This rule is still part of the IP filter's active rule set and will become effective when the node receives the address after a failover.
- All filtering rules must be the same for all NICs in the same IPMP group. In other words, if a rule is interface-specific, the same rule must also exist for all other interfaces in the same IPMP group.

For more information about IP Filter rules, see the [ipf\(4\)](#) man page.

3. Enable the `ipfilter` SMF service.

```
phys-schost# svcadm enable /network/ipfilter:default
```

Next Steps Configure Oracle Solaris Cluster software on the cluster nodes. Go to [“Establishing a New Global Cluster or New Global-Cluster Node”](#) on page 71.

Establishing the Global Cluster

This chapter provides information about how to establish a new global cluster or a new global-cluster node.

Note - To create a zone cluster, see [“Creating and Configuring a Zone Cluster” on page 188](#). You must establish a global cluster before you can create a zone cluster.

This chapter contains the following information:

- [“Overview of Establishing a New Cluster or Cluster Node” on page 69](#)
- [“Establishing a New Global Cluster or New Global-Cluster Node” on page 71](#)

Overview of Establishing a New Cluster or Cluster Node

The following task maps list the tasks to perform for either a new global cluster or a node added to an existing global cluster. Complete the procedures in the order that is indicated.

- [Table 3-1](#)
- [Table 3-2](#)

TABLE 3-1 Task Map: Establish a New Global Cluster

Method	Instructions
Use one of the following methods to establish a new global cluster:	
■ Use the <code>scinstall</code> utility to establish the cluster.	“Configuring Oracle Solaris Cluster Software on All Nodes (scinstall)” on page 71
■ Use an XML configuration file to establish the cluster.	“How to Configure Oracle Solaris Cluster Software on All Nodes (XML)” on page 81
■ Set up an Automated Installer (AI) install server. Then use the <code>scinstall</code> AI option to install the software on each node and establish the cluster.	“Establishing a New Oracle Solaris Cluster With the Automated Installer” on page 89
Assign quorum votes and remove the cluster from installation mode, if this operation was not already performed.	“How to Configure Quorum Devices” on page 137
Validate the quorum configuration.	“How to Verify the Quorum Configuration and Installation Mode” on page 142

Method	Instructions
(Optional) Change a node's private hostname.	"How to Change Private Hostnames" on page 143
Create or modify the NTP configuration file, if not already configured.	"Configuring Network Time Protocol (NTP)" on page 144
If using a volume manager, install the volume management software.	Chapter 4, "Configuring Solaris Volume Manager Software"
Create cluster file systems or highly available local file systems as needed.	Chapter 5, "Creating a Cluster File System" or "Enabling Highly Available Local File Systems" in "Oracle Solaris Cluster Data Services Planning and Administration Guide "
Install third-party applications, register resource types, set up resource groups, and configure data services.	"Oracle Solaris Cluster Data Services Planning and Administration Guide " Documentation that is supplied with the application software
Validate the cluster.	"How to Validate the Cluster" on page 148
Take a baseline recording of the finished cluster configuration.	"How to Record Diagnostic Data of the Cluster Configuration" on page 148

TABLE 3-2 Task Map: Add a Node to an Existing Global Cluster

Method	Instructions
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.	"How to Prepare the Cluster for Additional Global-Cluster Nodes" on page 114
Reconfigure the cluster interconnect and the private network address range as needed to accommodate the added node.	"How to Change the Private Network Configuration When Adding Nodes or Private Networks" on page 116
Use one of the following methods to add a node to an existing global cluster:	
<ul style="list-style-type: none"> ■ Configure Oracle Solaris Cluster software on the new node by using the <code>scinstall</code> utility. 	"Configuring Oracle Solaris Cluster Software on Additional Global-Cluster Nodes (<code>scinstall</code>)" on page 122
<ul style="list-style-type: none"> ■ Configure Oracle Solaris Cluster software on the new node by using an XML configuration file. 	"How to Configure Oracle Solaris Cluster Software on Additional Global-Cluster Nodes (XML File)" on page 129
Update the quorum configuration information.	"How to Update Quorum Devices After Adding a Node to a Global Cluster" on page 134
Validate the quorum configuration.	"How to Verify the Quorum Configuration and Installation Mode" on page 142
(Optional) Change a node's private hostname.	"How to Change Private Hostnames" on page 143
Modify the NTP configuration.	"Configuring Network Time Protocol (NTP)" on page 144
If using a volume manager, install the volume management software.	Chapter 4, "Configuring Solaris Volume Manager Software"
Create cluster file systems or highly available local file systems as needed.	Chapter 5, "Creating a Cluster File System" or "Enabling Highly Available Local File Systems" in "Oracle Solaris Cluster Data Services Planning and Administration Guide "
Install third-party applications, register resource types, set up resource groups, and configure data services.	"Oracle Solaris Cluster Data Services Planning and Administration Guide " Documentation that is supplied with the application software
Validate the cluster.	"How to Validate the Cluster" on page 148
Take a baseline recording of the finished cluster configuration.	"How to Record Diagnostic Data of the Cluster Configuration" on page 148

Establishing a New Global Cluster or New Global-Cluster Node

This section describes how to establish a new global cluster or add a node to an existing cluster. Global-cluster nodes can be physical machines, Oracle VM Server for SPARC I/O domains, or Oracle VM Server for SPARC guest domains. A cluster can consist of a combination of any of these node types. Before you start to perform these tasks, ensure that you installed software packages for the Oracle Solaris OS, Oracle Solaris Cluster framework, and other products as described in [“Installing the Software” on page 46](#).

This section contains the following information and procedures:

- [“Configuring Oracle Solaris Cluster Software on All Nodes \(scinstall\)” on page 71](#)
- [“How to Configure Oracle Solaris Cluster Software on All Nodes \(XML\)” on page 81](#)
- [“Establishing a New Oracle Solaris Cluster With the Automated Installer” on page 89](#)
- [“How to Prepare the Cluster for Additional Global-Cluster Nodes” on page 114](#)
- [“How to Change the Private Network Configuration When Adding Nodes or Private Networks” on page 116](#)
- [“Configuring Oracle Solaris Cluster Software on Additional Global-Cluster Nodes \(scinstall\)” on page 122](#)
- [“How to Configure Oracle Solaris Cluster Software on Additional Global-Cluster Nodes \(XML File\)” on page 129](#)
- [“How to Update Quorum Devices After Adding a Node to a Global Cluster” on page 134](#)
- [“How to Configure Quorum Devices” on page 137](#)
- [“How to Verify the Quorum Configuration and Installation Mode” on page 142](#)
- [“How to Change Private Hostnames” on page 143](#)
- [“Configuring Network Time Protocol \(NTP\)” on page 144](#)
- [“How to Validate the Cluster” on page 148](#)
- [“How to Record Diagnostic Data of the Cluster Configuration” on page 148](#)

Configuring Oracle Solaris Cluster Software on All Nodes (scinstall)

The `scinstall` utility runs in two modes of installation, Typical or Custom. For the Typical installation of Oracle Solaris Cluster software, `scinstall` automatically specifies the following configuration defaults.

Private-network address

```
172.16.0.0
```

Private-network netmask

255.255.240.0

Cluster-transport adapters

Exactly two adapters

Cluster-transport switches

switch1 and switch2

Global fencing

Enabled

Installation security (DES)

Limited

Complete one of the following cluster configuration worksheets to plan your Typical mode or Custom mode installation:

- 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?	
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?	First: Second:
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 as a quorum device.)</i>	Yes No
Check	Do you want to interrupt cluster creation for cluster check 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?	

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>	
Authenticating Requests to Add Nodes <i>(multiple-node cluster only)</i>	Do you need to use DES authentication?	No Yes
Minimum Number of Private Networks <i>(multiple-node cluster only)</i>	Should this cluster use at least two private networks?	Yes No
Point-to-Point Cables <i>(multiple-node cluster only)</i>	If this is a two-node cluster, does this cluster use switches?	Yes No
Cluster Switches <i>(multiple-node cluster only)</i>	Transport switch name: Defaults: switch1 and switch2	First: Second:
Cluster Transport Adapters and Cables <i>(multiple-node cluster only)</i>	Node name <i>(the node from which you run scinstall)</i> :	
	Transport adapter name:	First: Second:
	Will this be a dedicated cluster transport adapter? <i>(Answer No if using tagged VLAN adapters.)</i>	First: Yes No Second: Yes No
	If no, what is the VLAN ID for this adapter?	First: Second:
	Where does each transport adapter connect to <i>(a switch or another adapter)</i> ? Switch defaults: switch1 and switch2	First: Second:
	If a transport switch, do you want to use the default port name?	First: Yes No Second: Yes No
	If no, what is the name of the port that you want to use?	First: Second:
	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 <i>(multiple-node cluster only)</i>	Node name:	
	Transport adapter name:	First: Second:
	Will this be a dedicated cluster transport adapter? <i>(Answer No if using tagged VLAN adapters.)</i>	First: Yes No Second: Yes No

Component	Description/Example	Answer
	If no, what is the VLAN ID for this adapter?	First: Second:
	Where does each transport adapter connect to (<i>a switch or another adapter</i>)? Defaults: switch1 and switch2	First: Second:
	If a transport switch, do you want to use the default port name?	First: Yes No Second: Yes No
	If no, what is the name of the port that you want to use?	First: Second:
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 (255.255.240.0)?	Yes No
	If no, what are the maximum numbers of nodes, private networks, and zone clusters that you expect to configure in the cluster?	____ nodes ____ networks ____ zone clusters
	Which netmask do you want to use? (<i>Choose from the values calculated by scinstall or supply your own.</i>)	____.____.____.____
Shared-IP or Exclusive-IP zone clusters	How many exclusive-IP or shared-IP zone clusters do you want to create on the physical cluster? (<i>The combined number of exclusive-IP and shared-IP zone clusters must be less than or equal to the maximum number of zone clusters that can be created on the physical cluster. The maximum number of exclusive-IP zone clusters is three.</i>)	____ exclusive-IP zone clusters ____ shared-IP zone clusters
Global Fencing	Do you want to disable global fencing? (<i>Answer No unless the shared storage does not support SCSI reservations or unless you want systems that are outside the cluster to access the shared storage.</i>)	First: Yes No Second: Yes No
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 as a quorum device.</i>)	First: Yes No Second: Yes No
Check <i>(multiple-node cluster only)</i>	Do you want to interrupt cluster creation for cluster check errors?	Yes No
<i>(single-node cluster only)</i>	Do you want to run the cluster check utility to validate the cluster?	Yes No
Automatic Reboot <i>(single-node cluster only)</i>	Do you want scinstall to automatically reboot the node after installation?	Yes No

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

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

This procedure uses the interactive form of the `scinstall` command. For information about how to use the noninteractive forms of the `scinstall` command, such as when developing installation scripts, see the [`scinstall\(1M\)` man page](#).

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.

Before You Begin Perform the following tasks:

- Ensure that the Oracle Solaris OS is installed to support the Oracle Solaris Cluster software. If the Oracle Solaris software is already installed on the node, you must ensure that the Oracle Solaris installation meets the requirements for the Oracle Solaris Cluster software and any other software that you intend to install on the cluster. See [“How to Install Oracle Solaris Software” on page 48](#) for more information about installing the Oracle Solaris software to meet Oracle Solaris Cluster software requirements.
 - Ensure that NWAM is disabled. See [“How to Install Oracle Solaris Cluster Software Packages” on page 59](#) for instructions.
 - SPARC: If you are configuring Oracle VM Server for SPARC I/O domains or guest domains as cluster nodes, ensure that the Oracle VM Server for SPARC software is installed on each physical machine and that the domains meet Oracle Solaris Cluster requirements. See [“How to Install Oracle VM Server for SPARC Software and Create Domains” on page 58](#).
 - Ensure that Oracle Solaris Cluster software packages and updates are installed on each node. See [“How to Install Oracle Solaris Cluster Software Packages” on page 59](#).
 - Ensure that any adapters that you want to use as tagged VLAN adapters are configured and that you have their VLAN IDs.
 - Have available your completed Typical Mode or Custom Mode installation worksheet. See [“Configuring Oracle Solaris Cluster Software on All Nodes \(`scinstall`\)” on page 71](#).
1. **If you are using switches in the private interconnect of your new cluster, ensure that Neighbor Discovery Protocol (NDP) is disabled.**
Follow the procedures in the documentation for your switches to determine whether NDP is enabled and to disable NDP.

During cluster configuration, the software checks that there is no traffic on the private interconnect. If NDP sends any packages to a private adapter when the private interconnect is being checked for traffic, the software will assume that the interconnect is not private and cluster configuration will be interrupted. NDP must therefore be disabled during cluster creation.

After the cluster is established, you can re-enable NDP on the private-interconnect switches if you want to use that feature.

2. On each node to configure in a cluster, assume the root role.

Alternatively, if your user account is assigned the System Administrator profile, issue commands as non-root through a profile shell, or prefix the command with the `pfexec` command.

3. Ensure that TCP wrappers for RPC are disabled on all nodes of the cluster.

The Oracle Solaris TCP wrappers for RPC feature prevents internode communication that is necessary for cluster configuration.

a. On each node, display the status of TCP wrappers for RPC.

TCP wrappers are enabled if `config/enable_tcpwrappers` is set to `true`, as shown in the following example command output.

```
# svccfg -s rpc/bind listprop config/enable_tcpwrappers
config/enable_tcpwrappers boolean true
```

b. If TCP wrappers for RPC are enabled on a node, disable TCP wrappers and refresh the RPC bind service.

```
# svccfg -s rpc/bind setprop config/enable_tcpwrappers = false
# svcadm refresh rpc/bind
# svcadm restart rpc/bind
```

4. Prepare public-network interfaces.

a. Create static IP addresses for each public-network interface.

```
# ipadm create-ip interface
# ipadm create-addr -T static -a local=address/prefix-length addrobj
```

For more information, see [“How to Configure an IPv4 Interface”](#) in [“Configuring and Administering Network Components in Oracle Solaris 11.2”](#).

b. (Optional) Create IPMP groups for public-network interfaces.

During initial cluster configuration, unless non-link-local IPv6 public network interfaces exist in the cluster, IPMP groups are automatically created based on matching subnets.

These groups use transitive probes for interface monitoring and no test addresses are required.

If these automatically created IPMP groups would not meet your needs, or if IPMP groups would not be created because your configuration includes one or more non-link-local IPv6 public network interfaces, do one of the following:

- **Create the IPMP groups you need before you establish the cluster.**
- **After the cluster is established, use the `ipadm` command to edit the IPMP groups.**

For more information, see [“Configuring IPMP Groups”](#) in [“Administering TCP/IP Networks, IPMP, and IP Tunnels in Oracle Solaris 11.2”](#).

5. From one cluster node, start the `scinstall` utility.

```
phys-schost# scinstall
```

6. Type the option number 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) Print release information for this cluster node

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

```
Option: 1
```

The New Cluster and Cluster Node Menu is displayed.

**7. Type the option number for Create a New Cluster and press the Return key.**

The Typical or Custom Mode menu is displayed.

**8. Type the option number 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.

**9. 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. Oracle Solaris Cluster installation output is logged in a `/var/cluster/logs/install/scinstall.log.N` file.

**10. Verify on each node that multiuser services for the Service Management Facility (SMF) are online.**

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

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

**11. From one node, 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. Take the cluster out of installmode.**

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

**13. (Optional) Enable the automatic node reboot feature.**

This feature automatically reboots a node if all monitored shared-disk paths fail, provided that at least one of the disks is accessible from a different node in the cluster.

---

**Note** - At initial configuration time, disk-path monitoring is enabled by default for all discovered devices.

---

**a. Enable automatic reboot.**

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

-p

Specifies the property to set

```
reboot_on_path_failure=enable
```

Enables automatic node reboot if failure of all monitored shared-disk paths occurs.

**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 plan to enable RPC use of TCP wrappers, add all clprivnet0 IP addresses to the /etc/hosts.allow file on each cluster node.**

Without this addition to the /etc/hosts.allow file, TCP wrappers prevent internode communication over RPC for cluster administration utilities.

**a. On each node, display the IP addresses for all clprivnet0 devices on the node.**

```
/usr/sbin/ipadm show-addr
ADDROBJ TYPE STATE ADDR
clprivnet0/N static ok ip-address/netmask-length
...
```

**b. On each cluster node, add to the /etc/hosts.allow file the IP addresses of all clprivnet0 devices in the cluster.****15. If you intend to use the HA for NFS data service (HA for NFS) on a highly available local file system, exclude from the automounter map all shares that are part of the highly available local file system that is exported by HA for NFS.**

See [“Administrative Tasks Involving Maps”](#) in [“Managing Network File Systems in Oracle Solaris 11.2”](#) for more information about modifying the automounter map.

**Example 3-1** Configuring Oracle Solaris 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 utility in Typical Mode. The other cluster node is phys-schost-2. The adapter names are net2 and net3. The automatic selection of a quorum device is enabled.

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

```
Configuring global device using lofi on pred1: done
Starting discovery of the cluster transport configuration.
```

The following connections were discovered:

```
phys-schost-1:net2 switch1 phys-schost-2:net2
phys-schost-1:net3 switch2 phys-schost-2:net3
```

Completed discovery of the cluster transport configuration.

Started cluster check on "phys-schost-1".  
Started cluster check on "phys-schost-2".

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

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

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

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

**Troubleshooting** **Unsuccessful configuration** – If one or more nodes cannot join the cluster, or if the wrong configuration information was specified, first attempt to perform this procedure again. If that does not correct the problem, perform the procedure [“How to Unconfigure Oracle Solaris Cluster Software to Correct Installation Problems” on page 227](#) on each misconfigured node to remove it from the cluster configuration. You do not need to uninstall the Oracle Solaris Cluster software packages. Then perform this procedure again.

- Next Steps**
- If you installed a single-node cluster, cluster establishment is complete. Go to [“Creating Cluster File Systems” on page 165](#) to install volume management software and configure the cluster.
  - If you installed a multiple-node cluster and chose automatic quorum configuration, post installation setup is complete. Go to [“How to Verify the Quorum Configuration and Installation Mode” on page 142](#).
  - If you installed a multiple-node cluster and declined automatic quorum configuration, perform post installation setup. Go to [“How to Configure Quorum Devices” on page 137](#).

If you intend to configure any quorum devices in your cluster, go to [“How to Configure Quorum Devices” on page 137](#).

Otherwise, go to [“How to Verify the Quorum Configuration and Installation Mode” on page 142](#).



## ▼ How to Configure Oracle Solaris Cluster Software on All Nodes (XML)

Perform this procedure to configure a new global cluster by using an XML cluster configuration file. The new cluster can be a duplication of an existing cluster that runs Oracle Solaris Cluster 4.1 software.

This procedure configures the following cluster components:

- Cluster name
- Cluster node membership
- Cluster interconnect

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

- Ensure that the Oracle Solaris OS is installed to support the Oracle Solaris Cluster software. If the Oracle Solaris software is already installed on the node, you must ensure that the Oracle Solaris installation meets the requirements for the Oracle Solaris Cluster software and any other software that you intend to install on the cluster. See [“How to Install Oracle Solaris Software” on page 48](#) for more information about installing the Oracle Solaris software to meet Oracle Solaris Cluster software requirements.
- Ensure that NWAM is disabled. See [“How to Install Oracle Solaris Cluster Software Packages” on page 59](#) for instructions.
- SPARC: If you are configuring Oracle VM Server for SPARC I/O domains or guest domains as cluster nodes, ensure that the Oracle VM Server for SPARC software is installed on each physical machine and that the domains meet Oracle Solaris Cluster requirements. See [“How to Install Oracle VM Server for SPARC Software and Create Domains” on page 58](#).
- Ensure that any adapters that you want to use as tagged VLAN adapters are configured and that you have their VLAN IDs.
- Ensure that Oracle Solaris Cluster 4.2 software and updates are installed on each node that you will configure. See [“How to Install Oracle Solaris Cluster Software Packages” on page 59](#).

1. **Ensure that the Oracle Solaris Cluster 4.2 software is not yet configured on each potential cluster node.**
  - a. **Assume the root role on a potential node that you want to configure in the new cluster.**
  - b. **Determine whether the Oracle Solaris Cluster 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 a cluster: Operation not applicable
```

This message indicates that the Oracle Solaris 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 the Oracle Solaris Cluster software is already configured on the node.

If the cluster is running an older version of Oracle Solaris Cluster software and you want to install Oracle Solaris Cluster 4.2 software, instead perform upgrade procedures in [“Oracle Solaris Cluster Upgrade Guide”](#).

- c. **Repeat Step a and Step b on each remaining potential node that you want to configure in the new cluster.**

If the Oracle Solaris Cluster software is not yet configured on any of the potential cluster nodes, proceed to Step 2.

2. **Ensure that TCP wrappers for RPC are disabled on all nodes of the cluster.**

The Oracle Solaris TCP wrappers for RPC feature prevents internode communication that is necessary for cluster configuration.

- a. **On each node, display the status of TCP wrappers for RPC.**

TCP wrappers are enabled if `config/enable_tcpwrappers` is set to `true`, as shown in the following example command output.

```
svccfg -s rpc/bind listprop config/enable_tcpwrappers
config/enable_tcpwrappers boolean true
```

- b. **If TCP wrappers for RPC are enabled on a node, disable TCP wrappers and refresh the RPC bind service.**

```
svccfg -s rpc/bind setprop config/enable_tcpwrappers = false
svcadm refresh rpc/bind
svcadm restart rpc/bind
```

3. **If you are using switches in the private interconnect of your new cluster, ensure that Neighbor Discovery Protocol (NDP) is disabled.**

Follow the procedures in the documentation for your switches to determine whether NDP is enabled and to disable NDP.

During cluster configuration, the software checks that there is no traffic on the private interconnect. If NDP sends any packages to a private adapter when the private interconnect

is being checked for traffic, the software will assume that the interconnect is not private and cluster configuration will be interrupted. NDP must therefore be disabled during cluster creation.

After the cluster is established, you can re-enable NDP on the private-interconnect switches if you want to use that feature.

**4. If you are duplicating an existing cluster than runs the Oracle Solaris Cluster 4.2 software, use a node in that cluster to create a cluster configuration XML file.**

**a. Assume the `root` role 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.

**5. Assume the `root` role on the potential node from which you will configure the new cluster.**

**6. Modify or create the cluster configuration XML file as needed.**

Include or modify the values of the XML elements to reflect the cluster configuration that you want to create.

- 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.

- To establish a cluster, the following components must have valid values in the cluster configuration XML file:
  - Cluster name
  - Cluster nodes
  - Cluster transport
- 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.

**7. Validate the cluster configuration XML file.**

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

See the [xmllint\(1\)](#) man page for more information.

**8. 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.

**9. Verify on each node that multiuser services for the Service Management Facility (SMF) are online.**

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

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

**10. From one node, 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. Perform any necessary updates to the Oracle Solaris Cluster software.**

For instructions on updating your software, see [Chapter 11, “Updating Your Software,”](#) in [“Oracle Solaris Cluster System Administration Guide”](#).

**12. If you plan to enable RPC use of TCP wrappers, add all clprivnet0 IP addresses to the /etc/hosts.allow file on each cluster node.**

Without this addition to the /etc/hosts.allow file, TCP wrappers prevent internode communication over RPC for cluster administration utilities.

**a. On each node, display the IP addresses for all clprivnet0 devices on the node.**

```
/usr/sbin/ipadm show-addr
ADDROBJ TYPE STATE ADDR
clprivnet0/N static ok ip-address/netmask-length
...
```

**b. On each cluster node, add to the /etc/hosts.allow file the IP addresses of all clprivnet0 devices in the cluster.**

**13. If you intend to use the HA for NFS data service (HA for NFS) on a highly available local file system, exclude from the automounter map all shares that are part of the highly available local file system that is exported by HA for NFS.**

See [“Administrative Tasks Involving Maps”](#) in [“Managing Network File Systems in Oracle Solaris 11.2”](#) for more information about modifying the automounter map.

**14. 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 137.

**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 [“How to Install and Configure Oracle Solaris Cluster Quorum Server Software”](#) on page 55.

**b. If you are using a 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 [“Oracle Solaris Cluster With Network-Attached Storage Device Manual”](#).

- 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 device-name
```

*device-name*

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

- 15. **Remove the cluster from installation mode.**

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

- 16. **Close access to the cluster configuration by machines that are not configured cluster members.**

```
phys-schost# claccess deny-all
```

- 17. **(Optional) Enable automatic node reboot if all monitored shared-disk paths fail.**

---

**Note** - At initial configuration time, disk-path monitoring is enabled by default for all discovered devices.

---

- a. **Enable automatic reboot.**

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

*-p*

Specifies the property to set

```
reboot_on_path_failure=enable
```

Enables automatic node reboot if failure of all monitored shared-disk paths occurs.

- 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 Oracle Solaris 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 11.1 OS. 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 a cluster: 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-N# 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

```

### Configuring Additional Components

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 describes a list of the cluster components that you can create from a cluster configuration XML file after the cluster is established. The list includes the man page for the command that you use to duplicate the component:

- Device groups: Solaris Volume Manager: [cldevicegroup\(1CL\)](#)  
For Solaris Volume Manager, first create the disk sets that you specify in the cluster configuration XML file.
- Resource Group Manager components
  - Resources: [clresource\(1CL\)](#)
  - Shared address resources: [clressharedaddress\(1CL\)](#)
  - Logical hostname resources: [clreslogicalhostname\(1CL\)](#)
  - Resource types: [clresourcetype\(1CL\)](#)
  - Resource groups: [clresourcegroup\(1CL\)](#)

You can use the `-a` option of the `clresource`, `clressharedaddress`, or `clreslogicalhostname` 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.

- NAS devices: [clnasdevice\(1CL\)](#)  
You must first set up the NAS device as described in the device's documentation.
- SNMP hosts: [clsnmphost\(1CL\)](#)  
The `clsnmphost create -i` command requires that you specify a user password file with the `-f` option.
- SNMP users: [clsnmpuser\(1CL\)](#)
- Thresholds for monitoring system resources on cluster objects:  
[cltelemetryattribute\(1CL\)](#)

**Troubleshooting** **Unsuccessful configuration** – If one or more nodes cannot join the cluster, or if the wrong configuration information was specified, first attempt to perform this procedure again. If that does not correct the problem, perform the procedure [“How to Unconfigure Oracle Solaris Cluster Software to Correct Installation Problems” on page 227](#) on each misconfigured node to remove it from the cluster configuration. You do not need to uninstall the Oracle Solaris Cluster software packages. Then perform this procedure again.

**Next Steps** Go to [“How to Verify the Quorum Configuration and Installation Mode” on page 142](#).



## Establishing a New Oracle Solaris Cluster With the Automated Installer

The interactive `scinstall` utility runs in two modes of installation, Typical or Custom. For the Typical installation of Oracle Solaris Cluster software, `scinstall` automatically specifies the following configuration defaults.

Private-network address

172.16.0.0

Private-network netmask

255.255.240.0

Cluster-transport adapters

Exactly two adapters

Cluster-transport switches

switch1 and switch2

Global fencing

Enabled

Installation security (DES)

Limited

You can install and configure a new cluster by installing the Oracle Solaris and Oracle Solaris Cluster software packages from IPS repositories, or from an Oracle Solaris Unified Archive that is created on an existing cluster.

Besides forming a new cluster, you can also use the AI and the United Archives to replicate a cluster from the archive, and restore existing cluster nodes. You can also use the `clzonecluster` command to install a new zone cluster from the Unified Archives. For more information, see [“How to Install and Configure Oracle Solaris and Oracle Solaris Cluster Software \(Unified Archives\)”](#) on page 102, [“How to Replicate a Cluster from the Unified Archives”](#) on page 108, [“How to Restore a Node from the Unified Archive”](#) in [“Oracle Solaris Cluster System Administration Guide”](#), and [“How to Install a Zone Cluster from the Unified Archive”](#) in [“Oracle Solaris Cluster System Administration Guide”](#).

These nodes can be physical machines or (SPARC only) Oracle VM Server for SPARC I/O domains or guest domains, or a combination of any of these types of nodes.

AI uses a minimal boot image to boot the client. When you install the Oracle Solaris and Oracle Solaris Cluster software packages from IPS repositories, you must provide a source for the installation to obtain the boot image. The boot image content is published in the `install-`

image/solaris-auto-install package. The downloaded boot image ISO file also contains the boot image. You can either specify the repository from which the package can be retrieved, or you can specify the location of the downloaded boot image ISO file.

- To obtain the boot image from the repository, you will need to specify the publisher, the repository URL, and the architecture of the cluster nodes. If the repository uses HTTPS, you will also need to specify the SSL certificate and the private key, and provide the location of the files. You can request and download the key and certificate from the `http://pkg-register.oracle.com` site.
- To use the downloaded boot image ISO file, you must save it in a directory that can be accessed from the AI install server. The AI boot image must be the same version as the Oracle Solaris software release that you plan to install on the cluster nodes. Also, the boot image file must have the same architecture as that of the cluster nodes.

If you want to establish a new cluster from Oracle Unified Archives, either to install and configure a new cluster or to replicate a cluster from the archives, you do not need to provide the minimal boot image. The Unified Archive contains an image you can use. You do need to provide the path to access the Unified Archive.

When you install and configure a new cluster from either IPS repositories or Unified Archives, complete one of the following cluster configuration worksheets to plan your Typical mode or Custom mode installation:

- **Typical Mode Worksheet** – If you will use Typical mode and accept all defaults, complete the following worksheet.

| Component                                    | Description/Example                                                                                                                                                                                                                                                                                               | Answer |
|----------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|--------|
| Custom Automated Installer Boot Image Source | If you plan to use a downloaded AI ISO image file, you will need the following information:<br><br>What is the full path name of the Automated Installer boot image ISO file?                                                                                                                                     |        |
|                                              | If you plan to use a repository to get the AI boot image, you will need the following information:<br><br>What is the publisher for the boot image <code>install-image/solaris-auto-install</code> package?<br><br>What is the repository of the publisher?<br><br>What is the architecture of the cluster nodes? |        |
|                                              | For repositories that use HTTPS:<br><br>What is the full path of the certificate file for the repository?<br><br>What is the full path of the private key file for the repository?<br><br>You can request and download the key and certificate from the <code>http://pkg-register.oracle.com</code> site.         |        |
| Unified Archives                             | If you plan to use the Unified Archives to install, you will need the following information:                                                                                                                                                                                                                      |        |

| Component                                                                          | Description/Example                                                                                                                                                                                                                                                                                                                       | Answer                              |
|------------------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------------------------|
|                                                                                    | What is the location of the unified archive?                                                                                                                                                                                                                                                                                              |                                     |
| Custom Automated Installer<br>User root Password                                   | What is the password for the root account of the cluster nodes?                                                                                                                                                                                                                                                                           |                                     |
| Custom Automated Installer<br>Repositories (when not using<br>the Unified Archive) | What is the repository of publisher solaris?                                                                                                                                                                                                                                                                                              |                                     |
|                                                                                    | What is the repository of publisher ha-cluster?                                                                                                                                                                                                                                                                                           |                                     |
|                                                                                    | For repositories that use HTTPS:<br><br>What is the full path of the certificate file for the repository?<br><br>What is the full path of the private key file for the repository?<br><br>You can request and download the key and certificate from the <a href="http://pkg-register.oracle.com">http://pkg-register.oracle.com</a> site. |                                     |
|                                                                                    | Select the Oracle Solaris Cluster components that you want to install. <i>(Select one or more group packages to install.)</i>                                                                                                                                                                                                             |                                     |
|                                                                                    | Do you want to select any individual components that are contained in these group packages?                                                                                                                                                                                                                                               | Yes   No                            |
| 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>                                                                                                                                                                                    |                                     |
|                                                                                    | Confirm that the auto-discovered MAC address for each node is correct.                                                                                                                                                                                                                                                                    |                                     |
| Cluster Transport Adapters and<br>Cables                                           | First node name:                                                                                                                                                                                                                                                                                                                          |                                     |
|                                                                                    | Transport adapter names:                                                                                                                                                                                                                                                                                                                  | First:<br>Second:                   |
| VLAN adapters only                                                                 | Will this be a dedicated cluster transport adapter? <i>(Answer No if using tagged VLAN adapters.)</i>                                                                                                                                                                                                                                     | First: Yes   No<br>Second: Yes   No |
|                                                                                    | If no, what is the VLAN ID for this adapter?                                                                                                                                                                                                                                                                                              | First:<br>Second:                   |
| Specify for each additional<br>node                                                | Node name:                                                                                                                                                                                                                                                                                                                                |                                     |
|                                                                                    | Transport adapter names:                                                                                                                                                                                                                                                                                                                  | First:<br>Second:                   |
| Quorum Configuration<br><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 as a quorum device.)</i>                                                                                                                                | First: Yes   No<br>Second: 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   |
|------------------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|----------|
| Custom Automated Installer<br>Boot Image ISO source                                | If you plan to use a downloaded AI ISO image file, you will need the following information:<br><br>What is the full path name of the Automated Installer boot image ISO file?                                                                                                                                     |          |
|                                                                                    | If you plan to use a repository to get the AI boot image, you will need the following information:<br><br>What is the publisher for the boot image <code>install-image/solaris-auto-install</code> package?<br><br>What is the repository of the publisher?<br><br>What is the architecture of the cluster nodes? |          |
|                                                                                    | For repositories that use HTTPS:<br><br>What is the full path of the certificate file for the repository?<br><br>What is the full path of the private key file for the repository?<br><br>You can request and download the key and certificate from the <code>http://pkg-register.oracle.com</code> site.         |          |
| Unified Archives                                                                   | If you plan to use the Unified Archives to install, you will need the following information:<br><br>What is the location of the unified archive?                                                                                                                                                                  |          |
| Custom Automated Installer<br>User root Password                                   | What is the password for the root account of the cluster nodes?                                                                                                                                                                                                                                                   |          |
| Custom Automated Installer<br>Repositories (when not using<br>the Unified Archive) | What is the repository of publisher <code>solaris</code> ?                                                                                                                                                                                                                                                        |          |
|                                                                                    | What is the repository of publisher <code>ha-cluster</code> ?                                                                                                                                                                                                                                                     |          |
|                                                                                    | For repositories that use HTTPS:<br><br>What is the full path of the certificate file for the repository?<br><br>What is the full path of the private key file for the repository?<br><br>You can request and download the key and certificate from the <code>http://pkg-register.oracle.com</code> site.         |          |
|                                                                                    | Select the Oracle Solaris Cluster components that you want to install. ( <i>Select one or more group packages to install.</i> )                                                                                                                                                                                   |          |
|                                                                                    | Do you want to select any individual components that are contained in these group packages?                                                                                                                                                                                                                       | Yes   No |

| 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 cluster nodes that are planned for the initial cluster configuration. <i>(For a single-node cluster, press Control-D alone.)</i>                                                                                                                                                                                |                                                                                      |
|                                                                                                                 | Confirm that the auto-discovered MAC address for each node is correct.                                                                                                                                                                                                                                                                |                                                                                      |
| Authenticating Requests to Add Nodes<br><i>(multiple-node cluster only)</i>                                     | Do you need to use DES authentication?                                                                                                                                                                                                                                                                                                | No   Yes                                                                             |
| Network Address for the Cluster Transport<br><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?                                                                                                                                                                                                                                                                                            | Yes   No                                                                             |
|                                                                                                                 | If no, what are the maximum numbers of nodes, private networks, and zone clusters that you expect to configure in the cluster? <i>(The maximum number of exclusive-IP zone clusters is three. The combined number of exclusive-IP and shared-IP zone clusters must be less than or equal to the maximum number of zone clusters.)</i> | ____ nodes<br>____ networks<br>____ zone clusters<br>____ exclusive-IP zone clusters |
|                                                                                                                 | Which netmask do you want to use? <i>Choose from the values that are calculated by scinstall or supply your own.</i>                                                                                                                                                                                                                  | ____.____.____.____                                                                  |
| Minimum Number of Private Networks<br><i>(multiple-node cluster only)</i>                                       | Should this cluster use at least two private networks?                                                                                                                                                                                                                                                                                | Yes   No                                                                             |
| Point-to-Point Cables<br><i>(two-node cluster only)</i>                                                         | Does this cluster use switches?                                                                                                                                                                                                                                                                                                       | Yes   No                                                                             |
| Cluster Switches<br><i>(multiple-node cluster only)</i>                                                         | Transport switch name, if used:<br>Defaults: switch1 and switch2                                                                                                                                                                                                                                                                      | First:<br>Second:                                                                    |
| Cluster Transport Adapters and Cables<br><i>(multiple-node cluster only)</i><br><br><i>(VLAN adapters only)</i> | First node name:                                                                                                                                                                                                                                                                                                                      |                                                                                      |
|                                                                                                                 | Transport adapter name:                                                                                                                                                                                                                                                                                                               | First:<br>Second:                                                                    |
|                                                                                                                 | Will this be a dedicated cluster transport adapter? <i>(Answer No if using tagged VLAN adapters.)</i>                                                                                                                                                                                                                                 | First: Yes   No<br>Second: Yes   No                                                  |
|                                                                                                                 | If no, what is the VLAN ID for this adapter?                                                                                                                                                                                                                                                                                          | First:<br>Second:                                                                    |
|                                                                                                                 | Where does each transport adapter connect to <i>(a switch or another adapter)</i> ?<br>Switch defaults: switch1 and switch2                                                                                                                                                                                                           | First:<br>Second:                                                                    |

| Component                                                                          | Description/Example                                                                                                                                                                                        | Answer                              |
|------------------------------------------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------------------------|
|                                                                                    | If a transport switch, do you want to use the default port name?                                                                                                                                           | First: Yes   No<br>Second: Yes   No |
|                                                                                    | If no, what is the name of the port that you want to use?                                                                                                                                                  | First:<br>Second:                   |
| <i>Specify for each additional node</i><br><br><i>(multiple-node cluster only)</i> | Node name:                                                                                                                                                                                                 |                                     |
|                                                                                    | Transport adapter name:                                                                                                                                                                                    | First:<br>Second:                   |
|                                                                                    | Where does each transport adapter connect to ( <i>a switch or another adapter</i> )?<br>Switch defaults: switch1 and switch2                                                                               | First:<br>Second:                   |
|                                                                                    | If a transport switch, do you want to use the default port name?                                                                                                                                           | First: Yes   No<br>Second: Yes   No |
|                                                                                    | If no, what is the name of the port that you want to use?                                                                                                                                                  | First:<br>Second:                   |
| Global Fencing                                                                     | Do you want to disable global fencing? <i>Answer No unless the shared storage does not support SCSI reservations or unless you want systems that are outside the cluster to access the shared storage.</i> | First: Yes   No<br>Second: Yes   No |
| Quorum Configuration<br><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 as a quorum device.)</i> | First: Yes   No<br>Second: Yes   No |

---

**Note** - If your physically clustered machines are configured with Oracle VM Server for SPARC, install the Oracle Solaris Cluster software only in I/O domains or guest domains.

---

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.

Perform the following tasks:

- Ensure that the hardware setup is complete and connections are verified before you install Solaris software. See the [“Oracle Solaris Cluster 4.2 Hardware Administration Manual”](#) and your server and storage device documentation for details on how to set up the hardware.
- Ensure that an Automated Installer install server and a DHCP server are configured. See [Part III, “Installing Using an Install Server,”](#) in [“Installing Oracle Solaris 11.2 Systems”](#).

- Determine the Ethernet address of the cluster node and the length of the subnet mask of the subnet that the address belongs to.
- Determine the MAC address of each cluster node.
- Ensure that your cluster configuration planning is complete. See [“How to Prepare for Cluster Software Installation” on page 47](#) for requirements and guidelines.
- Have available the root user password for the cluster nodes.
- SPARC: If you are configuring Oracle VM Server for SPARC I/O domains or guest domains as cluster nodes, ensure that the Oracle VM Server for SPARC software is installed on each physical machine and that the domains meet Oracle Solaris Cluster requirements. See [“How to Install Oracle VM Server for SPARC Software and Create Domains” on page 58](#).
- If you plan to install from Unified Archives that are created on an existing cluster, have the path to the archive file and ensure that it can be accessed from the AI server.
- If you plan to install from IPS repositories, determine which Oracle Solaris Cluster software packages you want to install.

The following table lists the group packages for the Oracle Solaris Cluster 4.2 software that you can choose during an AI installation and the principal features that each group package contains. You must install at least the `ha-cluster-framework-minimal` group package.

| Feature                  | ha-cluster-framework-full | ha-cluster-data-services-full | ha-cluster-framework-minimal | ha-cluster-geo-full | manager |
|--------------------------|---------------------------|-------------------------------|------------------------------|---------------------|---------|
| Framework                | X                         | X                             | X                            | X                   | X       |
| Agents                   |                           | X                             |                              |                     |         |
| Localization             | X                         |                               |                              |                     |         |
| Framework man pages      | X                         |                               |                              |                     |         |
| Data Service man pages   |                           | X                             |                              |                     |         |
| Agent Builder            | X                         |                               |                              |                     |         |
| Generic Data Service     | X                         | X                             |                              |                     |         |
| Graphical User Interface |                           |                               |                              |                     | X       |
| Geographic Edition       |                           |                               |                              | X                   |         |

- Have available your completed Typical Mode or Custom Mode installation worksheet. See [“Establishing a New Oracle Solaris Cluster With the Automated Installer” on page 89](#).

## ▼ How to Install and Configure Oracle Solaris and Oracle Solaris Cluster Software (IPS Repositories)

You can set the AI server to install both Oracle Solaris OS and the Oracle Solaris Cluster framework and data service software from IPS repositories or the Unified Archives on all global-cluster nodes and establish the cluster. This procedure describes how to set up and use the `scinstall(1M)` custom Automated Installer installation method to install and configure the cluster from IPS repositories.

### 1. Set up your Automated Installer (AI) install server and DHCP server.

Ensure that the AI install server meets the following requirements.

- The install server is on the same subnet as the cluster nodes.
- The install server is not itself a cluster node.
- The install server runs a release of the Oracle Solaris OS that is supported by the Oracle Solaris Cluster software.
- Each new cluster node is configured as a custom AI installation client that uses the custom AI directory that you set up for Oracle Solaris Cluster installation.

Follow the appropriate instructions for your software platform and OS version to set up the AI install server and DHCP server. See [Chapter 8, “Setting Up an AI Server,”](#) in [“Installing Oracle Solaris 11.2 Systems ”](#) and [“Working With DHCP in Oracle Solaris 11.2 ”](#).

### 2. On the AI install server, assume the root role.

### 3. On the AI install server, install the Oracle Solaris Cluster AI support package.

#### a. Ensure that the `solaris` and `ha-cluster` publishers are valid.

```
installserver# pkg publisher
PUBLISHER TYPE STATUS URI
solaris origin online solaris-repository
ha-cluster origin online ha-cluster-repository
```

#### b. Install the cluster AI support package.

```
installserver# pkg install ha-cluster/system/install
```

### 4. On the AI install server, start the `scinstall` utility.

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

The `scinstall` Main Menu is displayed.

### 5. Select option 1 or option 2 from the Main Menu.

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



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

- \* 1) Install, restore, replicate, and configure a cluster from this Automated Installer install server
- \* 2) Securely install, restore, replicate, and configure a cluster from this Automated Installer install server
- \* 3) Print release information for this Automated Installer install server
  
- \* ?) Help with menu options
- \* q) Quit

Option:

6. **Follow the menu prompts to supply your answers from the configuration planning worksheet.**
  
7. **For each node, confirm the options you chose so that the `scinstall` utility performs the necessary configuration to install the cluster nodes from this AI server.**

The utility also prints instructions to add the DHCP macros on the DHCP server, and adds (if you chose secure installation) or clears (if you chose non-secure installation) the security keys for SPARC nodes. Follow those instructions.

8. **(Optional) To install extra software packages or to customize the target device, update the AI manifest for each node.**

The AI manifest is located in the following directory:

```
/var/cluster/logs/install/autosinstall.d/ \
cluster-name/node-name/node-name_aimanifest.xml
```

- a. **To install extra software packages, edit the AI manifest as follows:**

- Add the publisher name and the repository information. For example:

```
<publisher name="aie">
<origin name="http://aie.us.oracle.com:12345"/>
</publisher>
```

- Add the package names that you want to install, in the `software_data` item of the AI manifest.

- b. **To customize the target device, update the target element in the manifest file.**

`scinstall` assumes the existing boot disk in the manifest file to be the target device. To customize the target device, update the `target` element in the manifest file based on how you want to use the supported criteria to locate the target device for the installation. For example, you can specify the `disk_name` sub-element.

For more information, see [Part III, “Installing Using an Install Server,”](#) in [“Installing Oracle Solaris 11.2 Systems”](#) and the `ai_manifest(4)` man page.

**c. Run the `installadm` command for each node.**

```
installadm update-manifest -n cluster-name-{sparc|i386} \
-f /var/cluster/logs/install/autoscinstall.d/cluster-name/node-name/node-
name_aimanifest.xml \
-m node-name_manifest
```

Note that SPARC and i386 is the architecture of the cluster node.

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

- **If `pconsole` software is installed and configured on your administrative console, use the `pconsole` utility to display the individual console screens.**

As the root role, use the following command to start the `pconsole` utility:

```
adminconsole# pconsole host[:port] [...] &
```

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

**10. Shut down and boot each node to start the AI installation.**

The Oracle Solaris software is installed with the default configuration.

---

**Note** - You cannot use this method if you want to customize the Oracle Solaris installation. If you choose the Oracle Solaris interactive installation, the Automated Installer is bypassed and Oracle Solaris Cluster software is not installed and configured. To customize Oracle Solaris during installation, instead follow instructions in [“How to Install Oracle Solaris Software” on page 48](#), then install and configure the cluster by following instructions in [“How to Install Oracle Solaris Cluster Software Packages” on page 59](#).

---

- **SPARC:**

- a. Shut down each node.**

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

- b. Boot the node with the following command**

```
ok boot net:dhcp - install
```

---

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

---

■ **x86:**

**a. Reboot the node.**

```
reboot -p
```

**b. During PXE boot, press Control-N.**

The GRUB menu is displayed.

**c. Immediately select the Automated Install entry and press Return.**

---

**Note** - If you do not select the Automated Install entry within 20 seconds, installation proceeds using the default interactive text installer method, which will not install and configure the Oracle Solaris Cluster software.

---

On each node, a new boot environment (BE) is created and Automated Installer installs the Oracle Solaris OS and Oracle Solaris Cluster software. When the installation is successfully completed, each node is fully installed as a new cluster node. Oracle Solaris Cluster installation output is logged in the `/var/cluster/logs/install/scinstall.log.N` file and the `/var/cluster/logs/install/sc_ai_config.log` file on each node.

- 11. If you intend to use the HA for NFS data service (HA for NFS) on a highly available local file system, exclude from the automounter map all shares that are part of the highly available local file system that is exported by HA for NFS.**

See [“Administrative Tasks Involving Maps”](#) in [“Managing Network File Systems in Oracle Solaris 11.2”](#) for more information about modifying the automounter map.

- 12. (x86 only) Set the default boot file.**

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

```
grub edit> kernel /platform/i86pc/kernel/amd64/unix -B $ZFS-BOOTFS -k
```

For more information, see [“How to Boot a System With the Kernel Debugger \(kldb\) Enabled”](#) in [“Booting and Shutting Down Oracle Solaris 11.2 Systems”](#).

- 13. If you performed a task that requires a cluster reboot, reboot the cluster.**

The following tasks require a reboot:

- Installing software updates that require a node or cluster reboot
- Making configuration changes that require a reboot to become active

**a. On one node, assume the root role.**

**b. Shut down the cluster.**

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

---

**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 137](#).

---

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

■ **SPARC:**

```
ok boot
```

■ **x86:**

When the GRUB menu is displayed, select the appropriate Oracle Solaris entry and press Enter.

For more information about GRUB based booting, see [“Booting a System” in “Booting and Shutting Down Oracle Solaris 11.2 Systems”](#).

The cluster is established when all nodes have successfully booted into the cluster. Oracle Solaris Cluster installation output is logged in a `/var/cluster/logs/install/scinstall.log.N` file.

**14. From one node, 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.

**15. If you plan to enable RPC use of TCP wrappers, add all `clprivnet0` IP addresses to the `/etc/hosts.allow` file on each cluster node.**

Without this addition to the `/etc/hosts.allow` file, TCP wrappers prevent internode communication over RPC for cluster administration utilities.

**a. On each node, display the IP addresses for all `clprivnet0` devices on the node.**

```

/usr/sbin/ipadm show-addr
ADDROBJ TYPE STATE ADDR
clprivnet0/N static ok ip-address/netmask-length
...

```

**b. On each cluster node, add to the `/etc/hosts.allow` file the IP addresses of all `clprivnet0` devices in the cluster.**

**16. (Optional) On each node, enable automatic node reboot if all monitored shared-disk paths fail.**

---

**Note** - At initial configuration time, `disk-path` monitoring is enabled by default for all discovered devices.

---

**a. Enable automatic reboot.**

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

-p

Specifies the property to set

```
reboot_on_path_failure=enable
```

Enables automatic node reboot if failure of all monitored shared-disk paths occurs.

**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

```

...

**17. If you use the LDAP naming service, you must manually configure it on the cluster nodes after they boot.**

**Next Steps**

**1. Perform all of the following procedures that are appropriate for your cluster configuration.**

- [“How to Configure Internal Disk Mirroring” on page 57](#)
- [“How to Install Oracle VM Server for SPARC Software and Create Domains” on page 58](#)
- [“How to Install Sun QFS Software” on page 65](#)
- [“How to Set Up the Root Environment” on page 65](#)
- [“How to Configure IP Filter” on page 66](#)

**2. Configure quorum, if not already configured, and perform post installation tasks.**

- If you installed a multiple-node cluster and accepted automatic quorum configuration, post installation setup is complete. Go to [“How to Verify the Quorum Configuration and Installation Mode” on page 142](#).
- If you installed a multiple-node cluster and declined automatic quorum configuration, perform post installation setup. Go to [“How to Configure Quorum Devices” on page 137](#).
- If you added a node to an existing two-node cluster, go to [“How to Update Quorum Devices After Adding a Node to a Global Cluster” on page 134](#).
- If you added a new node to an existing cluster with at least three nodes that uses a quorum device, go to [“How to Update Quorum Devices After Adding a Node to a Global Cluster” on page 134](#).
- If you added a new node to an existing cluster with at least three nodes 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 142](#).
- If you installed a single-node cluster, cluster establishment is complete. Go to [“Creating Cluster File Systems” on page 165](#) to install volume management software and configure the cluster.

**Troubleshooting**

**Disabled scinstall option** – If the AI option of the `scinstall` command is not preceded by an asterisk, the option is disabled. This condition indicates that AI 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 8](#) to correct the AI setup, then restart the `scinstall` utility.

▼ **How to Install and Configure Oracle Solaris and Oracle Solaris Cluster Software (Unified Archives)**

You will use the AI server to install a cluster from the Unified Archives and configure its nodes. This procedure retains all the software packages that are contained in the Unified Archives, but you must provide the new cluster configuration that you designed in the worksheet. Before you

perform this procedure, you must first create the archive. See [Step 1](#) below for instructions on creating the recovery archive.

The AI server sets up installation of the nodes from the Unified Archives and creates the cluster with the new configuration. Only a Unified Archive created in the global zone is accepted. You can use either a *clone* archive or a *recovery* archive. The following list describes the differences between the two archives:

- When you install from a clone archive, only the global zone is installed. Any zones in the archive are not installed. When you install from a recovery archive, both the global zone and the zones contained in the archive are installed.
- A clone archive does not contain system configuration, including IPMPs, VLANs, and VNICs.
- A clone archive only contains the BE that is active when the archive is created, therefore only that BE is installed. A recovery archive can contain multiple BEs, but only the active BE is updated with the new cluster configuration.

This procedure prompts you for the cluster name, node names and their MAC addresses, the path to the Unified Archives, and the cluster configuration you designed in the worksheet.

**1. Assume the root role on a node of the global cluster and create an archive.**

```
phys-schost# archiveadm create -r archive-location
```

Use the `create` command to create a clone archive or the `create -r` option to create a recovery archive. For more information on using the `archiveadm` command, see the [archiveadm\(1M\)](#) man page.

**2. Set up your Automated Installer (AI) install server and DHCP server.**

Ensure that the AI install server meets the following requirements.

- The install server is on the same subnet as the cluster nodes.
- The install server is not itself a cluster node.
- The install server runs a release of the Oracle Solaris OS that is supported by the Oracle Solaris Cluster software.
- Each new cluster node is configured as a custom AI installation client that uses the custom AI directory that you set up for Oracle Solaris Cluster installation.

Follow the appropriate instructions for your software platform and OS version to set up the AI install server and DHCP server. See [Chapter 8, “Setting Up an AI Server,”](#) in [“Installing Oracle Solaris 11.2 Systems ”](#) and [“Working With DHCP in Oracle Solaris 11.2 ”](#).

**3. Log into the Automated Installer server and assume the root role.**

**4. On the AI install server, install the Oracle Solaris Cluster AI support package.**

- a. **Ensure that the `solaris` and `ha-cluster` publishers are valid.**

```
installserver# pkg publisher
PUBLISHER TYPE STATUS URI
solaris origin online solaris-repository
ha-cluster origin online ha-cluster-repository
```

**b. Install the cluster AI support package.**

```
installserver# pkg install ha-cluster/system/install
```

**5. On the AI install server, start the `scinstall` utility.**

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

The `scinstall` Main Menu is displayed.

**6. Type the option number and press Return.**

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

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

```
* 1) Install, restore, or replicate a cluster from this Automated Installer server
* 2) Securely install, restore, or replicate a cluster from this Automated Installer server
* 3) Print release information for this Automated Installer install server
```

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

```
* q) Quit
```

```
Option: 2
```

Choose Option 1 if you want to install a cluster using a non-secure AI server installation.

Choose Option 2 for a secure AI installation.

The Custom Automated Installer Menu or Custom Secure Automated Installer Menu is displayed.

**7. Type the option number to Install and Configure a New Cluster from Unified Archives and press Return.**

The Custom Automated Installer User screen is displayed.

**8. Type the password and press Return.**

Type the password a second time to confirm it. The Typical or Customer Mode screen is displayed.

**9. Type the option number for the install mode you will use.**

The Cluster Name screen is displayed.

**10. Type the name of the cluster you want to install and press Return.**

The Cluster Nodes screen is displayed.



**11. Types the names of the cluster nodes that you plan to install from the Unified Archives and press Return.**

If the `scinstall` utility is unable to find the MAC address of the nodes, type in each address when prompted and press Return. You can then choose to install all the nodes from the same archive, or use a different archive for each node.

**12. Type the full path to the archive and press Return.**

The archive can either be a recovery archive or a clone archive.

The Cluster Transport Adapters and Cables screen is displayed.

**13. Type the names of the cluster transport adapters and press Return.**

Select the type of each transport adapter. The Resource Security Configuration screen is displayed.

**14. Choose to enable to disable automatic quorum device selection and press Return.**

The Confirmation screen is displayed.

**15. For each node, confirm the options you chose so that the `scinstall` utility performs the necessary configuration to install the cluster nodes from this AI server.**

The utility also prints instructions to add the DHCP macros on the DHCP server, and adds (if you chose secure installation) or clears (if you chose non-secure installation) the security keys for SPARC nodes. Follow those instructions.

**16. (Optional) To customize the target device, update the AI manifest for each node.**

The AI manifest is located in the following directory:

```
/var/cluster/logs/install/autosinstall.d/ \
cluster-name/node-name/node-name_aimanifest.xml
```

**a. To customize the target device, update the target element in the manifest file.**

`scinstall` assumes the existing boot disk in the manifest file to be the target device. To customize the target device, update the target element in the manifest file based on how you want to use the supported criteria to locate the target device for the installation. For example, you can specify the `disk_name` sub-element.

For more information, see [Part III, “Installing Using an Install Server,”](#) in [“Installing Oracle Solaris 11.2 Systems”](#) and the `ai_manifest(4)` man page.

**b. Run the `installadm` command for each node.**

```
installadm update-manifest -n cluster-name-{sparc|i386} \

```

```
-f /var/cluster/logs/install/autoscinstall.d/cluster-name/node-name/node-
name_aimanifest.xml \
-m node-name_manifest
```

Note that SPARC and i386 is the architecture of the cluster node.

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

- **If pconsole software is installed and configured on your administrative console, use the pconsole utility to display the individual console screens.**

As the root role, use the following command to start the pconsole utility:

```
adminconsole# pconsole host[:port] [...] &
```

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

**18. Shut down and boot each node to start the AI installation.**

The Oracle Solaris software is installed with the default configuration.

---

**Note** - You cannot use this method if you want to customize the Oracle Solaris installation. If you choose the Oracle Solaris interactive installation, the Automated Installer is bypassed and Oracle Solaris Cluster software is not installed and configured. To customize Oracle Solaris during installation, instead follow instructions in [“How to Install Oracle Solaris Software” on page 48](#), then install and configure the cluster by following instructions in [“How to Install Oracle Solaris Cluster Software Packages” on page 59](#).

---

- **SPARC:**

- a. **Shut down each node.**

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

- b. **Boot the node with the following command**

```
ok boot net:dhcp - install
```

---

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

---

- **x86:**

**a. Reboot the node.**

```
reboot -p
```

**b. During PXE boot, press Control-N.**

The GRUB menu is displayed.

**c. Immediately select the Automated Install entry and press Return.**

---

**Note** - If you do not select the Automated Install entry within 20 seconds, installation proceeds using the default interactive text installer method, which will not install and configure the Oracle Solaris Cluster software.

---

Each node will be automatically rebooted a few times before the node completely joins the cluster. Ignore any error messages from SMF services on the console. On each node, the Automated Installer installs the software that is contained in the Unified Archives. When the installation is successfully completed, each node is fully installed as a new cluster node. Oracle Solaris Cluster installation output is logged in the `/var/cluster/logs/install/scinstall.log.N` file and the `/var/cluster/logs/install/sc_ai_config.log` file on each node.

**19. From one node, 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.

**20. If you plan to enable RPC use of TCP wrappers, add all `clprivnet0` IP addresses to the `/etc/hosts.allow` file on each cluster node.**

Without this addition to the `/etc/hosts.allow` file, TCP wrappers prevent internode communication over RPC for cluster administration utilities.

**a. On each node, display the IP addresses for all `clprivnet0` devices on the node.**

```
/usr/sbin/ipadm show-addr
ADDROBJ TYPE STATE ADDR
clprivnet0/N static ok ip-address/netmask-length
...
```

- b. **On each cluster node, add to the `/etc/hosts.allow` file the IP addresses of all `clprivnet0` devices in the cluster.**
21. **If you use the LDAP naming service, you must manually configure it on the cluster nodes after they boot.**

## ▼ How to Replicate a Cluster from the Unified Archives

You can use the Unified Archives to replicate a cluster and its nodes. This procedure retains all the software packages in the archives. Furthermore, this new cluster will have the exact configuration as the archive cluster or you can customize the private network properties and host identities, such as zone host names and logical host names in cluster resources.

Only the Unified Archive created in the global zone is accepted. You can use either a clone archive or a recovery archive. The following list describes the differences between the two archives:

- When you install from a clone archive, only the global zone is installed. Any zones in the archive are not installed. When you install from a recovery archive, both the global zone and the zones contained in the archive are installed.
- A clone archive does not contain system configuration, including IPMPs, VLANs, and VNICs.
- A clone archive only contains the BE that is active when the archive is created, therefore only that BE is installed. A recovery archive can contain multiple BEs, but only the active BE is updated with the new cluster configuration.

To replicate a cluster from the Unified Archives created on an existing cluster, the hardware configuration of the new cluster must be the same as the source cluster. The number of nodes in the new cluster must be the same as in the source cluster, and the transport adapters must also be the same as in the source cluster.

1. **Assume the root role on a node of the global cluster and create an archive.**

```
phys-schost# archiveadm create -r archive-location
```

Use the `create` command to create a clone archive or the `-r` option to create a recovery archive. When you create the archive, exclude the ZFS datasets that are on the shared storage. If you plan to migrate the data on the shared storage from the source cluster to the new cluster, use the traditional method.

For more information on using the `archiveadm` command, see the [archiveadm\(1M\)](#) man page.

## 2. Set up your Automated Installer (AI) install server and DHCP server.

Ensure that the AI install server meets the following requirements.

- The install server is on the same subnet as the cluster nodes.
- The install server is not itself a cluster node.
- The install server runs a release of the Oracle Solaris OS that is supported by the Oracle Solaris Cluster software.
- Each new cluster node is configured as a custom AI installation client that uses the custom AI directory that you set up for Oracle Solaris Cluster installation.

Follow the appropriate instructions for your software platform and OS version to set up the AI install server and DHCP server. See [Chapter 8, “Setting Up an AI Server,”](#) in [“Installing Oracle Solaris 11.2 Systems”](#) and [“Working With DHCP in Oracle Solaris 11.2”](#).

## 3. Log into the Automated Installer server and assume the root role.

## 4. On the AI install server, install the Oracle Solaris Cluster AI support package.

### a. Ensure that the `solaris` and `ha-cluster` publishers are valid.

```
installserver# pkg publisher
PUBLISHER TYPE STATUS URI
solaris origin online solaris-repository
ha-cluster origin online ha-cluster-repository
```

### b. Install the cluster AI support package.

```
installserver# pkg install ha-cluster/system/install
```

## 5. On the AI install server, start the `scinstall` utility.

```
phys-schost# scinstall
```

The `scinstall` Main Menu is displayed.

## 6. Type the option number and press Return.

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

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

- \* 1) Install, restore, or replicate a cluster from this Automated Installer server
- \* 2) Securely install, restore, or replicate a cluster from this Automated Installer server
- \* 3) Print release information for this Automated Installer install server

- \* ?) Help with menu options
- \* q) Quit

Option: 2

Choose Option 1 if you want to replicate a cluster using a non-secure AI server installation. Choose Option 2 for a secure AI replication.

The Custom Automated Installer Menu or Custom Secure Automated Installer Menu is displayed.

7. **Type the option number to Replicate a Cluster from Unified Archives and press Return.**

The Custom Automated Installer User screen is displayed.

8. **Type the password and press Return.**

Type the password a second time to confirm it.

The Cluster Name screen is displayed.

9. **Type the name of the cluster you want to replicate and press Return.**

The Cluster Nodes screen is displayed.

10. **Types the names of the cluster nodes that you plan to replicate from the Unified Archives.**

After you type the node names, press Control-D and then Return. If the `scinstall` utility is unable to find the MAC address of the nodes, type in each address when prompted and press Return.

11. **Type the full path to the archive for each node.**

A Unified Archive file must be created for each node in the source cluster, and only one archive can be specified per node in the new cluster. This 1:1 mapping ensures that one archive is mapped to one node in the source cluster. Similarly, the archive of one source node must be mapped to only one node in the new cluster.

Press Return to confirm the archive files.

12. **If you want to use a different private network address and netmask, specify them in the Network Address for the Cluster Transport menu.**

13. **Provide the path to the text file that contains the mapping from old host identities in the source cluster to the new host identities in the new cluster.**

To avoid using the same host identities in the new cluster as the source cluster, you can create and provide a text file that contains a 1:1 mapping from the old host identities in the source cluster to the new host identities that you intend to use in the new cluster. The text file can contain multiple lines, where each line has two columns. The first column is the hostname used in the source cluster, and the second column is the corresponding new hostname in the

new cluster. The hostnames are for the logical hostnames, shared address resources, and zone clusters. For example:

```
old-cluster-zc-host1 new-cluster-zc-host1
old-cluster-zc-host2 new-cluster-zc-host2
old-cluster-lh-1 new-cluster-lh1
old-cluster-lh-2 new-cluster-lh2
```

The Confirmation screen is displayed.

14. **For each node, confirm the options you chose so that the `scinstall` utility performs the necessary configuration to install the cluster nodes from this AI server.**

The utility also prints instructions to add the DHCP macros on the DHCP server, and adds or clears the security keys for SPARC nodes (if you chose secure installation). Follow those instructions.

15. **(Optional) To customize the target device, update the AI manifest for each node.**

The AI manifest is located in the following directory:

```
/var/cluster/logs/install/autosinstall.d/ \
cluster-name/node-name/node-name_aimanifest.xml
```

- a. **To customize the target device, update the target element in the manifest file.**

`scinstall` assumes the existing boot disk in the manifest file to be the target device. To customize the target device, update the target element in the manifest file based on how you want to use the supported criteria to locate the target device for the installation. For example, you can specify the `disk_name` sub-element.

For more information, see [Part III, “Installing Using an Install Server,”](#) in “Installing Oracle Solaris 11.2 Systems” and the `ai_manifest(4)` man page.

- b. **Run the `installadm` command for each node.**

```
installadm update-manifest -n cluster-name-{sparc|i386} \
-f /var/cluster/logs/install/autosinstall.d/cluster-name/node-name/node-
name_aimanifest.xml \
-m node-name_manifest
```

Note that SPARC and i386 is the architecture of the cluster node.

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

- **If `pconsole` software is installed and configured on your administrative console, use the `pconsole` utility to display the individual console screens.**

As the root role, use the following command to start the `pconsole` utility:

```
adminconsole# pconsole host[:port] [...] &
```

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

#### 17. Shut down and boot each node to start the AI installation.

The Oracle Solaris software is installed with the default configuration.

---

**Note** - You cannot use this method if you want to customize the Oracle Solaris installation. If you choose the Oracle Solaris interactive installation, the Automated Installer is bypassed and Oracle Solaris Cluster software is not installed and configured. To customize Oracle Solaris during installation, instead follow instructions in [“How to Install Oracle Solaris Software” on page 48](#), then install and configure the cluster by following instructions in [“How to Install Oracle Solaris Cluster Software Packages” on page 59](#).

---

- **SPARC:**

- a. **Shut down each node.**

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

- b. **Boot the node with the following command**

```
ok boot net:dhcp - install
```

---

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

---

- **x86:**

- a. **Reboot the node.**

```
reboot -p
```

- b. **During PXE boot, press Control-N.**

The GRUB menu is displayed.

- c. **Immediately select the Automated Install entry and press Return.**



---

**Note** - If you do not select the Automated Install entry within 20 seconds, installation proceeds using the default interactive text installer method, which will not install and configure the Oracle Solaris Cluster software.

---

Each node will be automatically rebooted a few times before the node completely joins the cluster. Ignore any error messages from SMF services on the console. Each node is installed with the software contained in the Unified Archives. When the installation is successfully completed, each node is booted as a member of the new cluster, with the same cluster configuration as the archive but with a different system identity and system configuration. Oracle Solaris Cluster installation output is logged in the `/var/cluster/logs/install/scinstall.log.N` file and the `/var/cluster/logs/install/sc_ai_config.log` file on each node.

**18. From one node, 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.

**19. The cluster objects, including resource groups and zone clusters, are offline after the last reboot. Check the configuration and make necessary changes in the new environment before bringing them online.**

If the source cluster uses another system as a cluster object (for example, using a system as a quorum device of the *quorum server* type), you must manually adjust the configuration both in the new cluster and on the quorum server in order for the device to work. For a quorum server, you can add a new quorum server quorum device and remove the one brought from the archive.

---

**Note** - If your source cluster uses Oracle Solaris Cluster Geographic Edition, follow the procedures in [Chapter 5, “Administering Cluster Partnerships,”](#) in [“Oracle Solaris Cluster Geographic Edition System Administration Guide”](#) to rename a cluster and reconstruct the partnership.

---

If you need to make any changes to the zone cluster configuration or the resource groups in the cluster, reboot the zone cluster to Offline Running mode:

```
phys-schost#clzonecluster reboot -o zoneclustername
```

If you do not plan to make changes to the zone cluster configuration, you can reboot the cluster to Online Running mode:

```
phys-schost #clzonecluster reboot zoneclustername
```

You can also check the log file, `/var/cluster/logs/install/sc_ai_config`, for more information.

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

Perform this procedure on existing global-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 [“Oracle Solaris Cluster 4.2 Hardware Administration Manual”](#).
  - Verify that any existing cluster interconnects can support the new node. See the [“Oracle Solaris Cluster 4.2 Hardware Administration Manual”](#).
  - Ensure that any additional storage is installed.

### 1. Add the name of the new node to the cluster's authorized-nodes list.

a. On any node, assume the root role.

b. Start the `clsetup` utility.

```
phys-schost# clsetup
```

The Main Menu is displayed.

c. Choose the **New Nodes** menu item.

d. Choose the **Specify the Name of a Machine Which May Add Itself** menu item.

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.**
2. **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 3](#).**
- **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 Cluster Interconnect menu item.**

- c. **Choose the Add a Transport Cable menu item.**

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 2.2.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.

3. **Ensure that the private-network configuration can support the nodes and private networks that you are adding.**

- a. **Display the maximum numbers of nodes, private networks, and zone clusters that the current private-network configuration supports.**

```
phys-schost# cluster show-netprops
```

The output looks similar to the following:

```
=== Private Network ===

private_netaddr: 172.16.0.0
private_netmask: 255.255.240.0
max_nodes: 64
max_privatenets: 10
max_zoneclusters: 12
```

**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 Oracle Solaris Cluster Software on Additional Global-Cluster Nodes \(scinstall\)”](#) on page 123.

- **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 116. 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 Oracle Solaris Cluster software on the new cluster nodes. Go to [“How to Configure Oracle Solaris Cluster Software on Additional Global-Cluster Nodes \(scinstall\)”](#) on page 123 or [“How to Configure Oracle Solaris Cluster Software on Additional Global-Cluster Nodes \(XML File\)”](#) on page 129.

## ▼ How to Change the Private Network Configuration When Adding Nodes or Private Networks

Perform this task to change the global-cluster's private IP address range to accommodate an increase in one or more of the following cluster components:

- The number of nodes or non-global zones
- The number of private networks
- The number of zone clusters

You can also use this procedure to decrease the private IP address range.

---

**Note** - This procedure requires you to shut down the entire cluster. If you need to change only the netmask, for example, to add support for zone clusters, do not perform this procedure. Instead, run the following command from a global-cluster node that is running in cluster mode to specify the expected number of zone clusters:

```
phys-schost# cluster set-netprops num_zoneclusters=N
```

This command does not require you to shut down the cluster.

---

1. **Assume the root role on a node of the cluster.**
2. **From one node, start the `clsetup` utility.**  

```
phys-schost# clsetup
```

The `clsetup` Main Menu is displayed.
3. **Switch each resource group offline.**
  - a. **Choose the Resource Groups menu item.**  
The Resource Group Menu is displayed.
  - b. **Choose the Online/Offline or Switchover a Resource Group menu item.**
  - 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. **Choose the Enable/Disable a Resource menu item.**
  - 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.

### ■ SPARC:

```
ok boot -x
```

### ■ x86:

- a. In the GRUB menu, use the arrow keys to select the appropriate Oracle Solaris entry and type `e` to edit its commands.

For more information about GRUB based booting, see [“Booting a System” in “Booting and Shutting Down Oracle Solaris 11.2 Systems”](#).

- b. In the boot parameters screen, use the arrow keys to select the kernel entry and type `e` to edit the entry.

- c. Add `-x` to the `multiboot` command to specify that the system boot into noncluster mode.

- d. Press Enter to accept the change and return to the boot parameters screen.

The screen displays the edited command.

**e. Type `y` 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 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. Choose the Change Network Addressing and Ranges for the Cluster Transport menu item.**

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.**

- **To change the default private network IP address:**

**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 `255.255.240.0`. This default IP address range supports up to 64 nodes, 12 zone clusters, and 10 private networks in the cluster.

- **To accept the default IP address range, type `yes` and press the Return key.**

- **To change the IP address range:**
  - 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, and zone clusters that you expect to configure in the cluster.
  - b. **Provide the number of nodes, private networks, and zone clusters 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, private networks, and zone clusters that you specified.
    - The second netmask supports twice the number of nodes, private networks, and zone clusters 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, private networks, and zone clusters.**
- 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.**
    - **SPARC:**

```
ok boot
```
    - **x86:**

When the GRUB menu is displayed, select the appropriate Oracle Solaris entry and press Enter.

For more information about GRUB based booting, see [“Booting a System”](#) in [“Booting and Shutting Down Oracle Solaris 11.2 Systems”](#).



**17. From one node, start the `clsetup` utility.**

```
clsetup
```

The `clsetup` Main Menu is displayed.

**18. Re-enable all disabled resources.**

**a. Choose the Resource Groups menu item.**

The Resource Group Menu is displayed.

**b. Choose the Enable/Disable a Resource menu item.**

**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. Choose the Online/Offline or Switchover a Resource Group menu item.**

**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 Control-C.

**Next Steps** To add a node to an existing cluster, go to one of the following procedures:

- [“How to Configure Oracle Solaris Cluster Software on Additional Global-Cluster Nodes \(`scinstall`\)” on page 123](#)
- [“How to Install and Configure Oracle Solaris and Oracle Solaris Cluster Software \(IPS Repositories\)” on page 96](#)
- [“How to Configure Oracle Solaris Cluster Software on Additional Global-Cluster Nodes \(XML File\)” on page 129](#)

## Configuring Oracle Solaris Cluster Software on Additional Global-Cluster Nodes (scinstall)

The `scinstall` utility runs in two modes of installation, Typical or Custom. For the Typical installation of Oracle Solaris Cluster software, `scinstall` automatically specifies the cluster transport switches as `switch1` and `switch2`.

Complete one of the following configuration planning worksheets. See [“Planning the Oracle Solaris OS” on page 12](#) and [“Planning the Oracle Solaris Cluster Environment” on page 15](#) 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? <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>cluster check</code> validation utility?	Yes   No
Autodiscovery of Cluster Transport	Do you want to use autodiscovery to configure the cluster transport? 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? Defaults: <code>switch1</code> and <code>switch2</code>	<i>First:</i> <i>Second:</i>
Cluster Transport Adapters and Cables	Transport adapter names:	<i>First:</i> <i>Second:</i>
	Where does each transport adapter connect to ( <i>a switch or another adapter</i> )? Switch defaults: <code>switch1</code> and <code>switch2</code>	<i>First:</i> <i>Second:</i>
	For transport switches, do you want to use the default port name?	<i>First:</i> Yes   No <i>Second:</i> Yes   No
	If no, what is the name of the port that you want to use?	<i>First:</i> <i>Second:</i>
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? <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>cluster check</code> validation utility?	Yes   No
Autodiscovery of Cluster Transport	Do you want to use autodiscovery to configure the cluster transport? 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	Transport switch name, if used: Defaults: <code>switch1</code> and <code>switch2</code>	<i>First:</i> <i>Second:</i>
Cluster Transport Adapters and Cables	Transport adapter name:	<i>First:</i> <i>Second:</i>
	Where does each transport adapter connect to ( <i>a switch or another adapter</i> )? Switch defaults: <code>switch1</code> and <code>switch2</code>	<i>First:</i> <i>Second:</i>
	If a transport switch, do you want to use the default port name?	<i>First:</i> Yes   No <i>Second:</i> Yes   No
	If no, what is the name of the port that you want to use?	<i>First:</i> <i>Second:</i>
Automatic Reboot	Do you want <code>scinstall</code> to automatically reboot the node after installation?	Yes   No

## ▼ How to Configure Oracle Solaris Cluster Software on Additional Global-Cluster Nodes (scinstall)

Perform this procedure to add a new node to an existing global cluster. To use Automated Installer to add a new node, follow the instructions in “[How to Install and Configure Oracle Solaris and Oracle Solaris Cluster Software \(IPS Repositories\)](#)” on page 96.

This procedure uses the interactive form of the `scinstall` command. For information about how to use the noninteractive forms of the `scinstall` command, such as when developing installation scripts, see the [scinstall\(1M\)](#) man page.

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.

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

- Ensure that the Oracle Solaris OS is installed to support the Oracle Solaris Cluster software. If the Oracle Solaris software is already installed on the node, you must ensure that the Oracle Solaris installation meets the requirements for the Oracle Solaris Cluster software and any other software that you intend to install on the cluster. See [“How to Install Oracle Solaris Software” on page 48](#) for more information about installing the Oracle Solaris software to meet Oracle Solaris Cluster software requirements.
- Ensure that NWAM is disabled. See [“How to Install Oracle Solaris Cluster Software Packages” on page 59](#) for instructions.
- SPARC: If you are configuring Oracle VM Server for SPARC I/O domains or guest domains as cluster nodes, ensure that the Oracle VM Server for SPARC software is installed on each physical machine and that the domains meet Oracle Solaris Cluster requirements. See [“How to Install Oracle VM Server for SPARC Software and Create Domains” on page 58](#).
- Ensure that Oracle Solaris Cluster software packages and updates are installed on the node. See [“How to Install Oracle Solaris Cluster Software Packages” on page 59](#).
- Ensure that the cluster is prepared for the addition of the new node. See [“How to Prepare the Cluster for Additional Global-Cluster Nodes” on page 114](#).
- Have available your completed Typical Mode or Custom Mode installation worksheet. See [“Configuring Oracle Solaris Cluster Software on Additional Global-Cluster Nodes \(scinstall\)” on page 122](#).

**1. On the cluster node to configure, assume the root role.**

**2. Ensure that TCP wrappers for RPC are disabled on all nodes of the cluster.**

The Oracle Solaris TCP wrappers for RPC feature prevents internode communication that is necessary for cluster configuration.

**a. On each node, display the status of TCP wrappers for RPC.**

TCP wrappers are enabled if `config/enable_tcpwrappers` is set to `true`, as shown in the following example command output.

```
svccfg -s rpc/bind listprop config/enable_tcpwrappers
config/enable_tcpwrappers boolean true
```

**b. If TCP wrappers for RPC are enabled on a node, disable TCP wrappers and refresh the RPC bind service.**

```
svccfg -s rpc/bind setprop config/enable_tcpwrappers = false
svcadm refresh rpc/bind
svcadm restart rpc/bind
```

### 3. Prepare public-network interfaces.

#### a. Create static IP addresses for each public-network interface.

```
ipadm create-ip interface
ipadm create-addr -T static -a local=address/prefix-length addrobj
```

For more information, see [“How to Configure an IPv4 Interface”](#) in [“Configuring and Administering Network Components in Oracle Solaris 11.2”](#).

#### b. (Optional) Create IPMP groups for public-network interfaces.

During initial cluster configuration, unless non-link-local IPv6 public network interfaces exist in the cluster, IPMP groups are automatically created based on matching subnets. These groups use transitive probes for interface monitoring and no test addresses are required.

If these automatically created IPMP groups would not meet your needs, or if IPMP groups would not be created because your configuration includes one or more non-link-local IPv6 public network interfaces, do one of the following:

- **Create the IPMP groups you need before you establish the cluster.**
- **After the cluster is established, use the `ipadm` command to edit the IPMP groups.**

For more information, see [“Configuring IPMP Groups”](#) in [“Administering TCP/IP Networks, IPMP, and IP Tunnels in Oracle Solaris 11.2”](#).

### 4. Start the `scinstall` utility.

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

The `scinstall` Main Menu is displayed.

### 5. Type the option number 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) Print release information for this cluster node

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

```
Option: 1
```

The New Cluster and Cluster Node Menu is displayed.

6. **Type the option number for Add This Machine as a Node in an Existing Cluster and press the Return key.**

7. **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.

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

9. **Verify on each node that multiuser services for the Service Management Facility (SMF) are online.**

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

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

10. **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 an Existing Cluster or Zone Cluster”](#) in [“Oracle Solaris Cluster System Administration Guide”](#) for procedures.

11. **From one node, 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. **If TCP wrappers are used in the cluster, ensure that the `clprivnet0` IP addresses for all added nodes are added to the `/etc/hosts.allow` file on each cluster node.**

Without this addition to the `/etc/hosts.allow` file, TCP wrappers prevent internode communication over RPC for cluster administration utilities.

- a. On each node, display the IP addresses for all `clprivnet0` devices.

```
# /usr/sbin/ipadm show-addr
ADDROBJ      TYPE      STATE      ADDR
clprivnet0/N static    ok         ip-address/netmask-length
...
```

- b. On each node, edit the `/etc/hosts.allow` file with the IP addresses of all `clprivnet0` devices in the cluster.

13. Verify that all necessary software updates are installed.

```
phys-schost# pkg list
```

14. (Optional) Enable automatic node reboot if all monitored shared-disk paths fail.

Note - At initial configuration time, disk-path monitoring is enabled by default for all discovered devices.

- a. Enable automatic reboot.

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

-p

Specifies the property to set

```
reboot_on_path_failure=enable
```

Enables automatic node reboot if failure of all monitored shared-disk paths occurs.

- 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
...
```

15. If you intend to use the HA for NFS data service (HA for NFS) on a highly available local file system, exclude from the automounter map all shares that are part of the highly available local file system that is exported by HA for NFS.

See [“Administrative Tasks Involving Maps”](#) in [“Managing Network File Systems in Oracle Solaris 11.2”](#) for more information about modifying the automounter map.

Example 3-3 Configuring Oracle Solaris 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 node "phys-schost-3" to the cluster configuration ... done
Adding adapter "net2" to the cluster configuration ... done
Adding adapter "net3" 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
Setting the node ID for "phys-schost-3" ... done (id=1)

Verifying the major number for the "did" driver from "phys-schost-1" ... done
Initializing NTP configuration ... done

Updating nsswitch.conf ... done

Adding cluster node 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

Ensure network routing is disabled ... done
Network routing has been disabled on this node by creating /etc/notrouter.
Having a cluster node act as a router is not supported by Oracle Solaris Cluster.
Please do not re-enable network routing.
Updating file ("ntp.conf.cluster") on node phys-schost-1 ... done
Updating file ("hosts") on node phys-schost-1 ... done

Log file - /var/cluster/logs/install/scinstall.log.6952

Rebooting ...
```

Troubleshooting **Unsuccessful configuration** – If one or more nodes cannot join the cluster, or if the wrong configuration information was specified, first attempt to perform this procedure again. If that does not correct the problem, perform the procedure [“How to Unconfigure Oracle Solaris Cluster Software to Correct Installation Problems” on page 227](#) on each misconfigured node to remove it from the cluster configuration. You do not need to uninstall the Oracle Solaris Cluster software packages. Then perform this procedure again.

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 Global Cluster” on page 134](#).

Otherwise, go to [“How to Verify the Quorum Configuration and Installation Mode” on page 142.](#)

▼ How to Configure Oracle Solaris Cluster Software on Additional Global-Cluster Nodes (XML File)

Perform this procedure to configure a new global-cluster node by using an XML cluster configuration file. The new node can be a duplication of an existing cluster node that runs the Oracle Solaris Cluster 4.2 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 Oracle Solaris OS is installed to support the Oracle Solaris Cluster software. If the Oracle Solaris software is already installed on the node, you must ensure that the Oracle Solaris installation meets the requirements for the Oracle Solaris Cluster software and any other software that you intend to install on the cluster. See [“How to Install Oracle Solaris Software” on page 48](#) for more information about installing the Oracle Solaris software to meet Oracle Solaris Cluster software requirements.
 - Ensure that NWAM is disabled. See [“How to Install Oracle Solaris Cluster Software Packages” on page 59](#) for instructions.
 - SPARC: If you are configuring Oracle VM Server for SPARC I/O domains or guest domains as cluster nodes, ensure that the Oracle VM Server for SPARC software is installed on each physical machine and that the domains meet Oracle Solaris Cluster requirements. See [“How to Install Oracle VM Server for SPARC Software and Create Domains” on page 58.](#)
 - Ensure that Oracle Solaris Cluster software packages and any necessary updates are installed on the node. See [“How to Install Oracle Solaris Cluster Software Packages” on page 59.](#)
 - Ensure that the cluster is prepared for the addition of the new node. See [“How to Prepare the Cluster for Additional Global-Cluster Nodes” on page 114.](#)
1. **Ensure that the Oracle Solaris Cluster software is not yet configured on the potential node that you want to add to a cluster.**
 - a. **Assume the root role on the potential node.**

b. Determine whether the Oracle Solaris Cluster software is configured on the potential node.

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

■ **If the command fails, go to [Step 2](#).**

The Oracle Solaris Cluster 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, the Oracle Solaris Cluster software is already a 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.

■ **SPARC:**

```
ok boot -x
```

■ **x86:**

i In the GRUB menu, use the arrow keys to select the appropriate Oracle Solaris entry and type `e` to edit its commands.

For more information about GRUB based booting, see [“Booting a System”](#) in [“Booting and Shutting Down Oracle Solaris 11.2 Systems”](#).

ii In the boot parameters screen, use the arrow keys to select the kernel entry and type `e` to edit the entry.

iii Add `-x` to the `multiboot` command to specify that the system boot into noncluster mode.

iv Press Enter to accept the change and return to the boot parameters screen.

The screen displays the edited command.

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 add the `-x` option to the kernel boot parameter command.

d. Unconfigure the Oracle Solaris Cluster software from the potential node.

```
phys-schost-new# /usr/cluster/bin/clnode remove
```

2. If you are duplicating a node that runs the Oracle Solaris Cluster 4.2 software, create a cluster configuration XML file.

a. Assume the root role 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. Assume the root role on the potential node.

4. Ensure that TCP wrappers for RPC are disabled on all nodes of the cluster.

The Oracle Solaris TCP wrappers for RPC feature prevents internode communication that is necessary for cluster configuration.

a. On each node, display the status of TCP wrappers for RPC.

TCP wrappers are enabled if `config/enable_tcpwrappers` is set to `true`, as shown in the following example command output.

```
# svccfg -s rpc/bind listprop config/enable_tcpwrappers
config/enable_tcpwrappers boolean true
```

- b. If TCP wrappers for RPC are enabled on a node, disable TCP wrappers and refresh the RPC bind service.**

```
# svccfg -s rpc/bind setprop config/enable_tcpwrappers = false
# svcadm refresh rpc/bind
# svcadm restart rpc/bind
```

5. Modify or create the cluster configuration XML file as needed.

- 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.

- 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.

6. Validate the cluster configuration XML file.

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

7. Configure the new cluster node.

```
phys-schost-new# clnode add -n sponsor-node -i clconfigfile
```

`-n sponsor-node`

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.

8. If TCP wrappers are used in the cluster, ensure that the `clprivnet0` IP addresses for all added nodes are added to the `/etc/hosts.allow` file on each cluster node.

Without this addition to the `/etc/hosts.allow` file, TCP wrappers prevent internode communication over RPC for cluster administration utilities.

- a. On each node, display the IP addresses for all `clprivnet0` devices.**

```
# /usr/sbin/ipadm show-addr
ADDROBJ      TYPE      STATE      ADDR
clprivnet0/N static    ok         ip-address/netmask-length
...
```

- b. On each node, edit the `/etc/hosts.allow` file with the IP addresses of all `clprivnet0` devices in the cluster.

9. (Optional) Enable automatic node reboot if all monitored shared-disk paths fail.

Note - At initial configuration time, disk-path monitoring is enabled by default for all discovered devices.

- a. Enable automatic reboot.

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

```
-p
```

Specifies the property to set

```
reboot_on_path_failure=enable
```

Enables automatic node reboot if failure of all monitored shared-disk paths occurs.

- 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 perform this procedure again. If that does not correct the problem, perform the procedure [“How to Unconfigure Oracle Solaris Cluster Software to Correct Installation Problems” on page 227](#) on each misconfigured node to remove it from the cluster configuration. You do not need to uninstall the Oracle Solaris Cluster software packages. Then perform this procedure again.

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 Global Cluster” on page 134](#).

Otherwise, go to [“How to Verify the Quorum Configuration and Installation Mode” on page 142](#).

▼ How to Update Quorum Devices After Adding a Node to a Global Cluster

If you added a node to a global cluster, you must update the configuration information of the quorum devices regardless of whether you use shared disks, 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 update 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 the Oracle Solaris Cluster software on the added node.

1. On any node of the cluster, assume the root role.

2. Ensure that all cluster nodes are online.

```
phys-schost# cluster status -t node
```

3. 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
...
```

4. Note the name of each quorum device that is listed.

5. Remove the original quorum device.

Perform this step for each quorum device that is configured.

```
phys-schost# clquorum remove device-name
```

device-name

Specifies the name of the quorum device.

6. 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
```

7. Update the global-devices namespace.

```
phys-schost# cldevice populate
```

Note - This step is necessary to prevent possible node panic.

8. 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
```

9. (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 and choose the shared device from the output.

```
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. Configure the shared device as a quorum device.

```
phys-schost# clquorum add -t type device-name
```

`-t type`

Specifies the type of quorum device. If this option is not specified, the default type `shared_disk` is used.

c. Repeat for each quorum device that you want to configure.

d. 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 142.

▼ How to Configure Quorum Devices

Note - You do not need to configure quorum devices in the following circumstances:

- You chose automatic quorum configuration during Oracle Solaris Cluster software configuration.
- You installed a single-node global cluster.
- You added a node to an existing global cluster and already have sufficient quorum votes assigned.

If you chose automatic quorum configuration when you established the cluster, do not perform this procedure. Instead, proceed to [“How to Verify the Quorum Configuration and Installation Mode” on page 142.](#)

Perform this procedure one time only, after the new cluster is fully formed. Use this procedure to assign quorum votes and then to remove the cluster from installation mode.

Before You Begin

- **Quorum servers** – To configure a quorum server as a quorum device, do the following:
 - Install the Oracle Solaris 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 [“How to Install and Configure Oracle Solaris Cluster Quorum Server Software” on page 55.](#)
 - Ensure that network switches that are directly connected to cluster nodes meet one of the following criteria:
 - The switch supports Rapid Spanning Tree Protocol (RSTP).
 - Fast port mode is enabled on the switch.

One of these features is required to ensure immediate communication between cluster nodes and the quorum server. If this communication is significantly delayed by the switch, the cluster interprets this prevention of communication 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
- **NAS devices** – To configure a network-attached storage (NAS) device as a quorum device, do the following:
 - Install the NAS device hardware and software. See [“Oracle Solaris Cluster With Network-Attached Storage Device Manual ”](#) and your device documentation for requirements and installation procedures for NAS hardware and software.

1. **If both of the following conditions apply, ensure that the correct prefix length is set for the public-network addresses.**

- You intend to use a quorum server.
- The public network uses variable-length subnet masking, also called classless inter domain routing (CIDR).

```
# ipadm show-addr
ADDROBJ      TYPE      STATE      ADDR
lo0/v4       static    ok         127.0.0.1/8
imp0/v4      static    ok         10.134.94.58/24
```

Note - If you use a quorum server but the public network uses classful subnets as defined in RFC 791, you do not need to perform this step.

2. On one node, assume the root role.

Alternatively, if your user account is assigned the System Administrator profile, issue commands as non-root through a profile shell, or prefix the command with the `pfexec` command.

3. Ensure that all cluster nodes are online.

```
phys-schost# cluster status -t node
```

4. To use a shared 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 the root role to run this command.

```
phys-schost-1# 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. Ensure that the output shows all connections between cluster nodes and storage devices.

c. Determine the global device ID 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 `cldevice` output from [Step 4.a](#) to identify the device ID of each shared disk that you are configuring as a quorum device. For example, the output in [Step 4.a](#) shows that global device `d3` is shared by `phys-schost-1` and `phys-schost-2`.

5. To use a shared disk that does not support the SCSI protocol, ensure that fencing is disabled for that shared disk.

a. Display the fencing setting for the individual disk.

```
phys-schost# cldevice show device

=== DID Device Instances ===
DID Device Name:                /dev/did/rdisk/dN
...
default_fencing:                nofencing
...
```

- If fencing for the disk is set to `nofencing` OR `nofencing-noscrub`, fencing is disabled for that disk. Go to [Step 6](#).

- If fencing for the disk is set to `pathcount` OR `scsi`, disable fencing for the disk. Skip to [Step 5.c](#).

- If fencing for the disk is set to `global`, determine whether fencing is also disabled globally. Proceed to [Step 5.b](#).

Alternatively, you can simply disable fencing for the individual disk, which overrides for that disk whatever value the `global_fencing` property is set to. Skip to [Step 5.c](#) to disable fencing for the individual disk.

b. Determine whether fencing is disabled globally.

```
phys-schost# cluster show -t global

=== Cluster ===
Cluster name:                    cluster
...
global_fencing:                 nofencing
...
```

- If global fencing is set to `nofencing` OR `nofencing-noscrub`, fencing is disabled for the shared disk whose `default_fencing` property is set to `global`. Go to [Step 6](#).

- **If global fencing is set to `pathcount` or `prefer3`, disable fencing for the shared disk. Proceed to [Step 5.c](#).**

Note - If an individual disk has its `default_fencing` property set to `global`, the fencing for that individual disk is disabled only while the cluster-wide `global_fencing` property is set to `nofencing` or `nofencing-noscrub`. If the `global_fencing` property is changed to a value that enables fencing, then fencing becomes enabled for all disks whose `default_fencing` property is set to `global`.

c. Disable fencing for the shared disk.

```
phys-schost# cldevice set \  
-p default_fencing=nofencing-noscrub device
```

d. Verify that fencing for the shared disk is now disabled.

```
phys-schost# cldevice show device
```

6. Start the `clsetup` utility.

```
phys-schost# clsetup
```

The Initial Cluster Setup screen is displayed.

Note - If the Main Menu is displayed instead, the initial cluster setup was already successfully performed. Skip to [Step 11](#).

7. Indicate whether you want to add any quorum devices.

- **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 10](#).**
 - **Type `Yes` to configure additional quorum devices.**

8. Specify what type of device you want to configure as a quorum device.

| Quorum Device Type | Description |
|--------------------|--|
| shared_disk | Shared LUNs from the following: <ul style="list-style-type: none"> ■ Shared SCSI disk ■ Serial Attached Technology Attachment (SATA) storage ■ Oracle ZFS Storage Appliance |
| quorum_server | Quorum server |

9. Specify the name of the device to configure as a quorum device and provide any required additional information.

- 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

10. Type `yes` to verify that it is okay to reset `installmode`.

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.

11. 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 142.

Troubleshooting **scinstall fails to perform an automatic configuration** – If `scinstall` fails to automatically configure a shared disk as a quorum device, or If the cluster's `installmode` state is still enabled, you can configure a quorum device and reset `installmode` by using the `clsetup` utility after the `scinstall` processing is completed.

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 adding 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 [“Oracle Solaris Cluster System Administration Guide”](#).

Unreachable quorum device – If you see messages on the cluster nodes that a quorum device is unreachable or if you see failures of cluster nodes with the message `CMM: Unable to acquire the quorum device`, there might be a problem with the quorum device or the path to it. Check that both the quorum device and the path to it are functional.

If the problem persists, use a different quorum device. Or, if you want to use the same quorum device, increase the quorum timeout to a high value, as follows:

Note - For Oracle RAC (Oracle RAC), do not change the default quorum timeout of 25 seconds. In certain split-brain scenarios, a longer timeout period might lead to the failure of Oracle RAC VIP failover, due to the VIP resource timing out. If the quorum device being used is not conforming with the default 25-second timeout, use a different quorum device.

1. **Assume the root role.**
2. **On each cluster node, edit the `/etc/system` file as the root role to set the timeout to a high value.**

The following example sets the timeout to 700 seconds.

```
phys-schost# pfedit /etc/system
...
set cl_haci:qd_acquisition_timer=700
```

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

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

4. **Boot each node back into the cluster.**

Changes to the `/etc/system` file are initialized after the reboot.

▼ How to Verify the Quorum Configuration and Installation Mode

Perform this procedure to verify that the quorum configuration was completed successfully and that cluster installation mode is disabled.

You do not need to be the root role to run these commands.

1. **From any global-cluster 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 and creation 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 143](#).
- If you want to install or modify the NTP configuration file, go to [“Configuring Network Time Protocol \(NTP\)” on page 144](#).
- If you want to install a volume manager, go to [Chapter 4, “Configuring Solaris Volume Manager Software”](#).
- If you want to create cluster file systems, go to [“How to Create Cluster File Systems” on page 165](#).
- To find out how to 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 [“Oracle Solaris Cluster Data Services Planning and Administration Guide”](#).
- When your cluster is fully configured, validate the configuration. Go to [“How to Validate the Cluster” on page 148](#).
- Before you put the cluster into production, make a baseline recording of the cluster configuration for future diagnostic purposes. Go to [“How to Record Diagnostic Data of the Cluster Configuration” on page 148](#).

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 “Oracle Solaris Cluster System Administration Guide”](#).

▼ How to Change Private Hostnames

Perform this task if you do not want to use the default private hostnames, `clusternodenodeID-priv`, that are assigned during Oracle Solaris 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. **Assume the root role on a global-cluster node.**

2. **Start the `clsetup` utility.**

```
phys-schost# clsetup
```

The `clsetup` Main Menu is displayed.

3. **Type the option number for Private Hostnames and press the Return key.**

The Private Hostname Menu is displayed.

4. **Type the option number for Change a Node 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 Update the NTP configuration with the changed private hostnames. Go to [“How to Update NTP After Changing a Private Hostname”](#) on page 147.

Configuring Network Time Protocol (NTP)

This section contains the following procedures:

- [“How to Use Your Own `/etc/inet/ntp.conf` File”](#) on page 145
- [“How to Install NTP After Adding a Node to a Single-Node Cluster”](#) on page 146
- [“How to Update NTP After Changing a Private Hostname”](#) on page 147

▼ How to Use Your Own /etc/inet/ntp.conf File

Note - If you installed your own /etc/inet/ntp.conf file **before** you installed the Oracle Solaris Cluster software, you do not need to perform this procedure. Proceed to [“How to Validate the Cluster” on page 148](#).

1. **Assume the root role on a cluster node.**
2. **Add your /etc/inet/ntp.conf file to each node of the cluster.**
3. **On each node, determine the state of the NTP service.**

```
phys-schost# svcs svc:/network/ntp:default
```

4. **Start the NTP daemon on each node.**
 - **If the NTP service is disabled, enable the service.**

```
phys-schost# svcadm enable svc:/network/ntp:default
```

- **If the NTP service is online, restart the service.**

```
phys-schost# svcadm restart svc:/network/ntp:default
```

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 install a volume manager, go to [Chapter 4, “Configuring Solaris Volume Manager Software”](#) .
- If you want to create cluster file systems, go to [“How to Create Cluster File Systems” on page 165](#).
- To find out how to 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 [“Oracle Solaris Cluster Data Services Planning and Administration Guide”](#) .
- When your cluster is fully configured, validate the configuration. Go to [“How to Validate the Cluster” on page 148](#).
- Before you put the cluster into production, make a baseline recording of the cluster configuration for future diagnostic purposes. Go to [“How to Record Diagnostic Data of the Cluster Configuration” on page 148](#).

▼ How to Install NTP After Adding a Node to a Single-Node Cluster

When you add 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.

1. **Assume the root role on a cluster node.**
2. **Copy the `/etc/inet/ntp.conf` and `/etc/inet/ntp.conf.sc` files from the added node to the original cluster node.**

These files were created on the added node when it was configured with the cluster.

3. **On the original cluster node, create a symbolic link named `/etc/inet/ntp.conf.include` that points to the `/etc/inet/ntp.conf.sc` file.**

```
phys-schost# ln -s /etc/inet/ntp.conf.sc /etc/inet/ntp.conf.include
```

4. **On each node, determine the state of the NTP service.**

```
phys-schost# svcs svc:/network/ntp:default
```

5. **Start the NTP daemon on each node.**

- **If the NTP service is disabled, enable the service.**

```
phys-schost# svcadm enable svc:/network/ntp:default
```

- **If the NTP service is onLine, restart the service.**

```
phys-schost# svcadm restart svc:/network/ntp:default
```

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 install a volume manager, go to [Chapter 4, “Configuring Solaris Volume Manager Software”](#).
- If you want to create cluster file systems, go to [“How to Create Cluster File Systems” on page 165](#).
- To find out how to 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 [“Oracle Solaris Cluster Data Services Planning and Administration Guide”](#).
- When your cluster is fully configured, validate the configuration. Go to [“How to Validate the Cluster” on page 148](#).

- Before you put the cluster into production, make a baseline recording of the cluster configuration for future diagnostic purposes. Go to [“How to Record Diagnostic Data of the Cluster Configuration” on page 148](#).

▼ How to Update NTP After Changing a Private Hostname

1. Assume the `root` role on a cluster node.
2. On each node of the cluster, update the `/etc/inet/ntp.conf.sc` file with the changed private hostname.
3. On each node, determine the state of the NTP service.

```
phys-schost# svcs svc:/network/ntp:default
```

4. Start the NTP daemon on each node.
 - If the NTP service is disabled, enable the service.

```
phys-schost# svcadm enable svc:/network/ntp:default
```

- If the NTP service is online, restart the service.

```
phys-schost# svcadm restart svc:/network/ntp:default
```

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 install a volume manager, go to [Chapter 4, “Configuring Solaris Volume Manager Software”](#).
- If you want to create cluster file systems, go to [“How to Create Cluster File Systems” on page 165](#).
- To find out how to 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 [“Oracle Solaris Cluster Data Services Planning and Administration Guide”](#).
- When your cluster is fully configured, validate the configuration. Go to [“How to Validate the Cluster” on page 148](#).
- Before you put the cluster into production, make a baseline recording of the cluster configuration for future diagnostic purposes. Go to [“How to Record Diagnostic Data of the Cluster Configuration” on page 148](#).

▼ How to Validate the Cluster

After you complete all configuration of the cluster, use the `cluster check` command to validate the cluster configuration and functionality. For more information, see the [cluster\(1CL\)](#) man page.

Before You Begin Ensure that you have completed the installation and configuration of all hardware and software components in the cluster, including firmware and software updates.

1. **Assume the root role on a node of the cluster.**

2. **Perform validation checks.**

Follow procedures in “[How to Validate a Basic Cluster Configuration](#)” in “[Oracle Solaris Cluster System Administration Guide](#)”.

Next Steps Before you put the cluster into production, make a baseline recording of the cluster configuration for future diagnostic purposes. Go to “[How to Record Diagnostic Data of the Cluster Configuration](#)” on page 148.

▼ How to Record Diagnostic Data of the Cluster Configuration

After you finish configuring the global cluster but before you put it into production, use the Oracle Explorer utility to record baseline information about the cluster. This data can be used if you need to troubleshoot the cluster in the future.

1. **Assume the root role.**

2. **Install the Oracle Explorer software if it is not already installed.**

The Services Tools Bundle contains the Oracle Explorer packages `SUNWexplo` and `SUNWexplu`. See <http://www.oracle.com/us/support/systems/premier/services-tools-bundle-sun-systems-163717.html> for software download and installation information.

3. **Run the explorer utility on each node in the cluster.**

Use the appropriate command for your platform. For example, to collect information on a Sun Fire T1000 server from Oracle, run the following command:

```
# explorer -i -w default,Tx000
```

For more information, see the `explorer(1M)` man page in the `/opt/SUNWexplo/man/man1m/` directory and *Oracle Explorer Data Collector User Guide* which is available through Note 1153444.1 on [My Oracle Support](https://support.oracle.com) (<https://support.oracle.com>).

The explorer output file is saved in the `/opt/SUNWexplo/output/` directory as `explorer.hostid.hostname-date.tar.gz`.

4. **Save the files to a location that you can access if the entire cluster is down.**
5. **Send all explorer files to the Oracle Explorer database for your geographic location.**

Follow the procedures in *Oracle Explorer Data Collector User's Guide* to use FTP or HTTPS to submit Oracle Explorer files.

The Oracle Explorer database makes your explorer output available to Oracle technical support if the data is needed to help diagnose a technical problem with your cluster.

◆◆◆ CHAPTER 4

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 41](#). See your Solaris Volume Manager documentation for additional details.

This chapter contains the following sections:

- [“Configuring Solaris Volume Manager Software” on page 151](#)
- [“Creating Disk Sets in a Cluster” on page 153](#)
- [“Configuring Dual-String Mediators” on page 161](#)

Configuring Solaris Volume Manager Software

The following table lists the tasks that you perform to configure Solaris Volume Manager software for Oracle Solaris Cluster configurations. Complete the procedures in the order that is indicated.

TABLE 4-1 Task Map: Configuring Solaris Volume Manager Software

| Task | Instructions |
|---|---|
| Plan the layout of your Solaris Volume Manager configuration. | “Planning Volume Management” on page 41 |
| Install the Solaris Volume Manager software. | “How to Install Solaris Volume Manager” on page 151 |
| Create state database replicas on the local disks. | “How to Create State Database Replicas” on page 152 |

▼ How to Install Solaris Volume Manager

Perform this procedure on each node in the global cluster.

1. **Assume the root role.**
2. **Ensure that the `solaris` publisher is valid.**

```
# pkg publisher
PUBLISHER          TYPE    STATUS  URI
solaris            origin  online  solaris-repository
```

For information about setting the `solaris` publisher, see [“Copying and Creating Package Repositories in Oracle Solaris 11.2”](#).

3. Install the Solaris Volume Manager package.

```
# pkg install system/svm
```

▼ How to Create State Database Replicas

Perform this procedure on each node in the global cluster.

1. Assume the root role.

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 [“Creating State Database Replicas”](#) in [“Solaris Volume Manager Administration Guide”](#) 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 Go to [“Creating Disk Sets in a Cluster” on page 153](#) to create Solaris Volume Manager disk sets.

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 an Oracle Solaris Cluster environment, the disk set is automatically registered with the Oracle Solaris 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.

The following table lists the tasks that you perform to create disk sets. Complete the procedures in the order that is indicated.

TABLE 4-2 Task Map: Configuring Solaris Volume Manager Disk Sets

| Task | Instructions |
|---|---|
| Create disk sets by using the <code>metaset</code> command. | “How to Create a Disk Set” on page 153 |
| Add drives to the disk sets. | “How to Add Drives to a Disk Set” on page 156 |
| <i>(Optional)</i> Repartition drives in a disk set to allocate space to different slices. | “How to Repartition Drives in a Disk Set” on page 157 |
| List DID pseudo-driver mappings and define volumes in the <code>/etc/lvm/md.tab</code> files. | “How to Create an <code>md.tab</code> File” on page 158 |
| Initialize the <code>md.tab</code> files. | “How to Activate Volumes” on page 160 |

▼ How to Create a Disk Set

Before You Begin The disk set that you intend to create must meet 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 two or three mediator hosts. These mediator hosts must include the two hosts attached to the enclosures containing the disk set. See [“Configuring Dual-String Mediators” on page 161](#) 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 $\text{count}(S1) + \text{count}(S2) > \text{count}(S3)$.

1. **On each node in the cluster, run the `devfsadm` command.**

You can run this command on all nodes in the cluster at the same time.

2. From one node of the cluster, update the global-devices namespace.

```
phys-schost# cldevice populate
```

See the [cldevice\(1CL\)](#) man page for more information.

3. 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
```

4. Ensure that the local state database replicas exist.

For instructions, see [“How to Create State Database Replicas”](#) on page 152.

5. Assume the root role on the cluster node that will master the disk set.

6. Create the disk set.

The following command creates the disk set and registers the disk set as an Oracle Solaris 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 [“Oracle Solaris Cluster System Administration Guide”](#) for more information about how to change the `numsecondaries` property.

7. **If you are configuring a replicated 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 [“Oracle Solaris Cluster System Administration Guide”](#).

8. **Verify the status of the new disk set.**

```
phys-schost# metaset -s setname
```

9. **As needed, set device group properties.**

```
phys-schost# cldevicegroup set -p name=value device-group
```

-p

Specifies a device-group property.

name

Specifies the name of a property.

value

Specifies the value or setting of the property.

device-group

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 4-2 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 155.

Adding Drives to a Disk Set

When you add a drive to a disk set, the volume management software repartitions the drive 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 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 153](#).

1. **Assume the root role.**
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# cldevicegroup 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 4-3 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 If you want to repartition drives for use in volumes, go to [“How to Repartition Drives in a Disk Set” on page 157](#).

Otherwise, go to [“How to Create an `md.tab` File” on page 158](#) to find out how 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 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 EFI slices 1 through 5, you can use these slices when you set up Solaris Volume Manager volumes.

1. Assume the root role.

2. Use the `format` command to change the disk partitioning for each drive in the disk set.

When you repartition a drive, take steps to prevent the `metaset` command from repartitioning the drive.

a. Create slice 6 for EFI starting at cylinder 0, large enough to hold a state database replica.

Do not allow the target slice to overlap any other slice on the drive.

See your Solaris Volume Manager administration guide to determine the size of a state database replica for your version of the volume-manager software.

b. Set the `Flag` field in the target slice to `wu` (read-write, unmountable).

Do not set this field to read-only.

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 158.](#)

▼ 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 IDs that are used to form disk sets. For example, if the device ID `/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. Assume the root role.

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 (`cNtXdY`). 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 that contains the volume definitions for the disk sets you created.

See [Example 4-4](#) for a sample `md.tab` file.

Note - If you have existing data on the drives that will be used for the submirrors, you must back up the data before 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`. 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 4-4 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/d10 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
```

3. 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 160.

▼ How to Activate Volumes

Perform this procedure to activate Solaris Volume Manager volumes that are defined in `md.tab` files.

1. **Assume the `root` role.**
2. **Ensure that `md.tab` files are located in the `/etc/lvm` directory.**
3. **Ensure that you have ownership of the disk set on the node where the command will be executed.**
4. **Take ownership of the disk set.**

```
phys-schost# cldevicegroup switch -n node device-group
```

`-n node`

Specifies the node that takes ownership.

`device-group`

Specifies the disk set name.

5. **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.

6. **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.

7. **Check the status of the volumes.**

```
phys-schost# metastat -s setname
```

See the [metastat\(1M\)](#) man page for more information.

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 [“Oracle Solaris Cluster System Administration Guide ”](#).

Example 4-5 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 161.

Otherwise, go to [“How to Create Cluster File Systems”](#) on page 165 to find out how to create a cluster file system.

Configuring Dual-String Mediators

This section provides information and procedures to configure dual-string mediator hosts.

A single *disk string* consists of a disk enclosure, its physical drives, cables from the enclosure to the node or nodes, and the interface adapter cards. A dual-string disk set includes disks in two disk strings, and is attached to exactly two nodes. If a single disk string in a dual-string disk set fails such that exactly half the Solaris Volume Manager replicas remain available, the disk set will stop functioning. Dual-string mediators are therefore required for all Solaris Volume Manager dual-string disk sets. The use of mediators enables the Oracle Solaris 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.

The following table lists the tasks that you perform to configure dual-string mediator hosts. Complete the procedures in the order that is indicated.

TABLE 4-3 Task Map: Configuring Solaris Volume Manager Dual-String Mediators

| Task | Instructions |
|---------------------------------------|---|
| Configure dual-string mediator hosts. | “Requirements for Dual-String Mediators” on page 162
“How to Add Mediator Hosts” on page 162 |

| Task | Instructions |
|---|--|
| Check the status of mediator data and, if necessary, fix bad mediator data. | “How to Check For and Fix Bad Mediator Data” on page 163 |

Requirements for Dual-String Mediators

The following rules apply to dual-string configurations that use mediators.

- Disk sets must be configured with either two or three mediator hosts. Two of these mediator hosts must be the same two cluster nodes that are used for the disk set. The third can be another node in the cluster or a non-clustered host on the cluster's public network, such as a quorum server.
- 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 consist of only 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. **If you will use a third mediator host for a dual-string disk set and that host does not already have disk sets configured, modify the `/etc/group` file and create a dummy disk set.**
 - a. **Add the entry `root` to the `sysadmin` group in the `/etc/group` file.**
 - b. **Create a dummy disk set by using the `metaset` command.**

```
phys-schost-3# metaset -s dummy-diskset-name -a -h hostname
```

```
-s dummy-diskset-net
```

Specifies the dummy disk set name.

```
-a
```

Adds to the disk set.

```
-h hostname
```

Specifies the name of the node.

2. Assume the `root` role on the node that currently masters the disk set to which you intend to add mediator hosts.
3. 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.

```
-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 4-6 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`. If needed, repeat the command a third time for a third mediator host. All commands are run from the node that masters the disk set for which you are adding mediator hosts, in this case `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
phys-schost# metaset -s dg-schost-1 -a -m phys-schost-3
```

Next Steps Check the status of mediator data. Go to [“How to Check For and Fix Bad Mediator Data”](#) on page 163.

▼ How to Check For and Fix Bad Mediator Data

Perform this procedure to repair bad mediator data.

Before You Begin Ensure that you have added mediator hosts as described in [“How to Add Mediator Hosts”](#) on page 162.

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. **Check the Status field of the medstat output for each mediator host.**
 - **If the status is ok, no error condition exists.**
 - **If the status is Bad, repair the affected mediator host.**
3. **Assume the root role on the node that owns the affected disk set.**
4. **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.

5. **Restore each mediator host that you removed in [Step 4](#).**

```
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 want to create cluster file systems, go to [“How to Create Cluster File Systems” on page 165](#).
- To find out how to 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 [“Oracle Solaris Cluster Data Services Planning and Administration Guide”](#).

Creating a Cluster File System

This chapter describes how to create a cluster file system to support data services.

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 [“Oracle Solaris Cluster Data Services Planning and Administration Guide”](#).

The following sections are in this chapter:

- [“Creating Cluster File Systems”](#) on page 165
- [“Creating Oracle ACFS File Systems”](#) on page 168

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 global cluster.

Before You Begin Perform the following tasks:

- Ensure that you installed software packages for the Oracle Solaris OS, Oracle Solaris Cluster framework, and other products as described in [“Installing the Software”](#) on page 46.
- Ensure that you established the new cluster or cluster node as described in [“Establishing a New Global Cluster or New Global-Cluster Node”](#) on page 71.
- If you are using a volume manager, ensure that volume management software is installed and configured.

- Determine the mount options to use for each cluster file system that you want to create. See [“Choosing Mount Options for UFS Cluster File Systems”](#) on page 38.

1. Assume the root role on any node in the cluster.

Tip - For faster file system creation, become the root role on the current primary of the global device for which you create a file system.

2. Create a UFS file system by using the `newfs` command.



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.

```
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 |
| None | <code>/dev/global/rdisk/d1s3</code> | Raw disk device d1s3 |

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/mount-point/
```

device-group

Name of the directory that corresponds to the name of the device group that contains the device.

mount-point

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.

- a. In each entry, specify the required mount options for the type of file system that you use.
- b. To automatically mount the cluster file system, set the `mount at boot` field to `yes`.
- c. For each cluster file system, ensure that 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# cluster check -k vfstab
```

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, no output is returned.

For more information, see the [cluster\(1CL\)](#) man page.

6. Mount the cluster file system from any node in the cluster.

```
phys-schost# mount /global/device-group/mountpoint/
```

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.

Example 5-1 Creating a UFS 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 `cluster check` 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# cluster check -k vfstab
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 To find out how to 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 [“Oracle Solaris Cluster Data Services Planning and Administration Guide”](#).

Creating Oracle ACFS File Systems

This section provides procedures to create an Oracle Automatic Storage Management Cluster File System (Oracle ACFS) file system to support data services. This file system can be used as a general-purpose file system or as the database home for an Oracle database. An Oracle ACFS file system is supported for use in the global cluster and in zone clusters.

Note - A minimum of Oracle ASM version 11g release 2 is required.

The following table lists the tasks to perform to create an Oracle ACFS file system. Complete the procedures in the order that is indicated.

TABLE 5-1 Task Map: Creating Oracle ACFS File Systems

| Task | Instructions |
|---|---|
| Review Oracle ACFS configuration diagrams. | “Sample Configurations of an Oracle ACFS File System” on page 169 |
| Register and configure the framework resource groups. | “How to Register and Configure the Framework Resource Groups” on page 173 |
| Create an Oracle ACFS file system. | “How to Create an Oracle ACFS File System” on page 176 |

| Task | Instructions |
|--|---|
| Register and configure the scalable device-group resource group. | “How to Register and Configure the Scalable Device-Group Resource Group” on page 178 |
| Register and configure the Oracle ASM resource groups. | “How to Register and Configure the Oracle ASM Resource Group” on page 179 |
| Configure interoperability between Oracle Grid Infrastructure and Oracle Solaris Cluster | “How to Create an Oracle Grid Infrastructure Resource for Interoperation With Oracle Solaris Cluster” on page 183 |
| Register and configure the Oracle ACFS proxy resource group. | “How to Register and Configure the Oracle ACFS Proxy Resource Group” on page 182 |

Sample Configurations of an Oracle ACFS File System

The following diagrams illustrate the configuration of an Oracle ACFS file system in the global zone and in a zone cluster, using Solaris Volume Manager or hardware RAID.

FIGURE 5-1 Configuration of Oracle ACFS in the Global Zone With Solaris Volume Manager

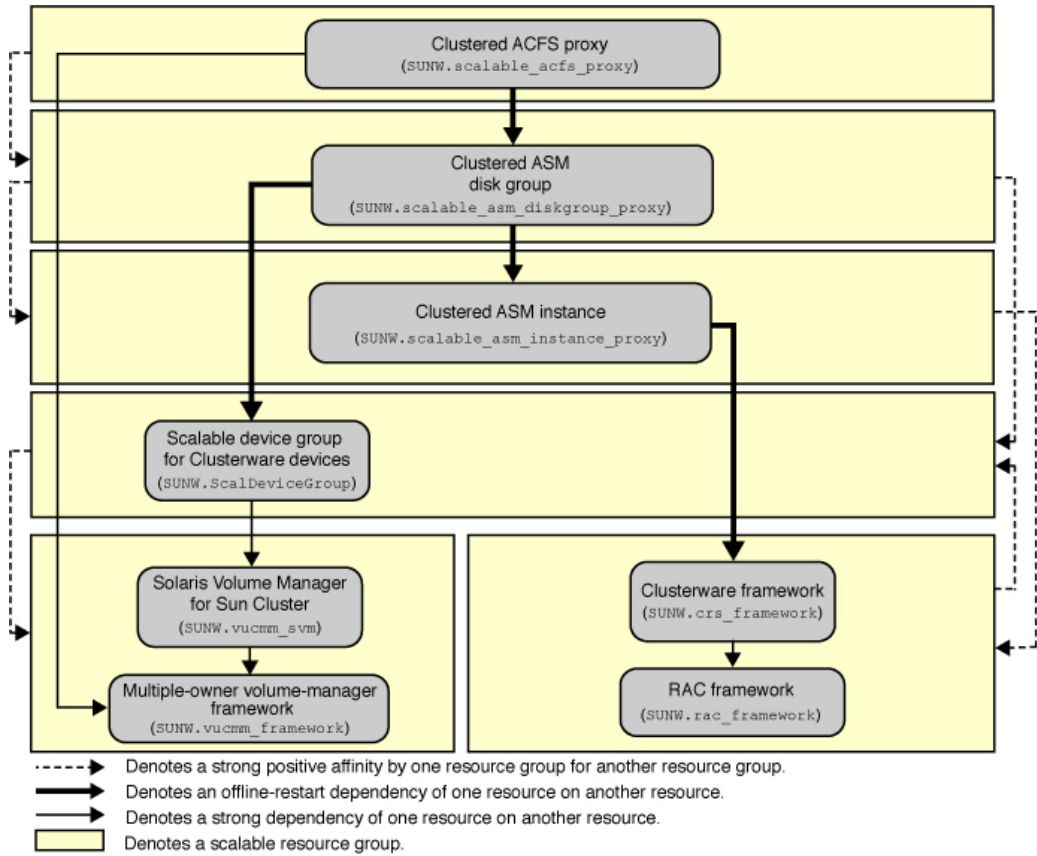


FIGURE 5-2 Configuration of Oracle ACFS in the Global Zone With Hardware RAID

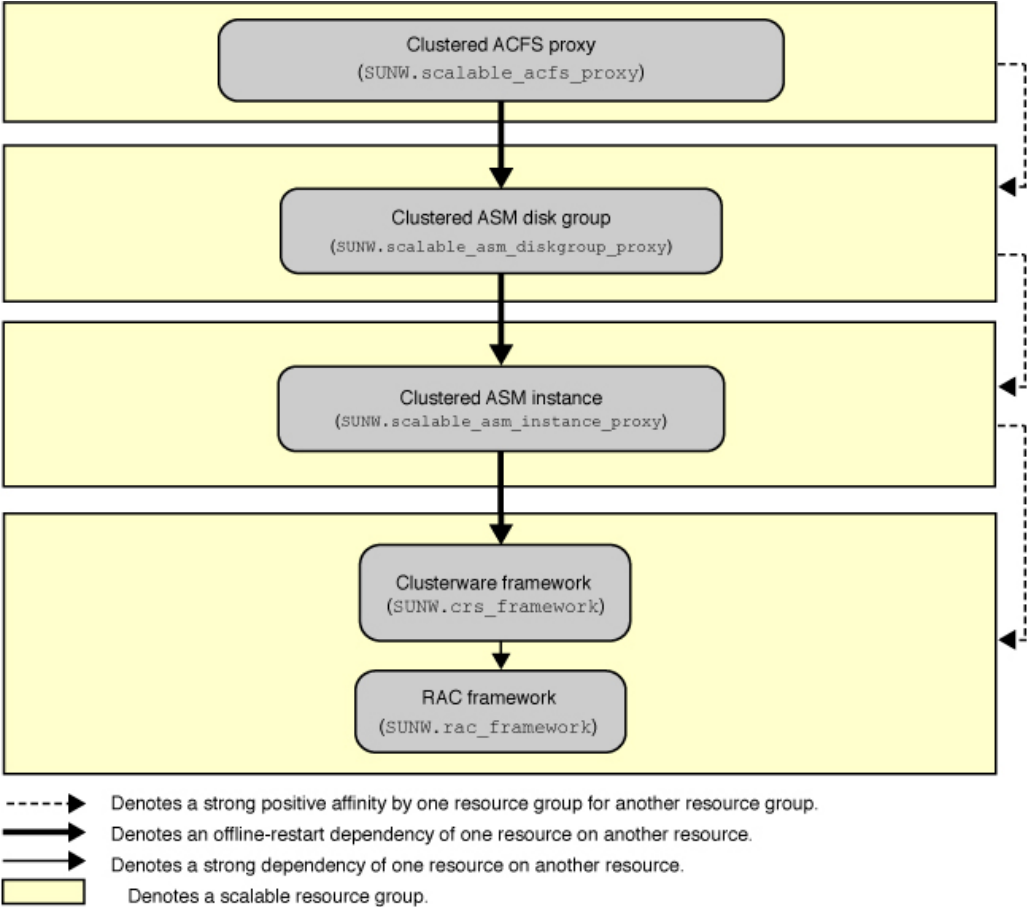


FIGURE 5-3 Configuration of Oracle ACFS in a Zone Cluster with Solaris Volume Manager

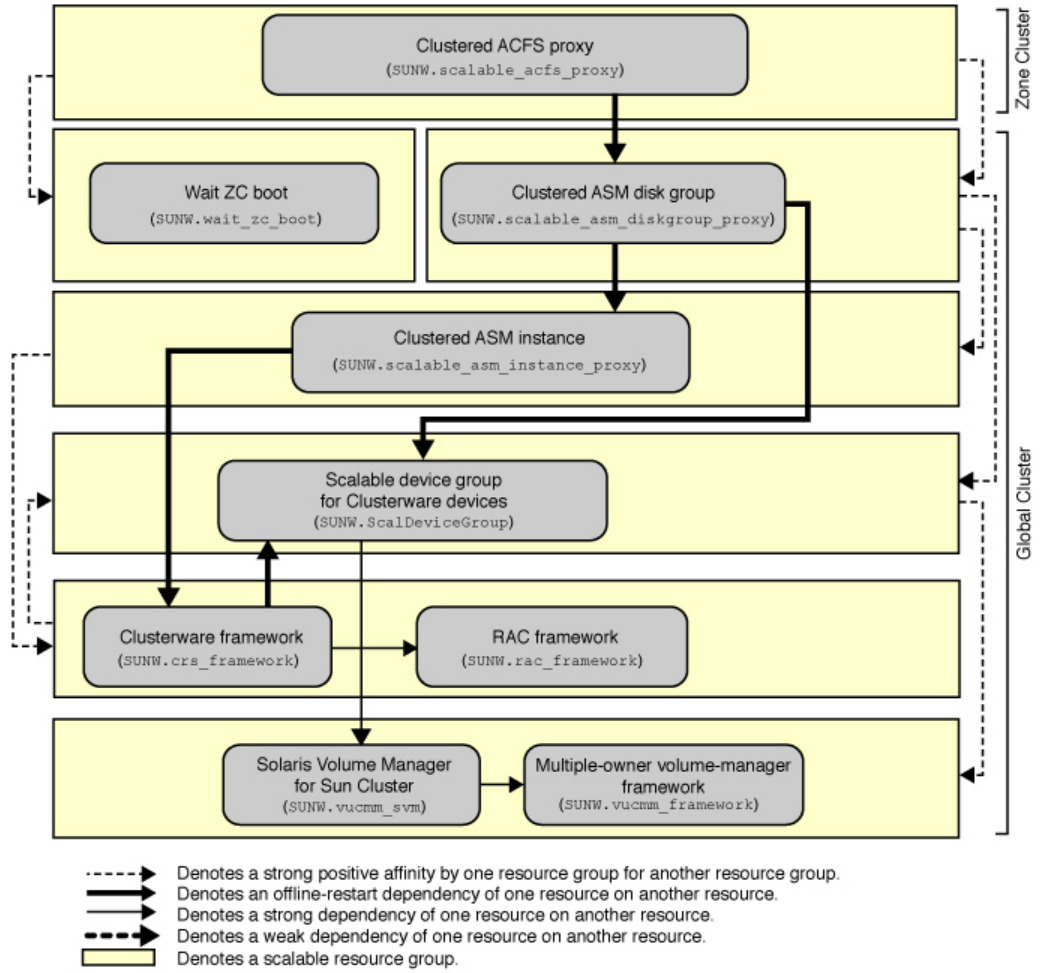
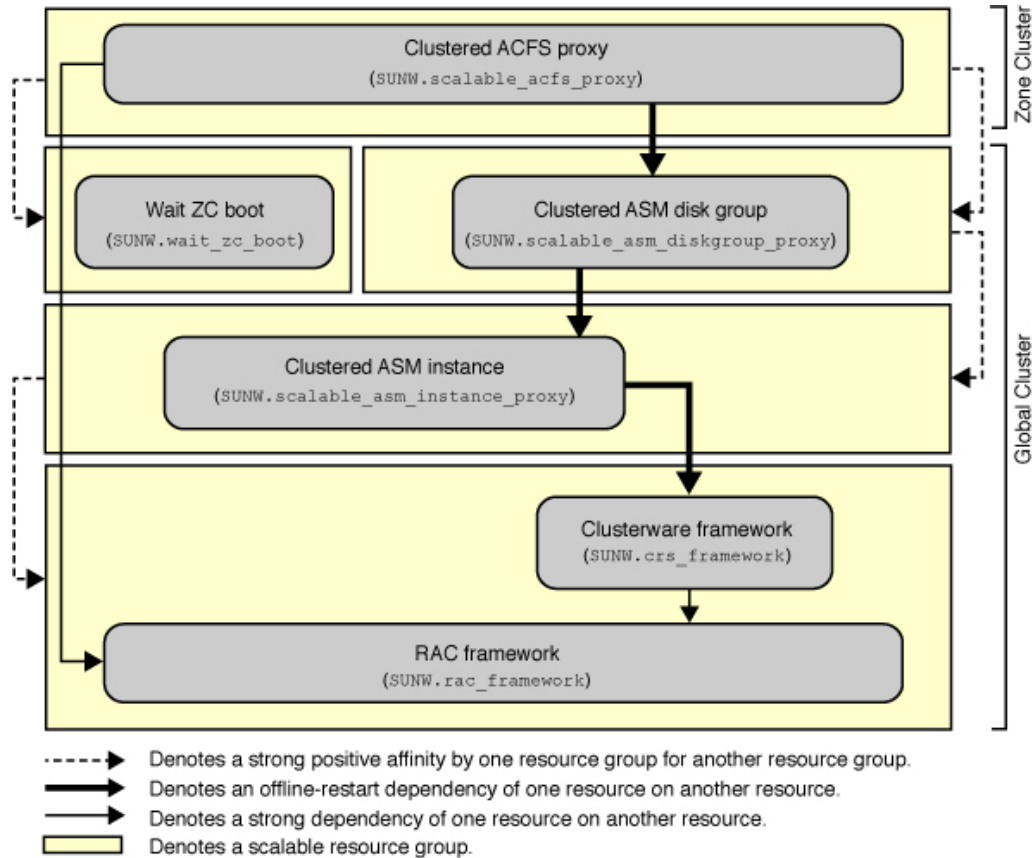


FIGURE 5-4 Configuration of Oracle ACFS in a Zone Cluster with Hardware RAID



▼ How to Register and Configure the Framework Resource Groups

Use this procedure to configure a scalable Oracle RAC framework resource group and, if Solaris Volume Manager for Sun Cluster is used, a multiple-owner volume-manager framework resource group.

Perform all steps from one node of the global cluster.

1. **Assume the root role or become an administrator that provides `solaris.cluster.admin` and `solaris.cluster.modify` authorizations.**

2. **Create a scalable Oracle RAC framework resource group.**

This framework resource group is used regardless of whether the file system is for general purpose or for a database home.

- a. **Create the Oracle RAC framework resource group.**

```
# clresourcegroup create -S \  
-p rg_description="description"] \  
rac-fmwk-rg
```

```
-p rg_description="description"
```

Specifies an optional brief description of the resource group. This description is displayed when you use Oracle Solaris Cluster maintenance commands to obtain information about the resource group.

```
rac-fmwk-rg
```

Specifies the name that you are assigning to the Oracle RAC framework resource group.

- b. **Register the `SUNW.rac_framework` resource type.**

```
# clresourcetype register SUNW.rac_framework
```

- c. **Add an instance of the `SUNW.rac_framework` resource type to the Oracle Real Application Clusters framework resource group.**

```
# clresource create -g rac-fmwk-rg -t SUNW.rac_framework rac-fmwk-rs
```

```
-g rac-fmwk-rg
```

Specifies the Oracle RAC framework resource group to which you are adding the resource.

```
rac-fmwk-rs
```

Specifies the name that you are assigning to the `SUNW.rac_framework` resource.

- d. **Register the Oracle Clusterware framework resource type**

```
# clresourcetype register SUNW.crs_framework
```

- e. **Add an instance of the `SUNW.crs_framework` resource type to the Oracle Real Application Clusters framework resource group.**

```
# clresource create -g rac-fmwk-rg \  
-p resource_dependencies=rac-fmwk-rs \  
-t SUNW.crs_framework \  

```

crs-fmwk-rs

crs-fmwk-rs Specifies the name that you are assigning to the SUNW.crs_framework resource.

3. If you are using Solaris Volume Manager for Sun Cluster, create a scalable multiple-owner volume-manager framework resource group.

This resource group manages Solaris Volume Manager for Sun Cluster.

If you are using hardware RAID, skip to [Step 4](#).

a. Create the resource group

```
# clresourcegroup create -n nodelist -S vucmm-fmwk-rg
```

-n nodelist=nodelist

Specifies the same node list that you configured for the scalable Oracle Real Application Clusters framework resource group.

vucmm-fmwk-rg

Specifies the name that you are assigning to the multiple-owner volume-manager framework resource group.

b. Register the SUNW.vucmm_framework resource type.

```
# clresourcetype register SUNW.vucmm_framework
```

c. Add an instance of the SUNW.vucmm_framework resource type to the multiple-owner volume-manager resource group.

```
# clresource create -g vucmm-fmwk-rg -t SUNW.vucmm_framework vucmm-fmwk-rs
```

-g vucmm-fmwk-rg Specifies the multiple-owner volume-manager resource group to which you are adding the resource.

vucmm-fmwk-rs Specifies the name that you are assigning to the SUNW.vucmm_framework resource.

d. Register the SUNW.vucmm_svm resource type.

```
# clresourcetype register SUNW.vucmm_svm
```

e. Add an instance of the SUNW.vucmm_svm resource type to the SUNW.vucmm_framework resource group.

Ensure that this instance depends on the SUNW.vucmm_framework resource that you created.

```
# clresource create -g vucmm-fmwk-rg \  
-t SUNW.vucmm_svm \  
-p resource_dependencies=vucmm-fmwk-rs svm-rs
```

-g vucmm-fmwk-rg

Specifies the multiple-owner volume-manager resource group.

-p resource_dependencies=vucmm-fmwk-rs

Specifies that this instance depends on the SUNW.vucmm_framework resource.

svm-rs

Specifies the name that you are assigning to the SUNW.vucmm_svm resource.

4. **Bring online and in a managed state the Oracle RAC framework resource group, the multiple-owner volume-manager framework resource group if used, and their resources.**

```
# clresourcegroup online -emM rac-fmwk-rg vucmm-fmwk-rg]
```

rac-fmwk-rg

Specifies that the Oracle RAC framework resource group is to be moved to the MANAGED state and brought online.

vucmm-fmwk-rg

Specifies that the multiple-owner volume-manager framework resource group is to be moved to the MANAGED state and brought online.

Next Steps Go to [“How to Create an Oracle ACFS File System” on page 176.](#)

▼ How to Create an Oracle ACFS File System

Use this procedure to create an Oracle ACFS file system. Perform all steps from the one node of the global cluster.

- Before You Begin**
- Ensure that the resource framework groups are configured. See [“How to Register and Configure the Framework Resource Groups” on page 173.](#)
 - Observe the following guidelines and restrictions for configuring an Oracle ACFS file system in an Oracle Solaris Cluster configuration.
 - Ensure that a minimum of Oracle ASM version 11g version 2 is installed.
 - An Oracle ACFS file system is supported in a global cluster and a zone cluster, but not in individual non-global zones.
 - An Oracle ACFS file system must be managed by an Oracle Clusterware resource.

1. **Create an Oracle ACFS file system.**

Follow procedures in “Creating an Oracle ACFS File System” in *Oracle Automatic Storage Management Administrator's Guide*.

Observe the following special instructions:

- Do not register the Oracle ACFS file-system mount point with the Oracle ACFS registry, if you will use the file system as a general purpose file system. Only register the mount point with the Oracle ACFS registry if the file system will be used as the database home.
- Configure the Oracle ACFS file system only in the global zone. To use the file system in a zone cluster, you direct mount the file system in the zone cluster.
- Configure the Oracle ACFS resource on the same nodes on which you configure the clustered Oracle ASM disk-group proxy resource for the Oracle ASM disk group that contains the Oracle ACFS volume.

2. Create a mount point for the Oracle ACFS file system.

Note - For a zone cluster, create the mount point under the zone root path:

```
# mkdir -p /zonepath/root/path-to-filesystem
```

3. For a zone cluster, ensure that the zone cluster is online.

```
# clzonecluster status zonecluster
```

4. Start and mount the Oracle ACFS file system.

```
# /Grid_home/bin/srvctl add filesystem -d /dev/asm/volume-dev-path
# /Grid_home/bin/srvctl start filesystem -d /dev/asm/volume-dev-path
```

5. For a zone cluster, add the file system to the zone cluster.

Perform these steps from the global zone of one node.

a. Add the Oracle ACFS file system to the zone cluster.

```
# clzonecluster configure zonecluster
clzc:zonecluster> add fs
clzc:zonecluster:fs> set dir=mountpoint
clzc:zonecluster:fs> set special=/dev/asm/volume-dev-path
clzc:zonecluster:fs> set type=acfs
clzc:zonecluster:fs> end
clzc:zonecluster> exit
```

b. Verify that the Oracle ACFS file system is added to the zone cluster.

```
# clzonecluster show zonecluster
...
Resource Name:          fs
```

```
dir:                mountpoint
special            volume
raw:
type:              acfs
options:           []
cluster-control:  true
...
```

Next Steps If you are using Solaris Volume Manager for Sun Cluster, go to [“How to Register and Configure the Scalable Device-Group Resource Group”](#) on page 178.

Otherwise, go to [“How to Register and Configure the Oracle ASM Resource Group”](#) on page 179.

▼ How to Register and Configure the Scalable Device-Group Resource Group

If you use Solaris Volume Manager for Sun Cluster, register and configure a scalable device-group resource group. Perform all steps from one node of the global cluster.

If your configuration does not use Solaris Volume Manager for Sun Cluster, do not perform this procedure. Proceed to [“How to Register and Configure the Oracle ASM Resource Group”](#) on page 179.

Before You Begin Ensure that the Oracle ACFS file system is created. See [“How to Create an Oracle ACFS File System”](#) on page 176.

1. **Assume the root role or become an administrator that provides `solaris.cluster.admin` and `solaris.cluster.modify` authorizations.**
2. **Create a scalable resource group to contain the scalable device-group resource.**
Set a strong positive affinity by the resource group for the multiple-owner volume-manager framework resource group.

```
# clresourcegroup create -S \  
-p rg_affinities=++vucmm-fwk-rg \  
-p rg_description="description"] \  
scal-dg-rg
```

3. **Register the `SUNW.ScalDeviceGroup` resource type.**

```
# clresourcetype register SUNW.ScalDeviceGroup
```
4. **Add an instance of the `SUNW.ScalDeviceGroup` resource type to the `SUNW.ScalDeviceGroup` resource group.**

Set a strong dependency for the instance of `SUNW.ScalDeviceGroup` on the `svm-rs` resource in the `SUNW.vucmm_svm` framework resource group. Limit the scope of this dependency to only the node where the `SUNW.ScalDeviceGroup` resource is running.

```
# clresource create -t SUNW.ScalDeviceGroup -g scal-dg-rg \  
-p resource_dependencies=svm-rs{local_node} \  
-p diskgroupname=disk-group scal-dg-rs
```

5. Bring online and in a managed state the scalable device-group resource group.

```
# clresourcegroup online -emM scal-dg-rg
```

6. Set an offline-restart dependency on `crs-fmwk-rs` by `scal-dg-rs`.

```
E clresource set -p resource_dependency_offline_restart=scal-dg-rs crs-fmwk-rs
```

Next Steps Go to [“How to Register and Configure the Oracle ASM Resource Group”](#) on page 179.

▼ How to Register and Configure the Oracle ASM Resource Group

Use this procedure to register and configure the Oracle Automatic Storage Management (Oracle ASM) resource group. Perform all steps from one node of the global cluster.

- Before You Begin**
- Ensure that the framework resource groups are created. See [“How to Register and Configure the Framework Resource Groups”](#) on page 173.
 - If you use Solaris Volume Manager for Sun Cluster, ensure that the scalable device-group resource group is created. See [“How to Register and Configure the Scalable Device-Group Resource Group”](#) on page 178.

1. Assume the root role or become an administrator that provides `solaris.cluster.admin` and `solaris.cluster.modify` authorizations.

2. Register the Oracle ASM resource types for the data service.

a. Register the scalable Oracle ASM instance proxy resource type.

```
# clresourcetype register SUNW.scalable_asm_instance_proxy
```

b. Register the Oracle ASM disk-group resource type.

```
# clresourcetype register SUNW.scalable_asm_diskgroup_proxy
```

3. Create resource groups `asm-inst-rg` and `asm-dg-rg`.

```
# clresourcegroup create -S asm-inst-rg asm-dg-rg
```

asm-inst-rg

Specifies the name of the Oracle ASM instance resource group.

asm-dg-rg

Specifies the name of the Oracle ASM disk-group resource group.

4. Set a strong positive affinity on *rac-fmwk-rg* by *asm-inst-rg*.

```
# clresourcegroup set -p rg_affinities=++rac-fmwk-rg asm-inst-rg
```

5. Set a strong positive affinity by *asm-dg-rg*.

- **If you are using hardware RAID, set the affinity on *asm-inst-rg***

```
# clresourcegroup set -p rg_affinities=++asm-inst-rg asm-dg-rg
```

- **If you are using Solaris Volume Manager for Sun Cluster, set the affinity on *scal-dg-rg* and *asm-inst-rg*.**

```
# clresourcegroup set -p rg_affinities=++asm-inst-rg,++scal-dg-rg asm-dg-rg
```

6. Create a `SUNW.scalable_asm_instance_proxy` resource and set the resource dependencies.

```
# clresource create -g asm-inst-rg \  
-t SUNW.scalable_asm_instance_proxy \  
-p ORACLE_HOME=Grid_home \  
-p CRS_HOME=Grid_home \  
-p "ORACLE_SID{node1}"=+ASM1 \  
-p "ORACLE_SID{node2}"=+ASM2 \  
-p resource_dependencies_offline_restart=crs-fmwk-rs \  
-d asm-inst-rs
```

```
-t SUNW.scalable_asm_instance_proxy
```

Specifies the type of the resource to add.

```
-p ORACLE_HOME=Grid_home
```

Sets the path to the Oracle Grid Infrastructure home directory, where Oracle ASM is installed.

```
-p CRS_HOME=Grid_home
```

Sets the path to the Oracle Grid Infrastructure home directory, where Oracle Clusterware is installed.

```
-p ORACLE_SID=+ASMn
```

Sets the Oracle ASM system identifier.

```
-d asm-inst-rs
```

Specifies the name of the Oracle ASM instance resource that you create.

7. Bring online the *asm-inst-rg* resource group in a managed state on a cluster node.

```
# clresourcegroup online -eM asm-inst-rg
```

8. Add an Oracle ASM disk-group resource to the *asm-dg-rg* resource group.

■ **For hardware RAID, use the following command:**

```
# clresource create -g asm-dg-rg \
-t SUNW.scalable_asm_diskgroup_proxy \
-p asm_diskgroups=dg[,dg...] \
-p resource_dependencies_offline_restart=asm-inst-rs \
-d asm-dg-rs
```

■ **For Solaris Volume Manager for Sun Cluster, use the following command:**

```
# clresource create -g asm-dg-rg \
-t SUNW.scalable_asm_diskgroup_proxy \
-p asm_diskgroups=dg[,dg...] \
-p resource_dependencies_offline_restart=asm-inst-rs,scal-dg-rs \
-d asm-dg-rs
```

9. Bring online the *asm-dg-rg* resource group in a managed state on a cluster node.

```
# clresourcegroup online -eM asm-dg-rg
```

10. For a zone cluster, from the global zone of one node, create a *SUNW.wait_zc_boot* resource group.

If the Oracle ACFS file system will not be used by a zone cluster, omit this step.

```
# clresourcetype register SUNW.wait_zc_boot
# clresourcegroup create -S scal-wait-zc-rg
# clresource create -g scal-wait-zc-rg \
-t SUNW.wait_zc-boot \
-p zcname=zonecluster \
wait-zc-rs
# clresourcegroup online -eM scal-wait-zc-rg
```

Next Steps Go to [“How to Create an Oracle Grid Infrastructure Resource for Interoperation With Oracle Solaris Cluster”](#) on page 183.

▼ How to Register and Configure the Oracle ACFS Proxy Resource Group

Use this procedure to register and configure the Oracle ACFS proxy resource group. Perform all steps from one node of the global cluster.

Before You Begin Ensure that you have registered and configured the Oracle Grid Infrastructure resource. See [“How to Create an Oracle Grid Infrastructure Resource for Interoperation With Oracle Solaris Cluster”](#) on page 183.

- 1. Assume the root role or become an administrator that provides `solaris.cluster.admin` and `solaris.cluster.modify` authorizations.**
- 2. Register the `SUNW.scalable_acfs_proxy` resource type.**
 - **If the file system is for use in the global cluster, use the following command:**

```
# clresource type register SUNW.scalable_acfs_proxy
```
 - **If the file system is for use by a zone cluster, use the following command:**

```
# clresource type register -Z zonecluster SUNW.scalable_acfs_proxy
```
- 3. Create the Oracle ACFS resource group with offline-restart resource dependencies.**
 - **If the file system is for use in the global cluster, use the following command:**

```
# clresourcegroup create -S -p rg_affinities=++asm-dg-rg /  
acfs-rg
```

`asm-dg-rg` Specifies the name of the Oracle ASM disk-group resource group.

`acfs-rg` Specifies the name of the Oracle ACFS resource group.
 - **If the file system is for use by a zone cluster, use the following command:**

```
# clresourcegroup create -Z zonecluster -S  
-p rg_affinities=++global:asm-dg-rg,++global:scal-wait-zc-rg /  
acfs-rg
```

`scal-wait-zc-rg` For a zone cluster, specifies the `SUNW.wait_zc_boot` resource group.
- 4. Add an instance of the `SUNW.scalable_acfs_proxy` resource type to the Oracle ACFS resource group.**
 - **If the file system is for use in the global cluster, use the following command:**

```
# clresource create -g acfs-rg \
-t SUNW.scalable_acfs_proxy \
-p acfs_mountpoint=/acfs-mountpoint \
-p resource_dependencies_offline_restart=asm-dg-rs \
-d acfs-rs
```

- If the file system is for use by a zone cluster, use the following command:

```
# clresource create -Z zonecluster -g acfs-rg \
-t SUNW.scalable_acfs_proxy \
-p acfs_mountpoint=/acfs-mountpoint \
-p resource_dependencies_offline_restart=global:asm-dg-rs \
-p resource_dependencies=global:wait-zc-rs \
-d acfs-rs
```

5. Bring online the *acfs-rg* resource group in a managed state on a cluster node.

```
# clresourcegroup online -eM acfs-rg
```

6. Verify the Oracle ACFS configuration.

```
# clresource status +
```

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 a zone cluster, go to [“Creating and Configuring a Zone Cluster”](#) on page 188.
- 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 [“Oracle Solaris Cluster Data Services Planning and Administration Guide”](#).
- To create an Oracle Grid Infrastructure resource, go to [“How to Create an Oracle Grid Infrastructure Resource for Interoperation With Oracle Solaris Cluster”](#) on page 183.

▼ How to Create an Oracle Grid Infrastructure Resource for Interoperation With Oracle Solaris Cluster

Use this procedure to create an Oracle Grid Infrastructure resource. This resource coordinates operations that are managed by Oracle Clusterware with operations that are managed by Oracle Solaris Cluster.

Before You Begin Ensure that you have registered and configured the Oracle ASM resource groups. See [“How to Register and Configure the Oracle ASM Resource Group”](#) on page 179.

1. Assume the `root` role or become an administrator that provides `solaris.cluster.admin` and `solaris.cluster.modify` authorizations.
2. If the file system will be used by a zone cluster, from the global zone of one node, create an Oracle Clusterware proxy resource.

- a. Create the resource.

```
# /Grid_home/bin/crsctl add type sun.zcboot_proxy.type -basetype local_resource
# /Grid_home/bin/crsctl add res sun.wait-zc-rs \
-type sun.zcboot_proxy.type \
-attr "ACTION_SCRIPT='/opt/SUNWscor/dsconfig/bin/scproxy_crs_action' \
ACL='owner:root:rw,pgroup:oinstall:rw,other::r--' \
SCRIPT_TIMEOUT='20' \
RESTART_ATTEMPTS='60' "
```

- b. Verify the resource.

```
# /Grid_home/bin/crsctl stat res sun.wait-zc-rs -p
NAME=sun.wait-zc-rs
TYPE=sun.zcboot_proxy.type
ACL=owner:root:rw,pgroup:oinstall:rw,other::r-
...
```

- c. Bring the resource online.

```
# /Grid_home/bin/crsctl start res sun.wait-zc-rs
```

3. Create the Oracle Grid Infrastructure `sun.storage_proxy.type` resource type.

Perform this step where you created the `sun.storage_proxy.type` resource type.

```
# /Grid_home/bin/crsctl \
add type sun.storage_proxy.type \
-basetype local_resource \
-attr \
"ATTRIBUTE=ACTION_SCRIPT,TYPE=string", \
"ATTRIBUTE=HOSTING_MEMBERS,TYPE=string", \
"ATTRIBUTE=CARDINALITY,TYPE=string", \
"ATTRIBUTE=PLACEMENT,TYPE=string", \
"ATTRIBUTE=SCRIPT_TIMEOUT,TYPE=int", \
"ATTRIBUTE=RESTART_ATTEMPTS,TYPE=int", \
"ATTRIBUTE=ACL,TYPE=string", \
"ATTRIBUTE=VERSION,TYPE=string"
```

4. Create the Oracle Grid Infrastructure `sun.resource` resource of type `sun.storage_proxy.type`.

Perform this step where you created the `sun.storage_proxy.type` resource type.

The Oracle Grid Infrastructure resource name uses the form `sun.resource`, where `resource` is the name of the `SUNW.ScalDeviceGroup`, `SUNW.ScalMountPoint` or `SUNW.scalable_acfs_proxy` resource.

```
# /Grid_home/bin/crsctl add resource sun.resource \
-type sun.storage_proxy.type \
-attr "ACTION_SCRIPT='/opt/SUNWscor/dsconfig/bin/scproxy_crs_action' \
CARDINALITY='number-nodes' \
SCRIPT_TIMEOUT='timeout' \
PLACEMENT='restricted' \
RESTART_ATTEMPTS='restarts' \
HOSTING_MEMBERS='nodelist' \
VERSION='1' "
```

CARDINALITY The number of nodes in the cluster membership

HOSTING_MEMBERS The list of nodes in the cluster membership

5. Bring online the Oracle Grid Infrastructure `storage_proxy` resource.

```
# /Grid_home/bin/crsctl start resource sun.resource
```

6. Create the Oracle Grid Infrastructure stop trigger for the Oracle Solaris Cluster ACFS proxy resource.

Perform this step from one node of the global cluster.

a. Create the stop trigger resource.

```
# /Grid_home/bin/crsctl add type sun.stoptrigger.type -basetype cluster_resource
# /Grid_home/bin/crsctl add res sun.acfs-rs -type sun.stoptrigger.type \
-attr "action_script='/opt/SUNWscor/dsconfig/bin/crs_stoptrigger_action' \
HOSTING_MEMBERS='node1 node2[...]' \
CARDINALITY='number-nodes' \
placement='restricted' \
ACL='owner:root:rwx,pgrp:oinstall:rwx,other::r--' \
SCRIPT_TIMEOUT='20' \
RESTART_ATTEMPTS='60' \
START_DEPENDENCIES='hard(ora.appdg.ASMvolume.acfs)
pullup:always(ora.appdg.ASMvolume.acfs)' \
STOP_DEPENDENCIES='hard(ora.appdg.ASMvolume.acfs)' "
```

b. Verify the stop trigger resource.

```
# /Grid_home/bin/crsctl stat res sun.acfs-rs -p
NAME=sun.resource
TYPE=sun.stoptrigger.type
...
```

c. Start the stop trigger resource.

```
# /Grid_home/bin/crsctl start res sun.acfs-rs
```

d. Verify that the resource is online on all nodes.

```
# /Grid_home/bin/crsctl stat res sun.acfs-rs
```

Creating Zone Clusters

This chapter provides the following information to create and configure a zone cluster:

- “Overview of Creating and Configuring a Zone Cluster” on page 187
- “Creating and Configuring a Zone Cluster” on page 188

Overview of Creating and Configuring a Zone Cluster

The following task map lists the tasks that you perform to configure a zone cluster. Complete the procedures in the order that is indicated.

TABLE 6-1 Task Map: Create a Zone Cluster

| Task | Instructions |
|--|--|
| Plan the zone cluster configuration. | “Zone Clusters” on page 32 |
| If using Trusted Extensions, install and configure the Trusted Extensions software. | “How to Install and Configure Trusted Extensions” on page 188 |
| Create the zone cluster. | “How to Create a Zone Cluster (clsetup)” on page 190 |
| <i>(Optional)</i> Add a file system for use by multiple nodes of the zone cluster. | “Adding File Systems to a Zone Cluster” on page 203 |
| <i>(Optional)</i> Add a file system dedicated to a single node of the zone cluster. | “Adding Local File Systems to a Specific Zone-Cluster Node” on page 217 |
| <i>(Optional)</i> Add a storage device for use by multiple nodes or a single node of the zone cluster. | “Adding Storage Devices to a Zone Cluster” on page 220 |
| <i>(Optional)</i> Migrate an Oracle Solaris 10 system to a node of a solaris10 brand zone cluster.
Note - The system to migrate must run a minimum of Oracle Solaris Cluster 3.3 patch 145333-15 for SPARC and 145334-15 for x86 before the migration. | “Creating the Image for Directly Migrating Oracle Solaris 10 Systems Into Zones” in “Creating and using Oracle Solaris 10 Zones” |
| <i>(solaris10 brand zone clusters)</i> Configure Oracle Solaris Cluster 3.3 data services. | Applicable data-service procedures for zone clusters. See Oracle Solaris Cluster 3.3 Documentation (http://www.oracle.com/technetwork/documentation/solaris-cluster-33-192999.html) . |
| <i>(Optional)</i> Validate the zone cluster. | “How to Validate a Basic Cluster Configuration” in “Oracle Solaris Cluster System Administration Guide” |

Creating and Configuring a Zone Cluster

This section provides the following information and procedures to create and configure a zone cluster.

- [“Creating a Zone Cluster” on page 188](#)
- [“Adding File Systems to a Zone Cluster” on page 203](#)
- [“Adding Local File Systems to a Specific Zone-Cluster Node” on page 217](#)
- [“Adding Storage Devices to a Zone Cluster” on page 220](#)

Creating a Zone Cluster

This section provides procedures on how to use the `clsetup` utility to create a zone cluster, and add a network address, file system, ZFS storage pool, and storage device to the new zone cluster.

If any node is in noncluster mode, changes that you make are propagated when the node returns to cluster mode. Therefore, you can create a zone cluster even if some global-cluster nodes are in noncluster mode. When those nodes return to cluster mode, the system performs zone-cluster creation tasks on those nodes.

You can alternatively use the `clzonecluster` utility to create and configure a cluster. See the [`clzonecluster\(1CL\)` man page](#) for more information.

Note - You cannot change the zone cluster name after the zone cluster is created.

This section contains the following procedures:

- [“How to Install and Configure Trusted Extensions” on page 188](#)
- [“How to Create a Zone Cluster \(`clsetup`\)” on page 190](#)
- [“How to Configure a Zone Cluster to Use Trusted Extensions” on page 201](#)

▼ How to Install and Configure Trusted Extensions

This procedure prepares the global cluster to use the Trusted Extensions feature of Oracle Solaris with zone clusters. If you do not plan to enable Trusted Extensions, proceed to [“Creating a Zone Cluster” on page 188](#).

Perform this procedure on each node in the global cluster.

Before You Begin Perform the following tasks:

- Ensure that the Oracle Solaris OS is installed to support Oracle Solaris Cluster and Trusted Extensions software. See [“How to Install Oracle Solaris Software” on page 48](#) for more information about installing Oracle Solaris software to meet Oracle Solaris Cluster software requirements.
- If an external name service is used, ensure that an LDAP naming service is configured for use by Trusted Extensions. See [Chapter 5, “Configuring LDAP for Trusted Extensions,” in “Trusted Extensions Configuration and Administration ”](#)
- Review requirements and guidelines for Trusted Extensions in a zone cluster. See [“Guidelines for Trusted Extensions in a Zone Cluster” on page 35.](#)

1. Assume the root role on a node of the global cluster.

2. Install and configure Trusted Extensions software.

Follow procedures in [Chapter 3, “Adding the Trusted Extensions Feature to Oracle Solaris,” in “Trusted Extensions Configuration and Administration ”](#).

3. Disable the Trusted Extensions zoneshare and zoneunshare scripts.

The Trusted Extensions zoneshare and zoneunshare scripts support the ability to export home directories on the system. An Oracle Solaris Cluster configuration does not support this feature.

Disable this feature by replacing each script with a symbolic link to the `/bin/true` utility.

```
phys-schost# ln -s /usr/lib/zones/zoneshare /bin/true
phys-schost# ln -s /usr/lib/zones/zoneunshare /bin/true
```

4. Configure all logical-hostname and shared-IP addresses that are to be used in the zone cluster.

See [“How to Create a Default Trusted Extensions System” in “Trusted Extensions Configuration and Administration ”](#).

5. (Optional) Enable remote login by the LDAP server to the global-cluster node.

a. In the `/etc/default/login` file, comment out the `CONSOLE` entry.

b. Enable remote login.

```
phys-schost# svcadm enable rlogin
```

c. Modify the `/etc/pam.conf` file.

Modify the account management entries by appending a Tab and typing `allow_remote` or `allow_unlabeled` respectively, as shown below.

```
other account requisite pam_roles.so.1      Tab allow_remote
other account required pam_unix_account.so.1 Tab allow_unlabeled
```

6. Modify the `admin_low` template.

- a. **Assign the `admin_low` template to each IP address that does *not* belong to a Trusted Extensions machine that is used by the global zone.**

```
# tncfg -t admin_low
tncfg:admin_low> add host=ip-address1
tncfg:admin_low> add host=ip-address2
...
tncfg:admin_low> exit
```

- b. **Remove the wildcard address `0.0.0.0/32` from the `tncfg` template.**

```
# tncfg -t admin_low remove host=0.0.0.0
```

7. Assign the `cipso` template to each IP address that *does* belong to a Trusted Extensions machine that is used by the global zone.

```
# tncfg -t cipso
tncfg:cipso> add host=ip-address1
tncfg:cipso> add host=ip-address2
...
tncfg:cipso> exit
```

8. Repeat [Step 1](#) through [Step 7](#) on each remaining node of the global cluster.

When all steps are completed on all global-cluster nodes, perform the remaining steps of this procedure on each node of the global cluster.

9. On each global-cluster node, add the IP address of the Trusted Extensions-enabled LDAP server to the `/etc/inet/hosts` file.

The LDAP server is used by the global zone and by the nodes of the zone cluster.

10. (Optional) Make the global-cluster node an LDAP client.

See “[Make the Global Zone an LDAP Client in Trusted Extensions](#)” in “[Trusted Extensions Configuration and Administration](#)”.

11. Add Trusted Extensions users.

See “[Creating Roles and Users in Trusted Extensions](#)” in “[Trusted Extensions Configuration and Administration](#)”.

Next Steps Create the zone cluster. Go to “[Creating a Zone Cluster](#)” on page 188.

▼ How to Create a Zone Cluster (clsetup)

Perform this procedure to create a zone cluster using the `clsetup` utility.

To modify the zone cluster after it is installed, see [“Performing Zone Cluster Administrative Tasks”](#) in [“Oracle Solaris Cluster System Administration Guide ”](#) and the [clzonecluster\(1CL\)](#) man page.

Note - You cannot change the zone cluster name after the zone cluster is created.

- Before You Begin**
- Create a global cluster. See [Chapter 3, “Establishing the Global Cluster”](#).
 - Read the guidelines and requirements for creating a zone cluster. See [“Zone Clusters”](#) on page 32.
 - If the zone cluster will use Trusted Extensions, ensure that you have installed, configured, and enabled Trusted Extensions as described in [“How to Install and Configure Trusted Extensions”](#) on page 188.
 - If the cluster does not have sufficient subnets available to add a zone cluster, you must modify the private IP address range to provide the needed subnets. For more information, see [“How to Change the Private Network Address or Address Range of an Existing Cluster”](#) in [“Oracle Solaris Cluster System Administration Guide ”](#).
 - Have available the following information:
 - The unique name to assign to the zone cluster.

Note - If Trusted Extensions is enabled, the zone cluster name must be the same name as a Trusted Extensions security label that has the security levels that you want to assign to the zone cluster. Create a separate zone cluster for each Trusted Extensions security label that you want to use.

- The zone path that the nodes of the zone cluster will use. For more information, see the description of the zonepath property in [“Resource Types and Properties”](#) in [“Introduction to Oracle Solaris Zones ”](#). By default, whole-root zones are created.
- The name of each node in the global cluster on which to create a zone-cluster node.
- The zone public hostname, or host alias, that you assign to each zone-cluster node.
- If applicable, the public-network IP address that each zone-cluster node uses. Specifying an IP address and NIC for each zone cluster node is required if the zone cluster will be used in a Geographic Edition configuration. Otherwise, this requirement is optional. For more information about this Geographic Edition requirement, see [“Geographic Edition”](#) on page 17.
- If applicable, the name of the public-network IPMP group that each zone-cluster node uses to connect to the public network.

Note - If you do not configure an IP address for each zone cluster node, two things will occur:

- That specific zone cluster will not be able to configure NAS devices for use in the zone cluster. The cluster uses the IP address of the zone cluster node when communicating with the NAS device, so not having an IP address prevents cluster support for fencing NAS devices.
 - The cluster software will activate any Logical Host IP address on any NIC.
-

Tip - While in the `clsetup` utility, you can press the `<` key to return to a previous screen.

You can also use the Oracle Solaris Cluster Manager GUI to create a zone cluster. For GUI log-in instructions, see [“How to Access Oracle Solaris Cluster Manager”](#) in [“Oracle Solaris Cluster System Administration Guide”](#).

1. Assume the `root` role on an active member node of a global cluster.

You perform all steps of this procedure from a node of the global cluster.

2. Ensure that the node of the global cluster is in cluster mode.

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

--- Node Status ---

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

3. Start the `clsetup` utility.

```
phys-schost# clsetup
```

The Main Menu is displayed.

4. Choose the Zone Cluster menu item.

5. Choose the Create a Zone Cluster menu item.

6. Type the name of the zone cluster you want to add.

A zone cluster name can contain ASCII letters (a-z and A-Z), numbers, a dash, or an underscore. The maximum length of the name is 20 characters.

7. Choose the property you want to change.

You can set the following properties:

| Property | Description |
|--|---|
| <code>zonepath=zone-cluster-node-path</code> | Specifies the path to the zone cluster node. For example, <code>/zones/sczone</code> . |
| <code>brand=brand-type</code> | <p>Specifies the <code>solaris</code>, <code>solaris10</code>, or <code>labeled</code> zones brand used in the zone cluster.</p> <p>Note - To use Trusted Extensions, you must use only the <code>labeled</code> brand. To create an exclusive-IP zone cluster, you must use only the <code>solaris</code> brand.</p> <p>To create an exclusive-IP <code>solaris10</code> brand zone cluster, set the properties using the <code>clzonecluster create</code> command as follows:</p> <pre> cz1> set brand=solaris10 cz1> set ip-type=exclusive </pre> |
| <code>ip-type=value</code> | <p>Specifies the type of network IP address used by the zone cluster. Valid <code>ip-type</code> values are <code>shared</code> and <code>exclusive</code>.</p> <p>The maximum number of exclusive-IP zone clusters is constrained by the <code>cluster</code> property <code>num_xip_zoneclusters</code>, which you can set during initial cluster installation. This value has a default of three. For more information, see the cluster(1CL) man page.</p> |
| <code>enable_priv_net=value</code> | <p>When set to <code>true</code>, Oracle Solaris Cluster private network communication is enabled between the nodes of the zone cluster. The Oracle Solaris Cluster private hostnames and IP addresses for the zone cluster nodes are automatically generated by the system. Private network communication is disabled if the value is set to <code>false</code>. The default value is <code>true</code>.</p> <p>When the <code>enable_priv_net</code> property is set to <code>true</code> along with the following properties, private communication occurs in the following ways:</p> <ul style="list-style-type: none"> ■ <code>ip-type=shared</code> – Communication between zone cluster nodes uses the private networks of the global cluster. ■ <code>ip-type=exclusive</code> (<code>solaris</code> brand only) – Communication between zone cluster nodes uses the specified <code>privnet</code> resources. The <code>privnet</code> resources are either Virtual Network Interfaces (VNICs) for the Ethernet type of private network adapters, or InfiniBand (IB) partitions for the IB type of private network adapters. The VNICs or IB partitions are automatically created by the wizard over each private network adapter of the global cluster, and used to configure a zone cluster. <p>The VNICs or IB partitions that the wizard generates use the following naming conventions:</p> <p>For the Ethernet type: <code>private-network-interface-name_zone-cluster-name_vnic0</code>.</p> <p>For the IB type: <code>private-network-interface-name_zone-cluster-name_ibp0</code>.</p> <p>For example, the private network interfaces of the global cluster are <code>net2</code> and <code>net3</code>, and the zone cluster name is <code>zone1</code>. If <code>net2</code> and <code>net3</code> are Ethernet type network interfaces, the two VNICs that are created for the zone cluster will have the names <code>net2_zone1_vnic0</code> and <code>net3_zone1_vnic0</code>.</p> |

| Property | Description |
|----------|--|
| | If net2 and net3 are IB type network interfaces, the two IB partitions created for the zone cluster will have the names net2_zone1_ibp0 and net3_zone1_ibp0. |

8. For a solaris10 brand zone cluster, enter a zone root password.

A root account password is required for a solaris10 brand zone.

9. (Optional) Choose the Zone System Resource Control property that you want to change.

You can set the following properties:

| Property | Description |
|------------------------------|---|
| max-lwps= <i>value</i> | Specifies the maximum number of lightweight processes (LWPs) simultaneously available to this zone cluster. |
| max-shm-memory= <i>value</i> | Specifies the maximum amount of shared memory in GBytes allowed for this zone cluster. |
| max-shm-ids= <i>value</i> | Specifies the maximum number of shared memory IDs allowed for this zone cluster. |
| max-msg-ids= <i>value</i> | Specifies the maximum number of message queue IDs allowed for this zone cluster. |
| max-sem-ids= <i>value</i> | Specifies the maximum number of semaphore IDs allowed for this zone cluster. |
| cpu-shares= <i>value</i> | Specifies the number of Fair Share Scheduler (FSS) shares to allocate to this zone cluster. |
| | |

10. (Optional) Choose the Zone CPU Resource Control property that you want to change.

You can set the following properties:

| Property | Description |
|--------------------------|---|
| scope= <i>scope-type</i> | Specifies whether the ncpus property used in a zone cluster is dedicated-cpu or capped-cpu. |
| ncpus= <i>value</i> | Specifies the limit for the scope type. <ul style="list-style-type: none"> ■ If the scope property is set to dedicated-cpu, the ncpus property sets a limit on the number of CPUs that should be assigned for this zone's exclusive use. The zone will create a pool and processor set when it boots. See the pooladm(1M) and poolcfg(1M) man pages for more information on resource pools. ■ If the scope property is set to capped-cpu, the ncpus property sets a limit on the amount of CPU time that can be used by a zone cluster. The unit used translates to the percentage of a single CPU that can be used |

| Property | Description |
|----------|---|
| | by all user threads in a zone, expressed as a fraction (for example, .75) or a mixed number (whole number and fraction, for example, 1.25). An <code>ncpus</code> value of 1 means 100% of a CPU. See the pooladm(1M) , pooladm(1M) , and poolcfg(1M) man pages for more information on resource pools. |

11. (Optional) Choose the capped-memory property that you want to change.

You can set the following properties:

| Property | Description |
|-----------------------------|--|
| <code>physical=value</code> | Specifies the GByte limit for physical memory. |
| <code>swap=value</code> | Specifies the GByte limit for swap memory. |
| <code>locked=value</code> | Specifies the GByte limit for locked memory. |

You can also use the Oracle Solaris Cluster Manager GUI to view the capped-cpu memory configuration of a zone cluster, as well as the dedicated-CPU configuration. For GUI log-in instructions, see “[How to Access Oracle Solaris Cluster Manager](#)” in “[Oracle Solaris Cluster System Administration Guide](#)”.

12. Choose a physical host from the list of available physical hosts.

You can select one or all of the available physical nodes (or hosts), and then configure one zone-cluster node at a time.

You can set the following properties:

| Property | Description |
|---|--|
| <code>hostname=hostname</code> | Specifies the zone-cluster node hostname. For example, <code>zc-host-1</code> . |
| <code>address=public-network-address</code> | Specifies the public network address for the zone-cluster node on a shared-IP type zone cluster. For example, <code>172.1.1.1</code> . |
| <code>physical=physical-interface</code> | Specifies a network physical interface for the public network from the available network interfaces that are discovered on the physical nodes. For example, <code>sc_igmp0</code> or <code>net0</code> . |
| <code>defrouter=default-router</code> | Specifies the default router for the network address, if your zone is configured in a different subnet. Each zone or set of zones that uses a different <code>defrouter</code> setting must be on a different subnet, for example, <code>192.168.0.1</code> . See the zonecfg(1M) man page for more information about the <code>defrouter</code> property. |

13. Specify the network addresses for the zone cluster.

The network addresses can be used to configure a logical hostname or shared IP cluster resources in the zone cluster. The network address is in the zone cluster global scope.

- 14. At the Review Configuration screen, press Return to continue and then type c to create the zone cluster.**

The results of your configuration change are displayed, similar to the following:

```
>>> Result of the Creation for the Zone Cluster(sczone) <<<
```

The zone cluster is being created with the following configuration

```
/usr/cluster/bin/clzonecluster configure sczone
create
set brand=solaris
set zonepath=/zones/sczone
set ip-type=shared
set enable_priv_net=true
add capped-memory
set physical=2G
end
add node
set physical-host=phys-schost-1
set hostname=zc-host-1
add net
set address=172.1.1.1
set physical=net0
end
add net
set address=172.1.1.2
end
```

Zone cluster, zc2 has been created and configured successfully.

Continue to install the zone cluster(yes/no) ?

- 15. Type yes to continue.**

The clsetup utility performs a standard installation of a zone cluster and you cannot specify any options.

- 16. When finished, exit the clsetup utility.**

- 17. Verify the zone cluster configuration.**

The verify subcommand checks for the availability of the specified resources. If the clzonecluster verify command succeeds, no output is displayed.

```
phys-schost-1# clzonecluster verify zone-cluster-name
phys-schost-1# clzonecluster status zone-cluster-name
=== Zone Clusters ===
```

```
--- Zone Cluster Status ---
```

| Name | Node Name | Zone HostName | Status | Zone Status |
|------|-----------|---------------|--------|-------------|
| ---- | ----- | ----- | ----- | ----- |

```
zone      basenode1  zone-1      Offline  Configured
          basenode2  zone-2      Offline  Configured
```

18. For Trusted Extensions, make the password files writable on each zone-cluster node.

From the global zone, launch the txzonemgr GUI.

```
phys-schost# txzonemgr
```

Select the global zone, then select the item, Configure per-zone name service.

19. Install the zone cluster.

```
phys-schost-1# clzonecluster install options zone-cluster-name
Waiting for zone install commands to complete on all the nodes
of the zone cluster "zone-cluster-name"...
```

- **For a solaris or labeled brand zone cluster, the following options are valid.**

| Option | Description |
|------------------------------|---|
| -c <i>config-profile.xml</i> | Includes system configuration information. The -c <i>config-profile.xml</i> option provides a configuration profile for all non-global zones of the zone cluster. Using this option changes only the hostname of the zone, which is unique for each zone in the zone cluster. All profiles must have a .xml extension. |
| -M <i>manifest.xml</i> | Specifies a custom Automated Installer manifest that you configure to install the necessary packages on all zone-cluster nodes. Use this option if the base global-cluster nodes for the zone-cluster are not all installed with the same Oracle Solaris Cluster packages but you do not want to change which packages are on the base nodes. If the <code>clzonecluster install</code> command is run without the -M option, zone-cluster installation fails on a base node if it is missing a package that is installed on the issuing base node. |

- **For a solaris10 brand zone cluster, the following options are valid.**

Use either the -a or -d option to install Geographic Edition software, core packages, and agents that are supported in the zone cluster. Also, you can use the -d and -p options in the same command.

Note - For a list of agents that are currently supported in a solaris10 brand zone cluster, see [Oracle Solaris Cluster 4 Compatibility Guide \(http://www.oracle.com/technetwork/server-storage/solaris-cluster/overview/solariscluster4-compatibilityguide-1429037.pdf\)](http://www.oracle.com/technetwork/server-storage/solaris-cluster/overview/solariscluster4-compatibilityguide-1429037.pdf).

| Option | Description |
|--|--|
| Oracle Solaris Cluster patch 145333-15 for SPARC and 145334-15 for x86 | <p>Note - These patches are only required when you are installing the zone cluster with either the Oracle Solaris Cluster 3.3 software or the Oracle Solaris Cluster 3.3 5/11 software.</p> <p>You must install a minimum of Oracle Solaris Cluster 3.3 patch 145333-15 for SPARC or 145334-15 for x86 before you install the solaris10 brand zone cluster. Log in to My Oracle Support to retrieve the patch. Then from the global zone, use the -p option to install the patch. The patchdir directory is required, and must be accessible from inside the solaris10 brand zone on all nodes of the zone cluster.</p> <pre># clzonecluster install-cluster \ -p patchdir=patchdir[,patchlistfile=filename] \ [-n phys-schost-1[,...]] \ [-v] \ zone-cluster-name</pre> <p>For additional instructions on installing patches, log in to My Oracle Support (https://support.oracle.com) and search for ID 1278636. 1, <i>How to Find and Download any Revision of a Solaris Patch</i>.</p> |
| -a <i>absolute_path_to_archive zone-cluster-name</i> | <p>Specifies the absolute path to an image archive to be used as the source image.</p> <pre># clzonecluster install \ [-n nodename] \ -a absolute_path_to_archive \ zone-cluster-name</pre> |
| -d <i>dvd-image zone-cluster-name</i> | <p>Specifies the full directory path to the root directory of an installed solaris10 non-global zone. The cluster software DVD directory must be accessible from the global zone of the node where you run the command.</p> <pre># clzonecluster install-cluster \ -d dvd-image \ zoneclustername</pre> |

For more information, see the [clzonecluster\(1CL\)](#) man page.

20. If you did not use the -c *config-profile.xml* option when you installed the zone cluster, perform sysid configuration.

Otherwise, skip to [Step 21](#).

Note - In the following steps, the non-global zone *znode* and *zone-cluster-name* share the same name.

- **For an exclusive-IP labeled brand zone cluster , perform the following steps.**

Configure only *one* zone-cluster node at a time.

a. Boot the non-global zone of one zone-cluster node.

```
phys-schost# zoneadm -z zcnode boot
```

b. Unconfigure the Oracle Solaris instance and reboot the zone.

```
phys-schost# zlogin zcnode
zcnode# sysconfig unconfigure
zcnode# reboot
```

The `zlogin` session terminates during the reboot.

c. Issue the `zlogin` command and progress through the interactive screens.

```
phys-schost# zlogin -C zcnode
```

d. When finished, exit the zone console.

For information about methods to exit from a non-global zone, see [“How to Exit a Non-Global Zone”](#) in [“Creating and Using Oracle Solaris Zones”](#).

e. From the global zone, halt the zone-cluster node.

```
phys-schost# zoneadm -z zcnode halt
```

f. Repeat the preceding steps for each remaining zone-cluster node.

■ **For a shared-IP labeled brand zone cluster, perform the following steps on each zone-cluster node.**

a. From one global-cluster node, boot the zone cluster.

```
phys-schost# clzonecluster boot zone-cluster-name
```

b. Unconfigure the Oracle Solaris instance and reboot the zone.

```
phys-schost# zlogin zcnode
zcnode# sysconfig unconfigure
zcnode# reboot
```

The `zlogin` session terminates during the reboot.

c. Issue the `zlogin` command and progress through the interactive screens.

```
phys-schost# zlogin -C znode
```

d. When finished, exit the zone console.

For information about methods to exit from a non-global zone, see [“How to Exit a Non-Global Zone”](#) in [“Creating and Using Oracle Solaris Zones”](#).

e. Repeat [Step 20.2.b](#) through [Step 20.2.d](#) for each remaining zone-cluster node.

■ **For a solaris or solaris10 brand zone cluster, perform the following steps on each zone-cluster node.**

a. From one global-cluster node, boot the zone cluster.

```
phys-schost# clzonecluster boot zone-cluster-name
```

b. Issue the zlogin command and progress through the interactive screens.

```
phys-schost# zlogin -C znode
```

c. When finished, exit the zone console.

For information about methods to exit from a non-global zone, see [“How to Exit a Non-Global Zone”](#) in [“Creating and Using Oracle Solaris Zones”](#).

d. Repeat [Step 20.3.b](#) through [Step 20.3.c](#) for each remaining zone-cluster node.

21. Boot the zone cluster

Installation of the zone cluster might take several minutes.

```
phys-schost# clzonecluster boot zone-cluster-name
```

22. (Exclusive-IP zone clusters) Manually configure an IPMP group.

The `clsetup` utility does not automatically configure IPMP groups for exclusive-IP zone clusters. You must create an IPMP group manually before you create a logical-hostname or shared-address resource.

```
phys-schost# ipadm create-ipmp -i interface sc_ipmp0
phys-schost# ipadm delete-addr interface/name
phys-schost# ipadm create-addr -T static -a IPaddress/prefix sc_ipmp0/name
```

Next Steps To configure Oracle Solaris Cluster 3.3 data services that you installed in a `solaris10` brand zone cluster, follow procedures for zone clusters in the applicable data-service manual. See

Oracle Solaris Cluster 3.3 Documentation (<http://www.oracle.com/technetwork/documentation/solaris-cluster-33-192999.html>).

To complete Trusted Extensions configuration, go to “How to Configure a Zone Cluster to Use Trusted Extensions” on page 201.

Otherwise, add file systems or storage devices to the zone cluster. See the following sections:

- “Adding File Systems to a Zone Cluster” on page 203
- “Adding Local File Systems to a Specific Zone-Cluster Node” on page 217
- “Adding Storage Devices to a Zone Cluster” on page 220

▼ How to Configure a Zone Cluster to Use Trusted Extensions

After you create a labeled brand zone cluster, perform the following steps to finish configuration to use Trusted Extensions.

1. Complete IP-address mappings for the zone cluster.

Perform this step on each node of the zone cluster.

a. From a node of the global cluster, display the node's ID.

```
phys-schost# cat /etc/cluster/nodeid
N
```

b. Log in to a zone-cluster node on the same global-cluster node.

Ensure that the SMF service has been imported and all services are up before you log in.

c. Determine the IP addresses used by this zone-cluster node for the private interconnect.

The cluster software automatically assigns these IP addresses when the cluster software configures a zone cluster.

In the `ifconfig -a` output, locate the `clprivnet0` logical interface that belongs to the zone cluster. The value for `inet` is the IP address that was assigned to support the use of the cluster private interconnect by this zone cluster.

```
zcl# ifconfig -a
lo0:3: flags=20010008c9<UP,LOOPBACK,RUNNING,NOARP,MULTICAST,IPv4,VIRTUAL> mtu 8232 index
  1
zone zcl
inet 127.0.0.1 netmask ff000000
net0: flags=1000843<UP,BROADCAST,RUNNING,MULTICAST,IPv4> mtu 1500 index 2
inet 10.11.166.105 netmask fffffff0 broadcast 10.11.166.255
groupname sc_imp0
ether 0:3:ba:19:fa:b7
```

```

ce0: flags=9040843<UP,BROADCAST,RUNNING,MULTICAST,DEPRECATED,IPv4,NOFAILOVER> mtu 1500
  index 4
inet 10.11.166.109 netmask ffffffff broadcast 10.11.166.255
groupname sc_imp0
ether 0:14:4f:24:74:d8
ce0:3: flags=1000843<UP,BROADCAST,RUNNING,MULTICAST,IPv4> mtu 1500 index 4
zone zc1
inet 10.11.166.160 netmask ffffffff broadcast 10.11.166.255
clprivnet0: flags=1009843<UP,BROADCAST,RUNNING,MULTICAST,MULTI_BCAST,PRIVATE,IPv4> mtu
  1500 index 7
inet 172.16.0.18 netmask ffffffff broadcast 172.16.0.23
ether 0:0:0:0:0:2
clprivnet0:3: flags=1009843<UP,BROADCAST,RUNNING,MULTICAST,MULTI_BCAST,PRIVATE,IPv4> mtu
  1500 index 7
zone zc1
inet 172.16.0.22 netmask ffffffff broadcast 172.16.0.23

```

d. Add to the zone-cluster node's /etc/inet/hosts file the following addresses of the zone-cluster node.

- The hostname for the private interconnect, which is `clusternodeN-priv`, where *N* is the global-cluster node ID

```
172.16.0.22    clusternodeN-priv
```

- Each net resource that was specified to the `clzonecluster` command when you created the zone cluster

e. Repeat on the remaining zone-cluster nodes.

2. Authorize communication with zone-cluster components.

Create new entries for the IP addresses used by zone-cluster components and assign each entry a CIPSO template. These IP addresses which exist in the zone-cluster node's `/etc/inet/hosts` file are as follows:

- Each zone-cluster node private IP address
- All `cl_privnet` IP addresses in the zone cluster
- Each logical-hostname public IP address for the zone cluster
- Each shared-address public IP address for the zone cluster

```

phys-schost# tncfg -t cipso
tncfg:cipso> add host=ipaddress1
tncfg:cipso> add host=ipaddress2
...
tncfg:cipso> exit

```

For more information about CIPSO templates, see [“How to Configure a Different Domain of Interpretation”](#) in [“Trusted Extensions Configuration and Administration”](#).

3. Set IP strict multihoming to weak.

Perform the following commands on each node of the zone cluster.

```
phys-schost# ipadm set-prop -p hostmodel=weak ipv4
phys-schost# ipadm set-prop -p hostmodel=weak ipv6
```

For more information about the `hostmodel` property, see [“hostmodel \(ipv4 or ipv6\)”](#) in [“Oracle Solaris 11.2 Tunable Parameters Reference Manual”](#).

Next Steps To add file systems or storage devices to the zone cluster. See the following sections:

- [“Adding File Systems to a Zone Cluster”](#) on page 203
- [“Adding Local File Systems to a Specific Zone-Cluster Node”](#) on page 217
- [“Adding Storage Devices to a Zone Cluster”](#) on page 220

See Also If you want to update the software on a zone cluster, follow procedures in [Chapter 11, “Updating Your Software,”](#) in [“Oracle Solaris Cluster System Administration Guide”](#). These procedures include special instructions for zone clusters, where needed.

Adding File Systems to a Zone Cluster

After a file system is added to a zone cluster and brought online, the file system is authorized for use from within that zone cluster. To mount the file system for use, configure the file system by using cluster resources such as `SUNW.HASStoragePlus` or `SUNW.ScalMountPoint`.

Note - To add a file system whose use is limited to a single zone-cluster node, see instead [“Adding Local File Systems to a Specific Zone-Cluster Node”](#) on page 217.

This section provides the following procedures to add file systems for use by the zone cluster:

- [“How to Add a Highly Available Local File System to a Zone Cluster \(clsetup\)”](#) on page 204
- [“How to Add a Highly Available Local File System to a Zone Cluster \(CLI\)”](#) on page 205
- [“How to Add a ZFS Storage Pool to a Zone Cluster \(clsetup\)”](#) on page 207
- [“How to Add a ZFS Storage Pool to a Zone Cluster \(CLI\)”](#) on page 209
- [“How to Add a Sun QFS Shared File System to a Zone Cluster \(CLI\)”](#) on page 214
- [“How to Add a Cluster File System to a Zone Cluster \(clsetup\)”](#) on page 210
- [“How to Add a UFS Cluster File System to a Zone Cluster \(CLI\)”](#) on page 213
- [“How to Add an Oracle ACFS File System to a Zone Cluster \(CLI\)”](#) on page 216

You can also use the Oracle Solaris Cluster Manager GUI to add a file system to a zone cluster. For GUI log-in instructions, see [“How to Access Oracle Solaris Cluster Manager”](#) in [“Oracle Solaris Cluster System Administration Guide”](#).

▼ How to Add a Highly Available Local File System to a Zone Cluster (c1setup)

Perform this procedure to configure a highly available local file system on the global cluster for use by a zone cluster. The file system is added to the zone cluster and is configured with an HAStoragePlus resource to make the local file system highly available.

Note - Alternatively, you can use the command line to perform this task. See [“How to Add a Highly Available Local File System to a Zone Cluster \(CLI\)”](#) on page 205.

Perform all steps of the procedure from a node of the global cluster.

1. **Assume the root role on a node of the global cluster that hosts the zone cluster.**
2. **On the global cluster, create a file system that you want to use in the zone cluster.**

Ensure that the file system is created on shared disks.

3. **Start the c1setup utility.**

```
phys-schost# c1setup
```

The Main Menu is displayed.

Tip - To return to a previous screen, type the < key and press Return.

4. **Choose the Zone Cluster menu item.**
The Zone Cluster Tasks Menu is displayed.
5. **Choose the Add File System/Storage Device to a Zone Cluster menu item.**
The Select Zone Cluster menu is displayed.

6. **Choose the zone cluster where you want to add the file system.**

The Storage Type Selection menu is displayed.

7. **Choose the File System menu item.**

The File System Selection for the Zone Cluster menu is displayed.

8. **Choose the file system you want to add to the zone cluster.**

The file systems in the list are those that are configured on the shared disks and can be accessed by the nodes where the zone cluster is configured. You can also type **e** to manually specify all properties for a file system.

The Mount Type Selection menu is displayed.

9. Choose the Loopback mount type.

The File System Properties for the Zone Cluster menu is displayed.

10. Change the properties that you are allowed to change for the file system you are adding.

Note - For UFS file systems, enable logging.

When, finished, type **d** and press Return.

11. Type c to save the configuration change.

The results of your configuration change are displayed.

12. When finished, exit the c1setup utility.

13. Verify the addition of the file system.

```
phys-schost# c1zonecluster show -v zone-cluster-name
```

Next Steps Configure the file system to be highly available by using an HAStoragePlus resource. The HAStoragePlus resource manages the mounting of the file system on the zone-cluster node that currently host the applications that are configured to use the file system. See [“Enabling Highly Available Local File Systems”](#) in [“Oracle Solaris Cluster Data Services Planning and Administration Guide”](#).

▼ How to Add a Highly Available Local File System to a Zone Cluster (CLI)

Perform this procedure to add a highly available local file system on the global cluster for use by the zone cluster.

Note - Alternatively, you can use the `c1setup` utility to perform this task. See [“How to Add a Highly Available Local File System to a Zone Cluster \(c1setup\)”](#) on page 204.

To add a ZFS pool to a zone cluster, instead perform procedures in [“How to Add a ZFS Storage Pool to a Zone Cluster \(c1setup\)”](#) on page 207. Or, to configure a ZFS storage pool to be highly available in a zone cluster, see [“How to Set Up the HAStoragePlus Resource Type to Make a Local Solaris ZFS File System Highly Available”](#) in [“Oracle Solaris Cluster Data Services Planning and Administration Guide”](#).

1. Assume the root role on a node of the global cluster that hosts the zone cluster.

You perform all steps of the procedure from a node of the global cluster.

2. On the global cluster, create a file system that you want to use in the zone cluster.

Ensure that the file system is created on shared disks.

3. Add the file system to the zone-cluster configuration.

```
phys-schost# clzonecluster configure zone-cluster-name
clzc:zone-cluster-name> add fs
clzc:zone-cluster-name:fs> set dir=mount-point
clzc:zone-cluster-name:fs> set special=disk-device-name
clzc:zone-cluster-name:fs> set raw=raw-disk-device-name
clzc:zone-cluster-name:fs> set type=FS-type
clzc:zone-cluster-name:fs> end
clzc:zone-cluster-name> verify
clzc:zone-cluster-name> commit
clzc:zone-cluster-name> exit
```

dir=mount-point

Specifies the file system mount point

special=disk-device-name

Specifies the name of the disk device

raw=raw-disk-device-name

Specifies the name of the raw disk device

type=FS-type

Specifies the type of file system

Note - Enable logging for UFS file systems.

4. Verify the addition of the file system.

```
phys-schost# clzonecluster show -v zone-cluster-name
```

Example 6-1 Adding a Highly Available Local File System to a Zone Cluster (CLI)

This example adds the local file system `/global/oracle/d1` for use by the `sczone` zone cluster.

```
phys-schost-1# clzonecluster configure sczone
clzc:sczone> add fs
clzc:sczone:fs> set dir=/global/oracle/d1
```

```

clzc:sczone:fs> set special=/dev/md/oracle/dsk/d1
clzc:sczone:fs> set raw=/dev/md/oracle/rdisk/d1
clzc:sczone:fs> set type=ufs
clzc:sczone:fs> add options [logging]
clzc:sczone:fs> end
clzc:sczone> verify
clzc:sczone> commit
clzc:sczone> exit

phys-schost-1# clzonecluster show -v sczone
...
Resource Name:                fs
dir:                          /global/oracle/d1
special:                       /dev/md/oracle/dsk/d1
raw:                           /dev/md/oracle/rdisk/d1
type:                          ufs
options:                       [logging]
cluster-control:              [true]
...

```

Next Steps Configure the file system to be highly available by using an HAStoragePlus resource. The HAStoragePlus resource manages the mounting of the file system on the zone-cluster node that currently host the applications that are configured to use the file system. See [“Enabling Highly Available Local File Systems”](#) in [“Oracle Solaris Cluster Data Services Planning and Administration Guide”](#).

▼ How to Add a ZFS Storage Pool to a Zone Cluster (clsetup)

Perform this procedure to add a ZFS storage pool to a zone cluster. The pool can be local to a single zone-cluster node or configured with HAStoragePlus to be highly available.

The `clsetup` utility discovers and displays all configured ZFS pools on the shared disks that can be accessed by the nodes where the selected zone cluster is configured. After you use the `clsetup` utility to add a ZFS storage pool in cluster scope to an existing zone cluster, you can use the `clzonecluster` command to modify the configuration or to add a ZFS storage pool in node-scope.

Note - Alternatively, you can use the command line to perform this task. See [“How to Add a ZFS Storage Pool to a Zone Cluster \(CLI\)”](#) on page 209.

Before You Begin Ensure that the ZFS pool is connected on shared disks that are connected to all nodes of the zone cluster. See [“Managing ZFS File Systems in Oracle Solaris 11.2”](#) for procedures to create a ZFS pool.

1. **Assume the root role on a node of the global cluster that hosts the zone cluster.**
You perform all steps of this procedure from a node of the global cluster.

2. Start the clsetup utility.

```
phys-schost# clsetup
```

The Main Menu is displayed.

Tip - To return to a previous screen, type the < key and press Return.

3. Choose the Zone Cluster menu item.

The Zone Cluster Tasks Menu is displayed.

4. Choose the Add File System/Storage Device to a Zone Cluster menu item.

The Select Zone Cluster menu is displayed.

5. Choose the zone cluster where you want to add the ZFS storage pool.

The Storage Type Selection menu is displayed.

6. Choose the ZFS menu item.

The ZFS Pool Selection for the Zone Cluster menu is displayed.

7. Choose the ZFS pool you want to add to the zone cluster.

The ZFS pools in the list are those that are configured on the shared disks and can be accessed by the nodes where the zone cluster is configured. You can also type **e** to manually specify properties for a ZFS pool.

The ZFS Pool Dataset Property for the Zone Cluster menu is displayed. The selected ZFS pool is assigned to the name property.

8. Type **d and press Return.**

The Review File Systems/Storage Devices for the Zone Cluster menu is displayed.

9. Type **c to save the configuration change.**

The results of your configuration change are displayed. For example:

```
>>> Result of Configuration Change to the Zone Cluster(sczone) <<<
```

```
Adding file systems or storage devices to sczone zone cluster...
```

```
The zone cluster is being created with the following configuration
```

```
/usr/cluster/bin/clzonecluster configure sczone
add dataset
set name=myzpool5
end
```

```
Configuration change to sczone zone cluster succeeded.
```


10. **When finished, exit the `clsetup` utility.**

11. **Verify the addition of the file system.**

```
phys-schost# clzonecluster show -v zoneclustername
```

12. **To make the ZFS storage pool highly available, configure the pool with an `HASStoragePlus` resource.**

The `HASStoragePlus` resource manages the mounting of file systems in the pool on the zone-cluster node that currently hosts the applications that are configured to use the file system. See [“Enabling Highly Available Local File Systems”](#) in [“Oracle Solaris Cluster Data Services Planning and Administration Guide”](#).

▼ How to Add a ZFS Storage Pool to a Zone Cluster (CLI)

Perform this procedure to add a ZFS storage pool to a zone cluster.

Note - Alternatively, you can use the `clsetup` utility to perform this task. See [“How to Add a ZFS Storage Pool to a Zone Cluster \(`clsetup`\)”](#) on page 207.

To configure a ZFS storage pool to be highly available in a zone cluster, see [“How to Set Up the `HASStoragePlus` Resource Type to Make a Local Solaris ZFS File System Highly Available”](#) in [“Oracle Solaris Cluster Data Services Planning and Administration Guide”](#).

1. **Assume the `root` role on a node of the global cluster that hosts the zone cluster.**
You perform all steps of this procedure from a node of the global zone.

2. **Create the ZFS storage pool on the global cluster.**

Ensure that the pool is connected on shared disks that are connected to all nodes of the zone cluster.

See [“Managing ZFS File Systems in Oracle Solaris 11.2”](#) for procedures to create a ZFS pool.

3. **Add the pool to the zone-cluster configuration.**

```
phys-schost# clzonecluster configure zone-cluster-name
clzc:zone-cluster-name> add dataset
clzc:zone-cluster-name:dataset> set name=ZFSpoolname
clzc:zone-cluster-name:dataset> end
clzc:zone-cluster-name> verify
clzc:zone-cluster-name> commit
clzc:zone-cluster-name> exit
```

4. **Verify the addition of the file system.**

```
phys-schost# clzonecluster show -v zone-cluster-name
```

Example 6-2 Adding a ZFS Storage Pool to a Zone Cluster (CLI)

The following example shows the ZFS storage pool `zpool1` added to the zone cluster `sczone`.

```
phys-schost-1# clzonecluster configure sczone
clzc:sczone> add dataset
clzc:sczone:dataset> set name=zpool1
clzc:sczone:dataset> end
clzc:sczone> verify
clzc:sczone> commit
clzc:sczone> exit

phys-schost-1# clzonecluster show -v sczone
...
Resource Name:                dataset
name:                          zpool1
...
```

Next Steps Configure the ZFS storage pool to be highly available by using an `HAStoragePlus` resource. The `HAStoragePlus` resource manages the mounting of file systems in the pool on the zone-cluster node that currently hosts the applications that are configured to use the file system. See [“Enabling Highly Available Local File Systems”](#) in [“Oracle Solaris Cluster Data Services Planning and Administration Guide”](#).

▼ How to Add a Cluster File System to a Zone Cluster (clsetup)

The `clsetup` utility discovers and displays the available file systems that are configured on the cluster nodes where the selected zone cluster is configured. When you use the `clsetup` utility to add a file system, the file system is added in cluster scope.

You can add the following types of cluster file systems to a zone cluster:

- UFS cluster file system - You specify the file system type in the `/etc/vfstab` file, using the `global` mount option. This file system can be located on the shared disk or on a Solaris Volume Manager device.
- Sun QFS shared file system - You specify the file system type in the `/etc/vfstab` file, using the `shared` mount option.
- ACFS - Discovered automatically, based on the `ORACLE_HOME` path you provide.

Note - Alternatively, you can use the command line to perform this task. See one of the following procedures:

- [“How to Add a UFS Cluster File System to a Zone Cluster \(CLI\)”](#) on page 213
 - [“How to Add a Sun QFS Shared File System to a Zone Cluster \(CLI\)”](#) on page 214
 - [“How to Add an Oracle ACFS File System to a Zone Cluster \(CLI\)”](#) on page 216
-

Before You Begin Ensure that the cluster file system you want to add to the zone cluster is configured. See “Planning Cluster File Systems” on page 37 and Chapter 5, “Creating a Cluster File System”.

1. **Assume the root role on a node of the global cluster that hosts the zone cluster.**
You perform all steps of this procedure from a node of the global cluster.
2. **On each node of the global cluster that hosts a zone-cluster node, add an entry to the /etc/vfstab file for the file system that you want to mount on the zone cluster.**

```
phys-schost# vi /etc/vfstab
```

- **For a UFS entry, include the global mount option, similar to the following example:**

```
/dev/md/datadg/dsk/d0 /dev/md/datadg/rdisk/d0 /global/fs ufs 2 no global, logging
```

- **For a shared QFS entry, include the shared mount option, similar to the following example:**

```
Data-cz1 - /db_qfs/Data1 samfs - no shared,notrace
```

3. **On the global cluster, start the clsetup utility.**

```
phys-schost# clsetup
```

The Main Menu is displayed.

Tip - To return to a previous screen, type the < key and press Return.

4. **Choose the Zone Cluster menu item.**
The Zone Cluster Tasks Menu is displayed.
5. **Choose the Add File System/Storage Device to a Zone Cluster menu item.**
The Select Zone Cluster menu is displayed.
6. **Choose the zone cluster where you want to add the file system.**
The Storage Type Selection menu is displayed.
7. **Choose the File System menu item.**
The File System Selection for the Zone Cluster menu is displayed.
8. **Choose a file system from the list.**
You can also type **e** to manually specify all properties for a file system. If you are using an ACFS file system, you can select Discover ACFS and then specify the ORACLE_HOME directory.

The Mount Type Selection menu is displayed.

9. Choose the Loopback file system mount type for the zone cluster.

If you chose ACFS in [Step 7](#), the `clsetup` utility skips this step because ACFS supports only the direct mount type.

For more information about creating loopback file systems, see [“How to Create and Mount an LOFS File System”](#) in [“Managing File Systems in Oracle Solaris 11.2”](#).

The File System Properties for the Zone Cluster menu is displayed.

10. Specify the mount point directory.

Type the number for the `dir` property and press Return. Then type the LOFS mount point directory name in the New Value field and press Return.

When finished, type `d` and press Return. The Review File Systems/Storage Devices for the Zone Cluster menu is displayed.

11. Type `c` to save the configuration change.

The results of your configuration change are displayed. For example:

```
>>> Result of Configuration Change to the Zone Cluster(sczone) <<<

Adding file systems or storage devices to sczone zone cluster...

The zone cluster is being created with the following configuration

/usr/cluster/bin/clzonecluster configure sczone
add fs
set dir=/zones/sczone/dsk/d0
set special=/global/fs
set type=lofs
end

Configuration change to sczone zone cluster succeeded.
```

12. When finished, exit the `clsetup` utility.

13. Verify the addition of the LOFS file system.

```
phys-schost# clzonecluster show -v zone-cluster-name
```

Next Steps (Optional) Configure the cluster file system to be managed by an HAStoragePlus resource. The HAStoragePlus resource manages the mounting of the file systems in the global cluster, and later performs a loopback mount on the zone-cluster nodes that currently host the applications that are configured to use the file system. For more information, see [“Configuring an HAStoragePlus Resource for Cluster File Systems”](#) in [“Oracle Solaris Cluster Data Services Planning and Administration Guide”](#).

▼ How to Add a UFS Cluster File System to a Zone Cluster (CLI)

Perform this procedure to add a UFS cluster file system for use by a zone cluster.

Note - Alternatively, you can use the `clsetup` utility to perform this task. See [“How to Add a Cluster File System to a Zone Cluster \(clsetup\)”](#) on page 210.

1. **Assume the root role on a voting node of the global cluster that hosts the zone cluster.**

You perform all steps of this procedure from a voting node of the global cluster.

2. **On the global cluster, configure the cluster file system that you want to use in the zone cluster.**
3. **On each node of the global cluster that hosts a zone-cluster node, add an entry to the `/etc/vfstab` file for the file system that you want to mount on the zone cluster.**

```
phys-schost# vi /etc/vfstab
...
/dev/global/dsk/d12s0 /dev/global/rdisk/d12s0/ /global/fs ufs 2 no global, logging
```

4. **Configure the cluster file system as a loopback file system for the zone cluster.**

```
phys-schost# clzonecluster configure zone-cluster-name
clzc:zone-cluster-name> add fs
clzc:zone-cluster-name:fs> set dir=zone-cluster-lofs-mountpoint
clzc:zone-cluster-name:fs> set special=global-cluster-mount-point
clzc:zone-cluster-name:fs> set type=lofs
clzc:zone-cluster-name:fs> end
clzc:zone-cluster-name> verify
clzc:zone-cluster-name> commit
clzc:zone-cluster-name> exit
```

`dir=zone-cluster-lofs-mount-point`

Specifies the file system mount point for LOFS to make the cluster file system available to the zone cluster.

`special=global-cluster-mount-point`

Specifies the file system mount point of the original cluster file system in the global cluster.

For more information about creating loopback file systems, see [“How to Create and Mount an LOFS File System”](#) in [“Managing File Systems in Oracle Solaris 11.2”](#).

5. **Verify the addition of the LOFS file system.**

```
phys-schost# clzonecluster show -v zone-cluster-name
```

Example 6-3 Adding a UFS Cluster File System to a Zone Cluster (CLI)

The following example shows how to add a cluster file system with mount point `/global/apache` to a zone cluster. The file system is available to a zone cluster using the loopback mount mechanism at the mount point `/zone/apache`.

```
phys-schost-1# 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/apache ufs 2 yes global, logging

phys-schost-1# clzonecluster configure zone-cluster-name
clzc:zone-cluster-name> add fs
clzc:zone-cluster-name:fs> set dir=/zone/apache
clzc:zone-cluster-name:fs> set special=/global/apache
clzc:zone-cluster-name:fs> set type=lofs
clzc:zone-cluster-name:fs> end
clzc:zone-cluster-name> verify
clzc:zone-cluster-name> commit
clzc:zone-cluster-name> exit

phys-schost-1# clzonecluster show -v sczone
...
Resource Name:                fs
dir:                          /zone/apache
special:                       /global/apache
raw:
type:                          lofs
options:                       []
cluster-control:              true
...
```

Next Steps Configure the cluster file system to be available in the zone cluster by using an `HAStoragePlus` resource. The `HAStoragePlus` resource manages the mounting of the file systems in the global cluster, and later performs a loopback mount on the zone-cluster nodes that currently host the applications that are configured to use the file system. For more information, see [“Configuring an HAStoragePlus Resource for Cluster File Systems”](#) in [“Oracle Solaris Cluster Data Services Planning and Administration Guide”](#).

▼ How to Add a Sun QFS Shared File System to a Zone Cluster (CLI)

Perform this task to add a Sun QFS shared file system for use by a zone cluster.

Note - Alternatively, you can use the `clsetup` utility to perform this task. See [“How to Add a Cluster File System to a Zone Cluster \(clsetup\)”](#) on page 210.

At this time, Sun QFS shared file systems are only supported for use in clusters that are configured with Oracle RAC. On clusters that are not configured with Oracle RAC, you can use a single-machine Sun QFS file system that is configured as a highly available local file system.

1. **Assume the root role on a node of the global cluster that hosts the zone cluster.**
You perform all steps of this procedure from a node of the global cluster.
2. **On the global cluster, configure the Sun QFS shared file system that you want to use in the zone cluster.**
Follow procedures for shared file systems in your Sun QFS documentation..
3. **On each node of the global cluster that hosts a zone-cluster node, add an entry to the `/etc/vfstab` file for the file system that you want to mount on the zone cluster.**
4. **Add the file system to the zone cluster configuration.**

```
phys-schost# clzonecluster configure zone-cluster-name
clzc:zone-cluster-name> add fs
clzc:zone-cluster-name:fs> set dir=mount-point
clzc:zone-cluster-name:fs> set special=QFS-file-system-name
clzc:zone-cluster-name:fs> set type=samfs
clzc:zone-cluster-name:fs> end
clzc:zone-cluster-name> verify
clzc:zone-cluster-name> commit
clzc:zone-cluster-name> exit
```

5. **Verify the addition of the file system.**

```
phys-schost# clzonecluster show -v zone-cluster-name
```

Example 6-4 Adding a Sun QFS Shared File System as a Direct Mount to a Zone Cluster (CLI)

The following example shows the Sun QFS shared file system `Data-cz1` added to the zone cluster `sczone`. From the global cluster, the mount point of the file system is `/zones/sczone/root/db_qfs/Data1`, where `/zones/sczone/root/` is the zone's root path. From the zone-cluster node, the mount point of the file system is `/db_qfs/Data1`.

```
phys-schost-1# vi /etc/vfstab
#device    device    mount    FS        fsck    mount    mount
#to mount  to fsck   point    type      pass    at boot  options
#
Data-cz1   -        /zones/sczone/root/db_qfs/Data1 samfs - no shared,notrace

phys-schost-1# clzonecluster configure sczone
clzc:sczone> add fs
```

```
clzc:sczone:fs> set dir=/db_qfs/Data1
clzc:sczone:fs> set special=Data-cz1
clzc:sczone:fs> set type=samfs
clzc:sczone:fs> end
clzc:sczone> verify
clzc:sczone> commit
clzc:sczone> exit

phys-schost-1# clzonecluster show -v sczone
...
Resource Name:                fs
dir:                          /db_qfs/Data1
special:                       Data-cz1
raw:
type:                          samfs
options:                       []
...
```

▼ How to Add an Oracle ACFS File System to a Zone Cluster (CLI)

Perform this procedure to add an Oracle ACFS file system for use by a zone cluster.

Note - Alternatively, you can use the `clsetup` utility to perform this task. See [“How to Add a Cluster File System to a Zone Cluster \(clsetup\)”](#) on page 210.

Before You Begin Ensure that the Oracle ACFS file system is created and ready for use by a zone cluster. See [“How to Create an Oracle ACFS File System”](#) on page 176.

1. **Assume the root role or become an administrator that provides `solaris.cluster.admin` and `solaris.cluster.modify` authorizations.**

2. **Add the Oracle ACFS file system to the zone cluster.**

Perform this step from the global zone of one node.

```
# clzonecluster configure zonecluster
clzc:zonecluster> add fs
clzc:zonecluster:fs> set dir=mountpoint
clzc:zonecluster:fs> set special=/dev/asm/volume-dev-path
clzc:zonecluster:fs> set type=acfs
clzc:zonecluster:fs> end
clzc:zonecluster> exit
```

3. **Verify that the file system is added to the zone cluster.**

```
# clzonecluster show zonecluster
...
```



```

Resource Name:      fs
dir:                mountpoint
special            /dev/asm/volume-dev-path
raw:
type:              acfs
options:           []
cluster-control:   true
...

```

Adding Local File Systems to a Specific Zone-Cluster Node

This section describes how to add file systems that are dedicated to a single zone-cluster node. To instead configure file systems for use by the entire zone cluster, go to [“Adding File Systems to a Zone Cluster” on page 203](#).

This section contains the following procedures:

- [“How to Add a Local File System to a Specific Zone-Cluster Node \(CLI\)” on page 217](#)
- [“How to Add a Local ZFS Storage Pool to a Specific Zone-Cluster Node \(CLI\)” on page 219](#)

▼ How to Add a Local File System to a Specific Zone-Cluster Node (CLI)

Perform this procedure to add a local file system to a single, specific zone-cluster node of a specific zone cluster. The file system is not managed by Oracle Solaris Cluster software but is instead passed to the underlying Oracle Solaris zone.

Note - To add a highly available local file system to a zone cluster, perform procedures in [“How to Add a Highly Available Local File System to a Zone Cluster \(clsetup\)” on page 204](#) or [“How to Add a Highly Available Local File System to a Zone Cluster \(CLI\)” on page 205](#).

1. **Assume the root role on a node of the global cluster that hosts the zone cluster.**

Note - Perform all steps of this procedure from a node of the global cluster.

2. **Create the local file system that you want to configure to a specific zone-cluster node.**

Use local disks of the global-cluster node that hosts the intended zone-cluster node.

3. Add the file system to the zone-cluster configuration in the node scope.

```
phys-schost# clzonecluster configure zoneclustername
clzc:zoneclustername> select node physical-host=baseclusternode
clzc:zoneclustername:node> add fs
clzc:zoneclustername:node:fs> set dir=mountpoint
clzc:zoneclustername:node:fs> set special=disk-device-name
clzc:zoneclustername:node:fs> set raw=raw-disk-device-name
clzc:zoneclustername:node:fs> set type=FS-type
clzc:zoneclustername:node:fs> end
clzc:zoneclustername:node> end
clzc:zoneclustername> verify
clzc:zoneclustername> commit
clzc:zoneclustername> exit
```

dir=mountpoint

Specifies the file-system mount point

special=disk-device-name

Specifies the name of the disk device

raw=raw-disk-device-name

Specifies the name of the raw-disk device

type=FS-type

Specifies the type of file system

Note - Enable logging for UFS file systems.

4. Verify the addition of the file system.

```
phys-schost# clzonecluster show -v zoneclustername
```

Example 6-5 Adding a Local File System to a Zone-Cluster Node (CLI)

This example adds a local UFS file system `/local/data` for use by a node of the `sczone` zone cluster. This zone-cluster node is hosted on `global`—cluster node `phys-schost-1`.

```
phys-schost-1# clzonecluster configure sczone
clzc:sczone> select node physical-host=phys-schost-1
clzc:sczone:node> add fs
clzc:sczone:node:fs> set dir=/local/data
clzc:sczone:node:fs> set special=/dev/md/localdg/dsk/d1
clzc:sczone:node:fs> set raw=/dev/md/localdg/rdisk/d1
clzc:sczone:node:fs> set type=ufs
clzc:sczone:node:fs> add options [logging]
clzc:sczone:node:fs> end
```

```

clzc:sczone:node> end
clzc:sczone> verify
clzc:sczone> commit
clzc:sczone> exit

phys-schost-1# clzonecluster show -v sczone
...
--- Solaris Resources for phys-schost-1 ---
...
Resource Name:                fs
dir:                          /local/data
special:                       /dev/md/localdg/dsk/d1
raw:                           /dev/md/localdg/rdisk/d1
type:                          ufs
options:                       [logging]
cluster-control:              false ...

```

▼ How to Add a Local ZFS Storage Pool to a Specific Zone-Cluster Node (CLI)

Perform this procedure to add a local ZFS storage pool to a specific zone-cluster node. The local ZFS pool is not managed by Oracle Solaris Cluster software but is instead passed to the underlying Oracle Solaris zone.

Note - To add a highly available local ZFS pool to a zone cluster, see [“How to Add a Highly Available Local File System to a Zone Cluster \(clsetup\)”](#) on page 204 or [“How to Add a ZFS Storage Pool to a Zone Cluster \(CLI\)”](#) on page 209.

Perform all steps of the procedure from a node of the global cluster.

1. **Assume the root role on a node of the global cluster that hosts the zone cluster.**
2. **Create the local ZFS pool that you want to configure to a specific zone-cluster node.**

Use local disks of the global-cluster node that hosts the intended zone-cluster node.

3. **Add the pool to the zone-cluster configuration in the node scope.**

```

phys-schost# clzonecluster configure zoneclustername
clzc:zoneclustername> select node physical-host=baseclusternode
clzc:zoneclustername:node> add dataset
clzc:zoneclustername:node:dataset> set name=localZFSpoolname
clzc:zoneclustername:node:dataset> end
clzc:zoneclustername:node> end
clzc:zoneclustername> verify
clzc:zoneclustername> commit
clzc:zoneclustername> exit

```

```
set name=localZFSpoolname  
    Specifies the name of the local ZFS pool
```

4. Verify the addition of the ZFS pool.

```
phys-schost# clzonecluster show -v zoneclustername
```

Example 6-6 Adding a Local ZFS Pool to a Zone-Cluster Node (CLI)

This example adds the local ZFS pool `local_pool` for use by a node of the `sczone` zone cluster. This zone-cluster node is hosted on global—cluster node `phys-schost-1`.

```
phys-schost-1# clzonecluster configure sczone  
clzc:sczone> select node physical-host=phys-schost-1  
clzc:sczone:node> add dataset  
clzc:sczone:node:dataset> set name=local_pool  
clzc:sczone:node:dataset> end  
clzc:sczone:node> end  
clzc:sczone> verify  
clzc:sczone> commit  
clzc:sczone> exit  
  
phys-schost-1# clzonecluster show -v sczone  
...  
--- Solaris Resources for phys-schost-1 ---  
...  
Resource Name:                dataset  
name:                          local_pool
```

Adding Storage Devices to a Zone Cluster

This section describes how to add the direct use of global storage devices by a zone cluster or add storage devices that are dedicated to a single zone-cluster node. Global devices are devices that can be accessed by more than one node in the cluster, either one node at a time or multiple nodes concurrently.

After a device is added to a zone cluster, the device is visible only from within that zone cluster.

You can also use the Oracle Solaris Cluster Manager GUI to add a storage device to a zone cluster. For GUI log-in instructions, see [“How to Access Oracle Solaris Cluster Manager”](#) in [“Oracle Solaris Cluster System Administration Guide”](#).

This section contains the following procedures:

- [“How to Add a Global Storage Device to a Zone Cluster \(clsetup\)”](#) on page 221

- [“How to Add a DID Device to a Zone Cluster \(CLI\)” on page 223](#)
- [“How to Add a Disk Set to a Zone Cluster \(CLI\)” on page 224](#)
- [“How to Add a Raw-Disk Device to a Specific Zone-Cluster Node \(CLI\)” on page 225](#)

▼ How to Add a Global Storage Device to a Zone Cluster (clsetup)

Perform this procedure to add one of the following types of storage devices in cluster scope:

- Raw-disk devices
- Solaris Volume Manager disk sets (not including multi-owner)

Note - Alternatively, you can use the command line to perform this task. See the following procedures:

- [“How to Add a DID Device to a Zone Cluster \(CLI\)” on page 223](#)
- [“How to Add a Disk Set to a Zone Cluster \(CLI\)” on page 224](#)

To add a raw-disk device to a specific zone-cluster node, go instead to [“How to Add a Raw-Disk Device to a Specific Zone-Cluster Node \(CLI\)” on page 225](#).

The `clsetup` utility discovers and displays the available storage devices that are configured on the cluster nodes where the selected zone cluster is configured. After you use the `clsetup` utility to add a storage device to an existing zone cluster, use the `clzonecluster` command to modify the configuration. For instructions on using the `clzonecluster` command to remove a storage device from a zone cluster, see [“How to Remove a Storage Device From a Zone Cluster”](#) in [“Oracle Solaris Cluster System Administration Guide”](#).

1. **Assume the `root` role on a node of the global cluster that hosts the zone cluster.**
You perform all steps of this procedure from a node of the global cluster.
2. **Identify the device to add to the zone cluster and determine whether it is online.**

```
phys-schost# cldevicegroup status
```

3. **If the device that you are adding is not online, bring it online.**

```
phys-schost# cldevicegroup online device
```

4. **Start the `clsetup` utility.**

```
phys-schost# clsetup
```

The Main Menu is displayed.

Tip - To return to a previous screen, type the < key and press Return.

5. Choose the Zone Cluster menu item.

The Zone Cluster Tasks Menu is displayed.

6. Choose the Add File System/Storage Device to a Zone Cluster menu item.

The Select Zone Cluster menu is displayed.

7. Choose the zone cluster where you want to add the storage device.

The Storage Type Selection menu is displayed.

8. Choose the Device menu item.

A list of the available devices is displayed.

9. Choose a storage device from the list.

You can also type **e** to manually specify properties for a storage device.

The Storage Device Property for the Zone Cluster menu is displayed.

10. Add or change any properties for the storage device you are adding.

Note - An asterisk (*) is used as a wildcard character in the path name.

When finished, type **d** and press Return. The Review File Systems/Storage Devices for the Zone Cluster menu is displayed.

11. Type c to save the configuration change.

The results of your configuration change are displayed. For example:

```
>>> Result of Configuration Change to the Zone Cluster(sczone) <<<
```

```
Adding file systems or storage devices to sczone zone cluster...
```

```
The zone cluster is being created with the following configuration
```

```
/usr/cluster/bin/clzonecluster configure sczone
add device
set match=/dev/md/ddg/*dsk/*
end
add device
set match=/dev/md/shared/1/*dsk/*
end
```

```
Configuration change to sczone zone cluster succeeded.
```

```
The change will become effective after the zone cluster reboots.
```

12. **When finished, exit the `clsetup` utility.**

13. **Verify the addition of the device.**

```
phys-schost# clzonecluster show -v zoneclustername
```

▼ How to Add a DID Device to a Zone Cluster (CLI)

Perform this procedure to add a DID device to a zone cluster.

Note - Alternatively, you can use the `clsetup` utility to perform this task. See [“How to Add a Global Storage Device to a Zone Cluster \(`clsetup`\)”](#) on page 221.

1. **Assume the root role on a node of the global cluster that hosts the zone cluster.**

You perform all steps of this procedure from a node of the global cluster.

2. **Identify the DID device to add to the zone cluster.**

The device you add must be connected to all nodes of the zone cluster.

```
phys-schost# cldevice list -v
```

3. **Add the DID device for use by the zone cluster.**

Note - An asterisk (*) is used as a wildcard character in the path name.

```
phys-schost# clzonecluster configure zone-cluster-name
clzc:zone-cluster-name> add device
clzc:zone-cluster-name:device> set match=/dev/did/*dsk/dNs*
clzc:zone-cluster-name:device> end
clzc:zone-cluster-name> verify
clzc:zone-cluster-name> commit
clzc:zone-cluster-name> exit
```

```
match=/dev/did/*dsk/dNs*
```

Specifies the full device path of the DID device

4. **Reboot the zone cluster.**

The change becomes effective after the zone cluster reboots.

```
phys-schost# clzonecluster reboot zone-cluster-name
```

Example 6-7 Adding a DID Device to a Zone Cluster (CLI)

The following example adds the DID device `d10` to the `sczone` zone cluster.

```
phys-schost-1# clzonecluster configure sczone
clzc:sczone> add device
clzc:sczone:device> set match=/dev/did/*dsk/d10s*
clzc:sczone:device> end
clzc:sczone> verify
clzc:sczone> commit
clzc:sczone> exit
```

```
phys-schost-1# clzonecluster reboot sczone
```

▼ How to Add a Disk Set to a Zone Cluster (CLI)

Perform this procedure to add an entire Solaris Volume Manager disk set to a zone cluster.

Note - Alternatively, you can use the `clsetup` utility to perform this task. See [“How to Add a Global Storage Device to a Zone Cluster \(clsetup\)”](#) on page 221.

1. **Assume the root role on a node of the global cluster that hosts the zone cluster.**
You perform all steps of this procedure from a node of the global cluster.
2. **Identify the disk set to add to the zone cluster and determine whether it is online.**

```
phys-schost# cldevicegroup status
```

3. **If the disk set that you are adding is not online, bring it online.**

```
phys-schost# cldevicegroup online diskset
```

4. **Determine the set number that corresponds to the disk set to add.**

```
phys-schost# ls -l /dev/md/diskset
lrwxrwxrwx 1 root root 8 Jul 22 23:11 /dev/md/diskset -> shared/set-number
```

5. **Add the disk set for use by the zone cluster.**

You must use a separate `add device` session for each `set match=` entry.

Note - An asterisk (*) is used as a wildcard character in the path name.

```
phys-schost# clzonecluster configure zone-cluster-name
clzc:zone-cluster-name> add device
clzc:zone-cluster-name:device> set match=/dev/md/diskset/*dsk/*
clzc:zone-cluster-name:device> end
clzc:zone-cluster-name> add device
clzc:zone-cluster-name:device> set match=/dev/md/shared/set-number/*dsk/*
clzc:zone-cluster-name:device> end
clzc:zone-cluster-name> verify
```



```
clzc:zone-cluster-name> commit
clzc:zone-cluster-name> exit
```

```
match=/dev/md/diskset/*disk/*
```

Specifies the full logical device path of the disk set

```
match=/dev/md/shared/N/*disk/*
```

Specifies the full physical device path of the disk set number

6. Reboot the zone cluster.

The change becomes effective after the zone cluster reboots.

```
phys-schost# clzonecluster reboot zone-cluster-name
```

Example 6-8 Adding a Disk Set to a Zone Cluster (CLI)

The following example adds the disk set oraset to the sczone zone cluster. The set number of the disk set is 3.

```
phys-schost-1# clzonecluster configure sczone
clzc:sczone> add device
clzc:sczone:device> set match=/dev/md/oraset/*disk/*
clzc:sczone:device> end
clzc:sczone> add device
clzc:sczone:device> set match=/dev/md/shared/3/*disk/*
clzc:sczone:device> end
clzc:sczone> verify
clzc:sczone> commit
clzc:sczone> exit
```

```
phys-schost-1# clzonecluster reboot sczone
```

▼ How to Add a Raw-Disk Device to a Specific Zone-Cluster Node (CLI)

Perform this procedure to add a raw-disk device to a specific zone-cluster node. This device would not be under Oracle Solaris Cluster control. Perform all steps of the procedure from a node of the global cluster.

Note - To add a raw-disk device for use by the full zone cluster, go instead to [“How to Add a Global Storage Device to a Zone Cluster \(clsetup\)”](#) on page 221 or [“How to Add a DID Device to a Zone Cluster \(CLI\)”](#) on page 223.

1. Assume the root role on a node of the global cluster that hosts the zone cluster.

You perform all steps of this procedure from a node of the global cluster.

2. **Identify the device (cNtXdYsZ) to add to the zone cluster and determine whether it is online.**
3. **Add the device to the zone-cluster configuration in the node scope.**

Note - An asterisk (*) is used as a wildcard character in the path name.

```
phys-schost# clzonecluster configure zone-cluster-name
clzc:zone-cluster-name> select node physical-host=baseclusternode
clzc:zone-cluster-name:node> add device
clzc:zone-cluster-name:node:device> set match=/dev/*dsk/cNtXdYs*
clzc:zone-cluster-name:node:device> end
clzc:zone-cluster-name:node> end
clzc:zone-cluster-name> verify
clzc:zone-cluster-name> commit
clzc:zone-cluster-name> exit
```

```
match=/dev/*dsk/cNtXdYs*
```

Specifies the full device path of the raw-disk device

4. **Verify the addition of the device.**

```
phys-schost# clzonecluster show -v zoneclustername
```

Example 6-9 Adding a Raw-Disk Device to a Specific Zone-Cluster Node (CLI)

The following example adds the raw-disk device c1t1d0s0 for use by a node of the sczone zone cluster. This zone-cluster node is hosted on global—cluster node phys-schost-1.

```
phys-schost-1# clzonecluster configure sczone
clzc:sczone> select node physical-host=phys-schost-1
clzc:sczone:node> add device
clzc:sczone:node:device> set match=/dev/*dsk/c1t1d0s0
clzc:sczone:node:device> end
clzc:sczone:node> end
clzc:sczone> verify
clzc:sczone> commit
clzc:sczone> exit

phys-schost-1# clzonecluster show -v sczone
...
--- Solaris Resources for phys-schost-1 ---
...
Resource Name:                device
name:                          /dev/*dsk/c1t1d0s0
```

Uninstalling Software From the Cluster

This chapter provides procedures for uninstalling or removing certain software from an Oracle Solaris Cluster configuration.

Note - If you want to uninstall a node from an established cluster, see [“Removing a Node From a Cluster”](#) in [“Oracle Solaris Cluster System Administration Guide”](#).

Uninstalling the Software

This section provides the following procedures to uninstall or remove certain software products from a global cluster.

- [“How to Unconfigure Oracle Solaris Cluster Software to Correct Installation Problems”](#) on page 227
- [“How to Uninstall Oracle Solaris Cluster Quorum Server Software”](#) on page 230
- [“How to Unconfigure a Zone Cluster”](#) on page 231
- [“How to Uninstall the Availability Suite Feature of the Oracle Solaris Software”](#) on page 233

▼ **How to Unconfigure Oracle Solaris 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 142, do not perform this procedure. Instead, go to “[How to Uninstall Oracle Solaris Cluster Software From a Cluster Node](#)” in “[Oracle Solaris Cluster System Administration Guide](#)”.

Before You Begin Attempt to rerun cluster configuration of the node by using the `scinstall` utility. You can correct certain cluster node configuration failures by repeating Oracle Solaris 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, assume the root role.**
- 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 an Existing Cluster or Zone Cluster](#)” in “[Oracle Solaris Cluster System Administration Guide](#)” for procedures.

2. On a node that you intend to unconfigure, assume the root role.

3. Shut down the node.

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

4. Reboot the node into noncluster mode.

■ **SPARC:**

```
ok boot -x
```

■ **x86:**

- a. **In the GRUB menu, use the arrow keys to select the appropriate Oracle Solaris entry and type `e` to edit its commands.**

For more information about GRUB based booting, see [“Booting a System”](#) in [“Booting and Shutting Down Oracle Solaris 11.2 Systems”](#).

- b. **In the boot parameters screen, use the arrow keys to select the kernel entry and type `e` to edit the entry.**
- c. **Add `-x` to the `multiboot` command to specify that the system boot into noncluster mode.**
- d. **Press Enter to accept the change and return to the boot parameters screen.**
The screen displays the edited command.
- 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 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 Oracle Solaris Cluster packages.**

```
phys-schost# cd /
```

6. **Remove the node from the cluster configuration.**

- **To unconfigure the node but leave Oracle Solaris Cluster software installed, run the following command:**

```
phys-schost# /usr/cluster/bin/clnode remove
```

The node is removed from the cluster configuration but Oracle Solaris Cluster software is not removed from the node.

See the [clnode\(1CL\)](#) man page for more information.

- **To unconfigure the node and also remove Oracle Solaris Cluster software, run the following command:**

```
phys-schost# /usr/cluster/bin/scinstall -r [-b BE-name]
```

-r

Removes cluster configuration information and uninstalls Oracle Solaris Cluster framework and data-service software from the cluster node. You can then reinstall the node or remove the node from the cluster.

-b *BE-name*

Specifies the name of a new boot environment, which is where you boot into after the uninstall process completes. Specifying a name is optional. If you do not specify a name for the boot environment, one is automatically generated.

See the [scinstall\(1M\)](#) man page for more information.

7. Repeat [Step 2](#) through [Step 6](#) for each additional node to unconfigure.

Troubleshooting If the cluster node that you are removing is at least partially configured with the cluster, running the `clnode remove` command might exit with errors such as `Node is still enabled`. If such errors occur, add the `-F` option to the `clnode remove` command.

Next Steps Before you reinstall or reconfigure Oracle Solaris 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 “[Oracle Solaris Cluster 4.2 Hardware Administration Manual](#)” and the removal procedure in the Oracle Solaris Cluster manual for your storage array.

▼ How to Uninstall Oracle Solaris Cluster Quorum Server Software

Before You Begin Before you uninstall Oracle Solaris Cluster Quorum Server software, ensure that you have completed the following tasks:

- On each cluster that uses the quorum server, remove the quorum server as a quorum device. Perform the steps in “[How to Remove a Quorum Device](#)” in “[Oracle Solaris Cluster System Administration Guide](#)”.

In normal operation, this step also removes the quorum server information on the quorum-server host. If communications between the cluster and the quorum-server host computer are down during this step, you must clean up the invalid quorum server information on the quorum-server host computer. Perform the steps in “[Cleaning Up Stale Quorum Server Cluster Information](#)” in “[Oracle Solaris Cluster System Administration Guide](#)”.

- On each quorum-server host computer, stop the quorum server by following the steps in “[How to Stop a Quorum Server](#)” in “[Oracle Solaris Cluster System Administration Guide](#)”.

1. **(Optional) To use the interactive graphical interface, ensure that the display environment of the host server to uninstall is set to display the GUI.**

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

2. **Assume the root role on the quorum server host computer to uninstall.**
Alternatively, if your user account is assigned the System Administrator profile, issue commands as non-root through a profile shell, or prefix the command with the pfexec command.
3. **Uninstall the quorum server software.**

```
quorumserver# pkg uninstall ha-cluster/group-package/ha-cluster-quorum-server* \
ha-cluster/service/quorum-server*
```

4. **(Optional) Clean up or remove the quorum server directories.**
By default, this directory is /var/scqsd.

▼ How to Unconfigure a Zone Cluster

1. **Assume the root role on a node of the global cluster.**
You perform all steps of this procedure from a node of the global cluster.
2. **Take offline each resource group in the zone cluster and disable its resources.**

Note - The following steps are performed from a global-cluster node. To instead perform these steps from a node of the zone cluster, log in to the zone-cluster node and omit “-Z zone-cluster” from each command.

- a. **Take each resource offline.**

```
phys-schost# clresource offline -Z zone-cluster resource-group
```

- b. **List all enabled resources in the zone cluster.**

```
phys-schost# clresource show -Z zone-cluster -p Enabled
=== Resources ===
```

```
Resource:                                resource
Enabled{nodename1}:                       True
Enabled{nodename2}:                       True
...
```

- c. **Identify those resources that depend on other resources.**

```
phys-schost# clresource show -Z zone-cluster -p resource_dependencies
=== Resources ===
```

```
Resource:                               node
Resource_dependencies:                   node
...
```

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# clresource disable -Z zone-cluster resource
```

See the [clresource\(1CL\)](#) man page for more information.

e. Verify that all resources are disabled.

```
phys-schost# clresource show -Z zone-cluster -p Enabled
=== Resources ===
```

```
Resource:                               resource
Enabled{nodename1}:                     False
Enabled{nodename2}:                     False
...
```

f. Move each resource group to the unmanaged state.

```
phys-schost# clresourcegroup unmanage -Z zone-cluster resource-group
```

g. Verify that all resources on all nodes are offline and that all resource groups are in the Unmanaged state.

```
phys-schost# cluster status -Z zone-cluster -t resource,resourcegroup
```

h. Delete all resource groups and their resources from the zone cluster.

```
phys-schost# clresourcegroup delete -F -Z zone-cluster +
```

3. Halt the zone cluster.

```
phys-schost# clzonecluster halt zone-cluster-name
```

4. Uninstall the zone cluster.

```
phys-schost# clzonecluster uninstall zone-cluster-name
```

5. Unconfigure the zone cluster.

```
phys-schost# clzonecluster delete zone-cluster-name
```


▼ How to Uninstall the Availability Suite Feature of the Oracle Solaris Software

Perform the following steps on each node where you want to uninstall the Availability Suite feature of Oracle Solaris software.

1. **Assume the root role.**
2. **Uninstall all Availability Suite feature packages.**

```
phys-schost# pkg uninstall group/feature/storage-avs \  
storage/avs/avs-cache-management \  
storage/avs/avs-point-in-time-copy \  
storage/avs/avs-remote-mirror \  
driver/storage/sv
```


Index

Numbers and Symbols

- `/etc/hosts.allow` setting, 79, 85, 101, 107
- `/etc/inet/hosts` file
 - configuring, 51
 - planning, 17
- `/etc/inet/ntp.conf` file
 - configuring, 144
- `/etc/inet/ntp.conf.include` file
 - creating, 146, 147
- `/etc/inet/ntp.conf.sc` file
 - configuring, 146, 147
- `/etc/lvm/md.tab` file, 158
- `/etc/vfstab` file
 - adding mount points, 166
 - verifying the configuration, 167
- `/global` directory, 40
- `/opt/SUNWcluster/bin/` directory, 54
- `/opt/SUNWcluster/bin/pconsole` utility, 54
 - installing the software, 52
- `/usr/cluster/bin/` directory, 66
- `/usr/cluster/bin/claccess` command
 - adding nodes to the authorized-node list, 228
 - removing nodes from the authorized-node list, 126
- `/usr/cluster/bin/cldevice` command
 - determining device-ID names, 138
 - updating the global-devices namespace, 154
 - verifying command processing, 154
- `/usr/cluster/bin/clnode` command
 - viewing private hostnames, 144
- `/usr/cluster/bin/clquorumserver` command
 - starting the quorum server, 56
- `/usr/cluster/bin/clresource` command
 - disabling resources, 232
 - listing resources, 231
 - taking resource groups offline, 231
- `/usr/cluster/bin/clresourcegroup` command
 - deleting resource groups, 232
 - unmanaging resource groups, 232
- `/usr/cluster/bin/clsetup` utility
 - adding
 - cluster interconnects, 115
 - file systems to a zone cluster, 210
 - global storage devices to a zone cluster, 221
 - highly available local file systems to a zone cluster, 204
 - ZFS storage pools to a zone cluster, 207
 - changing private hostnames, 144
 - creating a zone cluster, 190
 - navigation tip, 192
 - post installation setup, 140
- `/usr/cluster/bin/cluster` check command
 - validating the cluster, 148
 - vfstab file check, 167
- `/usr/cluster/bin/cluster` command
 - adding
 - nodes, 129
 - zone clusters to the private IP-address range, 25, 117
 - creating new global clusters, 81
 - displaying private-network settings, 115
 - verifying installation mode, 143
- `/usr/cluster/bin/clzonecluster` command
 - booting the zone cluster, 200
 - halting the zone cluster, 232
 - installing the zone cluster, 197
- `/usr/cluster/bin/scinstall` command
 - adding nodes, 122
 - creating the global cluster, 71
 - creating the global cluster by using Automated Installer, 89
 - unconfiguring Oracle Solaris Cluster software, 227

A

- adapters
 - local MAC address, 19
 - Oracle VM Server for SPARC guest domains, 28
 - tagged VLAN
 - cluster interconnect guidelines, 27
 - public-network guidelines, 18
 - VLAN
 - private-network guidelines, 25
- adding, 46, 46
 - See also* configuring
 - See also* installing
 - cluster-file-system mount points to new nodes, 50
 - drives to a disk set, 155
 - file systems to a zone cluster, 203
 - local file systems to a zone cluster, 217
 - mediator hosts, 162
 - storage devices to zone clusters, 220
- address property, 195
- admin_low template
 - Trusted Extensions, 190
- administrative console
 - installing pconsole software, 52
 - IP addresses, 18
 - PATH, 54
- affinities
 - scalable device-group resource groups, 178
- AI *See* Automated Installer
- applications
 - installing on a cluster file system, 38
- archiveadm
 - create an archive, 102, 108
- authentication *See* authorized-node list
- authorized-node list
 - adding nodes, 228
 - removing nodes, 126
- Automated Installer
 - installing Oracle Solaris and Oracle Solaris Cluster software, 89
 - manifest, 97, 105, 111
 - troubleshooting installation, 102
- automatic power-saving shutdown
 - restriction, 13
- automatic reboot on disk-path failures, 78
- Availability Suite feature of Oracle Solaris software

- installing, 64
- uninstalling, 233

B

- booting
 - into noncluster mode, 228
- brand property, 193
- brand zones
 - zone clusters, 34

C

- capped-cpu scope type
 - zone clusters, 194
- changing
 - private hostnames, 143
 - private IP-address range, 116
- cipso template
 - global cluster IP addresses, 190
 - zone cluster IP addresses, 202
- claccess command
 - adding nodes to the authorized-node list, 228
 - removing nodes from the authorized-node list, 126
- cldevice command
 - determining device-ID names, 138
 - updating the global-devices namespace, 154
 - verifying command processing, 154
- clnode command
 - viewing private hostnames, 144
- clquorumserver command
 - starting the quorum server, 56
- clresource command
 - disabling resources, 232
 - listing resources, 231
 - taking resource groups offline, 231
- clresourcegroup command
 - bringing online resource groups, 176
 - creating
 - multiple-owner volume-manager framework resource group, 175
 - Oracle ACFS resource group, 182
 - Oracle RAC framework resource group, 174
 - scalable resource group, 178
 - SUNW.wait_zc_boot resource group, 181

- deleting resource groups, 232
- unmanaging resource groups, 232
- clsetup utility
 - adding
 - cluster interconnects, 115
 - file systems to a zone cluster, 210
 - global storage devices to a zone cluster, 221
 - highly available local file systems to a zone cluster, 204
 - ZFS storage pools to a zone cluster, 207
 - changing private hostnames, 144
 - creating a zone cluster, 190
 - navigation tip, 192
 - post installation setup, 140
- cluster check command
 - validating the cluster, 148
 - vfstab file check, 167
- cluster command
 - adding
 - nodes, 129
 - zone clusters to the private IP-address range, 25, 117
 - creating new global clusters, 81
 - displaying private-network settings, 115
 - verifying installation mode, 143
- cluster file systems
 - adding to a new node, 50
 - adding to a zone cluster, 213
 - adding to a zone cluster with clsetup, 210
 - configuring, 165
 - installing applications, 38
 - LOFS restriction, 37
 - mount options, 167
 - planning, 36, 40
 - restrictions
 - communication end points, 38
 - fattach command, 38
 - forcedirectio, 41
 - LOFS, 13
 - quotas, 37
 - zone clusters, 37
 - verifying the configuration, 167
- cluster interconnects
 - changing the private IP-address range, 116
 - configuring on a single-host cluster, 115
 - planning, 26
- cluster name, 23
- cluster nodes *See* global-cluster nodes *See* zone clusters
- clzonecluster command
 - booting the zone cluster, 200
 - halting the zone cluster, 232
 - installing the zone cluster, 197
- communication end points
 - restriction on cluster file systems, 38
- Compatibility Guide, 47
- config/enable_tcpwrappersdisabling, 76, 82, 124, 131
- config/enable_tcpwrappersenabling, 126, 132
- configuration guide, 47
- configuring
 - additional global-cluster nodes
 - by using an XML file, 129
 - by using scinstall, 122
 - cluster file systems, 165
 - cluster interconnects on a single-host cluster, 115
 - disk sets, 153
 - IP Filter feature of Oracle Solaris, 66
 - IPMP groups, 51
 - md.tab file, 158
 - Network Time Protocol (NTP), 144
 - new global clusters
 - by using an XML file, 81
 - by using Automated Installer, 89
 - by using scinstall, 71
 - Oracle ACFS file systems, 168
 - quorum devices, 137
 - quorum server software, 55
 - Solaris Volume Manager, 151
 - state database replicas, 152
 - troubleshooting, 80, 88, 128, 133
 - Trusted Extensions for use with zone clusters, 188, 201
 - Trusted Extensions software, 189
 - user work environment, 65
 - zone clusters, 188
- console-access devices
 - IP addresses, 18
 - planning, 18
- cpu-shares property, 194

- create
 - archive, 102
- creating
 - zone clusters, 190

D

- data services
 - support in solaris10 brand zone clusters, 197
- data-services installation
 - by using the pkg command, 59
- dedicated-cpu scope type
 - zone clusters, 194
- default router, 195
- defrouter property, 195
- dependencies
 - ScalDeviceGroup resource type, 179
 - SUNW.ScalDeviceGroup resource type, 179
 - SUNW.vucmm_svm resource type, 175
 - vucmm_svm resource type, 175
- device groups
 - planning, 37
 - replicated disks, 37
 - resource group affinities, 178
 - setting the replication property, 155
- device-ID names
 - determining, 138
- DID devices
 - adding to a zone cluster, 223
 - listing mappings, 156
- disabling
 - fencing protocols, 139
 - installation mode, 141
 - NWAM, 61
 - resources, 231, 232
- disk drives *See* drives
- disk sets
 - adding drives, 155
 - adding to a zone cluster, 224
 - configuring, 153
 - repartitioning drives, 157
 - reserved names, 42
- disk strings
 - dual-string mediator requirements, 162
- disk-path failures

- enabling automatic reboot, 78
- disk-path monitoring
 - default, 78
- disks *See* drives
- disksets *See* disk sets
- domain console network interfaces
 - IP addresses, 18
- drives
 - adding to disk sets, 155
 - mirroring differing device sizes, 43
 - repartitioning, 157
- dual-string mediators
 - adding hosts, 162
 - overview, 161
 - planning, 42
 - repairing data, 163
 - status, 163

E

- enable_priv_net property, 193
- enabling
 - automatic reboot on disk-path failures, 78
 - LOFS requirements, 13
 - NTP, 145
 - Oracle Solaris I/O multipathing, 51
 - remote login, 189
- examples
 - activating volumes in the md.tab file, 161
 - adding
 - drives to a disk set, 157
 - local file systems to a zone-cluster node, 218
 - local ZFS pools to a zone-cluster node, 220
 - mediator hosts, 163
 - raw-disk devices to a zone-cluster node, 226
 - adding devices to a zone cluster
 - DID devices, 223
 - disk sets, 225
 - adding file systems to a zone cluster
 - cluster file systems, 214
 - highly available local file systems, 206
 - Sun QFS shared file systems, 215
 - ZFS storage pools, 210
 - configuring Oracle Solaris Cluster software
 - on additional nodes by using scinstall, 128
 - on all nodes by using an XML file, 87

- on all nodes by using `scinstall`, 79
- creating a cluster file system, 167
- creating a disk set, 155
- creating state database replicas, 152
- `md.tab` file, 159
- updating SCSI quorum devices after adding a node, 136

exclusive IP type

- restriction for zone clusters, 34

explorer command, 148

F

failover file systems *See* highly available local file systems

Fair Share Scheduler shares, 194

`fattach` command restriction, 38

fencing protocols, 28

- disabling, 139
- NAS devices, 21
- SCSI quorum devices, 30
- software quorum, 30

file system logging

- planning, 43

file systems

- adding to a zone cluster, 203
 - highly available local file systems, 204
 - ZFS storage pools, 207

`forcedirectio` command

- restriction, 41

G

global devices

- file system limitation to UFS, 15
- namespace
 - updating, 154
- planning, 36
- restrictions on `lofi` device, 14
- space requirement for `lofi` device, 14

global fencing, 28

global file systems *See* cluster file systems

global storage devices

- adding to a zone cluster, 221

global-cluster nodes

adding nodes

- by using an XML file, 129
- by using `scinstall`, 122
- updating quorum devices after, 134

creating the global cluster

- by using an XML file, 81
- by using Automated Installer, 89
- by using `scinstall`, 71

IP addresses, 18

naming, 23

node ID number assignment, 23

Oracle VM Server for SPARC domains, 47

planning, 23

planning for zone clusters, 33

verifying

- installation mode, 143
- status, 78, 84, 100, 107, 113, 126

`Global_zone` resource-type property, 34

GUI

- tasks you can perform
 - add a file system to a zone cluster, 203
 - add storage a zone cluster, 220
 - create a zone cluster, 192
 - view capped memory for a zone cluster, 195

H

HA for NFS

- coexistence with highly available local file systems, 79, 85, 99, 127
- restriction with LOFS, 13, 37

hardware RAID

- internal disk mirroring, 57

HAStoragePlus, 34

- See also* highly available local file systems
- mounting file systems in zone clusters, 34

high-priority processes

- restriction, 22

highly available local file systems, 34

- See also* HAStoragePlus
- adding to a zone cluster, 204
- alternative to cluster file systems, 37
- coexistence with HA for NFS, 79, 85, 99, 127
- LOFS restriction with HA for NFS, 13
- quotas support, 37

highly available ZFS storage pools

- adding to a zone cluster, 207
- hostmodel property, 202
- hostname property, 195
- hosts file
 - configuring, 51
 - planning, 17
- hot-spare disks
 - planning, 42

I

- I/O multipathing software
 - Oracle VM Server for SPARC restriction, 32
- initialization files, 65
- installation mode
 - disabling, 141
 - verifying, 143
- installing, 46, 46
 - See also* adding
 - See also* configuring
 - Availability Suite feature of Oracle Solaris software, 64
 - data services
 - by using the pkg command, 59
 - man pages, 53
 - Oracle Solaris Cluster software, 59
 - Oracle Solaris software
 - alone, 48
 - with Oracle Solaris Cluster software, 89
 - Oracle VM Server for SPARC software, 58
 - pconsole software, 52
 - quorum server software, 55
 - Solaris Volume Manager software, 151
 - Sun QFS software, 65
 - Trusted Extensions software, 189
- installing and configuring
 - Oracle Solaris and Oracle Solaris Cluster from unified archives, 102
- integrated mirroring, 57
- internal hardware disk mirroring, 57
- IP addresses
 - adding to a naming service, 49
 - changing the private IP-address range, 116
 - guidelines for zone clusters, 33
 - planning the private network, 24

- planning the public network, 17
- IP Filter
 - configuring, 66
 - restriction on scalable data services, 66
 - restrictions, 13
- IP network multipathing (IPMP) *See* IPMP
- IP strict multihoming, 202
- IP type
 - zone clusters, 34
- ip-type property, 193
- ipadm listing private IP addresses, 79, 85, 101, 107, 126, 132
- IPMP
 - automatic group creation during installation, 19
 - configuring groups, 51
 - planning the public network, 19
- IPv6 addresses
 - private network restriction, 26, 27
 - public-network use, 19

J

- java
 - minimum version, 14

L

- labeled brand zone clusters
 - configuring with Trusted Extensions, 201
 - options, 197
 - planning, 32
- LDAP configuration for Trusted Extensions, 190
- licenses
 - planning, 16
- lightweight processes
 - zone clusters, 194
- link-based IPMP groups, 19
- live migration
 - restriction for Oracle VM Server for SPARC software, 32
- local
 - file systems
 - adding to a zone cluster, 205
 - MAC address, 19
- local file systems
 - adding to a specific zone-cluster node, 217

adding to a zone cluster, 217, 217
 See also file systems, adding to a zone cluster
 local ZFS storage pools
 adding to a specific zone-cluster node, 219
 locked memory, 195
 locked property, 195
 lofi device
 space requirement, 14
 use restrictions, 14
 LOFS
 restriction, 13, 37
 log files
 Oracle Solaris Cluster installation, 77
 logging for cluster file systems
 planning, 43
 logical addresses
 planning, 19
 logical network interfaces
 restriction, 28
 loopback file system (LOFS)
 restriction, 13, 37
 LWPs
 zone clusters, 194

M

man pages
 installing, 53
 manifest
 Automated Installer, 97, 105, 111
 max-lwps property, 194
 max-msg-ids property, 194
 max-sem-ids
 property, 194
 max-shm-ids property, 194
 max-shm-memory property, 194
 md.tab file
 configuring, 158
 mediators *See* dual-string mediators
 message queue IDs
 zone clusters, 194
 mirroring
 differing device sizes, 43
 internal disks, 57
 multihost disks, 43

 planning, 43
 root pool planning, 43
 mount options for cluster file systems
 requirements, 167
 Sun QFS software, 211
 UFS, 39, 40, 205, 211
 mount points
 adding to new nodes, 50
 cluster file systems, 40
 modifying the `/etc/vfstab` file, 166
 nested, 41
 MPxIO *See* Oracle Solaris I/O multipathing
 multihost disks
 mirroring, 43
 planning, 42
 multiple-owner volume-manager framework resource
 group
 configuring, 175
 multiported disks *See* multihost disks
 multiuser services
 verifying, 78, 84, 126

N

naming conventions
 cluster, 23
 global-cluster nodes, 23
 private hostnames, 26
 raw-disk devices, 166
 tagged VLAN adapters, 27
 zone clusters, 34
 naming service
 adding IP-address mappings, 49
 NAS devices
 configuring as quorum devices, 137
 fencing, 21
 troubleshooting in a zone cluster installation, 192
 NAT and IP Filter feature of Oracle Solaris
 software, 13
 ncpus property, 194
 netmasks
 changing the private netmask, 116
 displaying the private netmask, 115
 planning the private network, 24
 quorum server requirements, 137
 Network Auto-Magic (NWAM)

- disabling, 61
 - restriction, 13
 - Network File System (NFS)
 - guidelines for cluster nodes, 21
 - HA for NFS restriction with LOFS, 21
 - network interface card (NIC) support for local MAC addresses, 19
 - Network Time Protocol (NTP)
 - configuring, 144
 - enabling, 145
 - guidelines, 22
 - restarting, 145
 - starting, 145
 - network-attached storage *See* NAS
 - NFS *See* Network File System (NFS)
 - NIC support for local MAC addresses, 19
 - NIS servers
 - restriction for cluster nodes, 22
 - node IDs, 23
 - number assignment, 23
 - node lists
 - device groups, 42
 - nodes *See* global-cluster nodes
 - noncluster mode
 - booting into, 228
 - NTP
 - configuring, 144
 - enabling, 145
 - guidelines, 22
 - restarting, 145
 - starting, 145
 - ntp.conf file
 - configuring, 144
 - ntp.conf.include file
 - creating, 146, 147
 - ntp.conf.sc file
 - configuring, 146, 147
 - num_xip_zoneclusters property, 193
- O**
- Oracle ACFS file systems
 - adding to a zone cluster, 216
 - configuration diagrams, 169
 - creating, 168, 176
 - framework resource groups, 173
 - Oracle Explorer software, 148
 - Oracle RAC framework resource group
 - configuring for Oracle ACFS, 174
 - Oracle Solaris
 - minimum package, 13
 - Oracle Solaris Cluster Manager *See* GUI
 - Oracle Solaris Cluster software
 - group packages, 60, 95
 - publisher, 50, 53, 53, 55, 55, 63, 63, 64, 151
 - uninstalling software, 227
 - Oracle Solaris I/O multipathing software
 - enabling, 51
 - Oracle VM Server for SPARC restriction, 32
 - Oracle Solaris software
 - installing
 - alone, 48
 - Availability Suite feature, 64
 - Trusted Extensions feature, 188
 - with Oracle Solaris Cluster software, 89
 - planning, 12
 - /varfile system, 14
 - partitions, 14
 - root (/) file system, 14
 - volume managers, 15
 - publisher, 50, 64, 151
 - restrictions
 - automatic power-saving shutdown, 13
 - IP Filter feature, 13
 - Network Auto-Magic (NWAM), 13
 - SMF, 78, 84, 126
 - Trusted Extensions feature *See* Trusted Extensions
 - uninstalling
 - Availability Suite feature, 233
 - Oracle Solaris Unified Archive
 - installing Oracle Solaris and Oracle Solaris Cluster software, 89
 - how to, 96
 - replicate, restore, 89
 - Oracle Solaris unified archive
 - installing and configuring Oracle Solaris and Oracle Solaris Cluster software, 102
 - Oracle Solaris zones
 - fencing for NAS devices, 21
 - quorum servers, 20
 - Oracle VM Server for SPARC software
 - domains as cluster nodes, 47

- guest domains
 - Parallel Console Access (`pconsole`)
 - restriction, 52
 - planning the address range, 25
 - virtual adapter names, 28
 - guidelines, 31
 - I/O multipathing software restriction, 32
 - installing, 58
 - live migration restriction, 32
 - virtual disk multipathing restriction, 32
- P**
- package installation
 - Oracle Solaris Cluster man pages, 53
 - Oracle Solaris Cluster software, 59
 - `pconsole` software, 52
 - Parallel Console Access (`pconsole`) software
 - installing, 52
 - Oracle VM Server for SPARC restriction, 52
 - partitions
 - `/var` file system, 14
 - repartitioning drives, 157
 - root (`/`) file system, 14
 - swap, 14
 - volume manager, 15
 - patches *See* software updates
 - PATH
 - administrative console, 54
 - cluster nodes, 66
 - `pconsole` utility
 - installing the software, 52
 - starting, 54
 - using, 61, 98, 106, 111
 - physical memory
 - zone clusters, 195
 - physical property, 195, 195
 - `pkg` program
 - installing software packages, 59
 - private hostnames
 - changing, 143
 - planning, 26
 - verifying, 144
 - zone clusters, 34
 - private IP addresses
 - listing with `ipadm`, 79, 85, 101, 107, 126, 132
 - private network
 - adding zone clusters to the IP-address range, 25, 117
 - changing the IP-address range, 116
 - displaying settings, 115
 - IPv6 address restriction, 27
 - Oracle VM Server for SPARC guest domains, 25
 - planning, 24
 - use in zone clusters, 33
 - probe-based IPMP groups, 19
 - public network
 - adding IP addresses to a naming service, 49
 - IPv6 support, 19
 - planning, 18
 - publisher
 - Oracle Solaris Cluster software, 50, 53, 53, 55, 55, 63, 63, 64, 151
 - Oracle Solaris software, 50, 64, 151
- Q**
- QFS *See* Sun QFS software
 - quorum devices, 29
 - See also* quorum servers
 - initial configuration, 137
 - and mirroring, 44
 - NAS devices, 137
 - planning, 29
 - quorum servers, 137
 - restriction for replicated devices, 30
 - SATA disks, 30
 - SCSI protocol setting, 30
 - software quorum protocol, 30
 - troubleshooting
 - `clsetup` failure, 141
 - `scinstall` failure, 141
 - vote count, 141
 - updating after adding a node, 134
 - verifying, 142
 - ZFS restriction, 30
 - quorum servers, 55
 - See also* quorum devices
 - `/etc/scqsd/scqsd.conf` file, 56
 - configuring as quorum devices, 137
 - guidelines, 20

- installation directory, 57
- installing quorum server software, 55
- non-global zones, 20
- packages installed, 57
- removing, 230
- requirements as a quorum device, 137
- starting, 56
- troubleshooting, 57
- uninstalling, 230

quotas

- restriction on cluster file systems, 37

R

rac_framework resource type

- instantiating, 174
- registering, 174

RAID

- restriction, 41

raidctl command, 57

rarpd service

- restriction for cluster nodes, 22

raw-disk devices

- adding to a zone cluster, 225
- naming conventions, 166

rebooting

- into noncluster mode, 228

recovering

- unsuccessful cluster node creation, 80, 88, 128, 133

registering

- ScalDeviceGroup resource type, 178
- SUNW.rac_framework resource type, 174
- SUNW.ScalDeviceGroup resource type, 178
- SUNW.vucmm_framework resource type, 175
- SUNW.vucmm_svm resource type, 175

remote login

- Trusted Extensions, 189

removing, 227, 227

- See also* unconfiguring
- See also* uninstalling
- Oracle Solaris Cluster software, 227
- partially configured node, 227
- quorum servers, 230
- zone clusters, 231

repairing

- mediator data, 163

replicated devices

- disk requirement, 37
- restriction as quorum devices, 30
- setting the replication property, 155

replicating

- a zone cluster
 - from unified archives, 108
- a zone cluster node
 - from unified archives, 108

resource groups

- making unmanaged, 232
- multiple-owner volume-manager framework, 175
- scalable device group
 - affinities, 178

resource types

- ScalDeviceGroup
 - dependencies, 179
 - instantiating, 178
 - registering, 178
- SUNW.crs_framework, 174
- SUNW.rac_framework, 174, 174
- SUNW.ScalDeviceGroup
 - dependencies, 179
 - instantiating, 178
 - registering, 178
- SUNW.vucmm_framework, 175, 175
- SUNW.vucmm_svm, 175, 175
 - dependencies, 175

resources

- disabling, 231, 232
- listing, 231

restarting

- NTP, 145

restrictions

- live migration, 32
- lofi device, 14
- virtual disk multipathing, 32

root environment

- configuring, 65

root pool mirroring

- planning, 43

routers

- restriction for cluster nodes, 21

- RPC service
 - disabling TCP wrappers, 76, 82, 124, 131
 - enabling TCP wrappers, 126, 132
 - restoring external access, 61
 - restricted program numbers, 22

- S**
- SATA disks
 - configuring as quorum devices, 30
- scalable data services
 - IP Filter restriction, 66
- scalable device groups
 - resource group affinities, 178
- ScalDeviceGroup resource type
 - dependencies, 179
 - instantiating, 178
 - registering, 178
- scinstall command
 - adding nodes, 122
 - creating the global cluster, 71
 - creating the global cluster by using Automated Installer, 89
 - unconfiguring Oracle Solaris Cluster software, 227
- scope property, 194
- SCSI devices
 - correcting reservations after adding a third node, 134
 - fencing protocol setting in quorum devices, 30
- semaphore IDs
 - zone clusters, 194
- shared devices
 - installing quorum devices, 137
- shared IP type
 - zone clusters, 34
- shared memory
 - zone clusters, 194
- shared memory IDs
 - zone clusters, 194
- single-host clusters
 - configuring cluster interconnects, 115
 - naming, 23
- single-node clusters *See* single-host clusters
- SMF
 - verifying online services, 78, 84, 126
- software quorum protocol, 30
- software RAID
 - restriction, 41
- software updates
 - planning, 17
 - zone clusters, 203
- solaris brand zone clusters
 - options, 197
 - planning, 32
- Solaris Volume Manager
 - configuring, 151
 - disk set reserved names, 42
 - disk sets
 - adding drives, 155
 - adding to a zone cluster, 224
 - configuring, 153
 - repartitioning drives, 157
 - dual-string mediators
 - adding hosts, 162
 - overview, 161
 - repairing bad data, 163
 - status, 163
 - installing the software, 151
 - md.tab file, 158
 - mediators *See* dual-string mediators
 - planning, 42
 - raw-disk device names, 166
 - state database replicas, 152
 - unique naming, 42
 - volumes
 - activating, 160
- solaris10 brand zone clusters
 - options, 197
 - planning, 32
 - supported versions, 197
 - zone root password, 194
- spanning tree algorithm
 - disabling for quorum servers, 20
- starting
 - clsetup utility, 115
 - NTP, 145, 146
 - pconsole utility, 54, 61, 98, 106, 111
 - quorum servers, 56
 - scinstall utility, 77
- state database replicas
 - configuring, 152

- status
 - dual-string mediators, 163
 - verifying, 142
- storage devices
 - adding global storage devices to a zone cluster, 221
 - adding raw-disk devices to a zone cluster, 225
 - adding to zone clusters, 220
- Sun Explorer software *See* Oracle Explorer software
- Sun QFS
 - adding shared file systems to a zone cluster, 214
- Sun QFS software
 - installing the software, 65
 - limited use, 37
 - mount options, 211
 - restriction, 41
- Sun StorageTek Availability Suite *See* Availability Suite feature of Oracle Solaris software
- SUNW.crs_framework resource type
 - registering, 174
- SUNW.rac_framework resource group
 - configuring for Oracle ACFS, 174
- SUNW.rac_framework resource type
 - instantiating, 174
 - registering, 174
- SUNW.ScalDeviceGroup resource type
 - dependencies, 179
 - instantiating, 178
 - registering, 178
- SUNW.vucmm_framework resource type
 - instantiating, 175
 - registering, 175
- SUNW.vucmm_svm resource type
 - dependencies, 175
 - instantiating, 175
 - registering, 175
- supported configurations, 47
- swap
 - memory for zone clusters, 195
 - planning, 14
 - swap property, 195
- switches
 - planning, 28
- sysid configuration for zone clusters, 198

T

- tagged VLAN adapters
 - cluster interconnect guidelines, 27
 - public-network guidelines, 18
- TCP wrappers for RPC
 - disabling, 76, 82, 124, 131
 - enabling, 126, 132
 - modifying `/etc/hosts.allow`, 79, 85, 101, 107
- three-way mirroring, 43
- tncfg template, 190
- transport adapters *See* adapters
- transport switches
 - planning, 28
- troubleshooting
 - Automated Installer installation, 102
 - configuring
 - additional nodes, 128
 - new global clusters, 80, 88, 128, 133
 - explorer baseline record of the configuration, 148
 - NAS devices in a zone cluster, 192
 - quorum devices
 - clsetup failure, 141
 - scinstall failure, 141
 - vote count, 141
 - quorum server installation, 57
 - removing partially configured node, 230
- Trusted Extensions
 - admin_low template, 190
 - cipso template
 - global cluster, 190
 - zone cluster, 202
 - configuring, 188
 - configuring a zone cluster, 201
 - guidelines, 35
 - installation, 189
 - installing, 188
 - IP strict multihoming, 202
 - labeled brand, 32
 - LDAP, 190
 - password files, 197
 - remote login, 189
 - tncfg template, 190
 - txzonemgr GUI, 197
 - zonestare script, 189

U

- UFS logging
 - planning, 43
 - unconfiguring, 227, 227
 - See also* removing
 - See also* uninstalling
 - Oracle Solaris Cluster software, 227
 - zone clusters, 231
 - unified archives
 - archiveadm, 102
 - replicating a zone cluster, 108
 - uninstalling, 227, 227
 - See also* removing
 - See also* unconfiguring
 - Availability Suite feature of Oracle Solaris software, 233
 - Oracle Solaris Cluster software, 227
 - quorum servers, 230
 - unique naming
 - Solaris Volume Manager, 42
 - user-initialization files
 - modifying, 65
- V**
- validating the cluster, 148
 - verifying
 - automatic reboot on disk-path failures, 78
 - cldevice command processing, 154
 - cluster configuration, 148
 - cluster node status, 78, 84, 100, 107, 113, 126
 - installation mode, 143
 - Oracle Solaris Cluster software installation, 142
 - private hostnames, 144
 - quorum configurations, 142
 - SMF, 78, 84, 126
 - vfstab configuration, 167
 - vfstab file
 - adding mount points, 166
 - verifying the configuration, 167
 - virtual disk multipathing
 - restriction, 32
 - VLAN adapters
 - cluster interconnect guidelines, 27
 - private-network guidelines, 25

- public-network guidelines, 18
 - volume managers, 41
 - See also* Solaris Volume Manager
 - partitions for, 15
 - planning
 - general, 41
 - Solaris Volume Manager, 42
 - volumes
 - Solaris Volume Manager
 - activating, 160
 - vucmm_framework resource type
 - instantiating, 175
 - registering, 175
 - vucmm_svm resource type
 - dependencies, 175
 - instantiating, 175
 - registering, 175
- Z**
- ZFS root pools
 - internal disk mirroring, 57
 - planning mirroring, 43
 - ZFS storage pools
 - adding a local pool to a zone-cluster node, 219
 - adding to a zone cluster, 207, 209
 - restrictions, 30
 - zone clusters
 - adding an Oracle ACFS file system, 216
 - adding devices
 - DID devices, 223
 - disk sets, 224
 - global storage devices, 221
 - raw-disk devices, 225
 - adding file systems, 203
 - cluster file systems, 210, 213
 - highly available local file systems, 204
 - highly available ZFS storage pools, 207
 - local file systems, 205
 - local file systems for a specific zone-cluster node, 217
 - local ZFS storage pools, 219
 - Sun QFS, 214
 - ZFS storage pools, 207, 209
 - adding local file systems, 217
 - adding storage devices, 220

- adding to the private IP-address range, 25
- address property, 195
- brand property, 193
- brands, 32
- capped-cpu, 194
- configuring, 188, 188
- configuring to use Trusted Extensions, 201
- creating, 190
- dedicated-cpu, 194
- default IP type, 34
- default router, 195
- defrouter property, 195
- enable_priv_net property, 193
- Fair Share Scheduler shares, 194
- halting, 232
- hostmodel property, 202
- hostname property, 195
- IP addresses, 18
- IP strict multihoming, 202
- ip-type property, 193
- labeled brand options, 197
- lightweight processes, 194
- locked memory, 195
- locked property, 195
- message queue IDs, 194
- mounting file systems with HAStoragePlus, 34
- naming, 34
- ncpus property, 194
- num_xip_zoneclusters property, 193
- Oracle ACFS file systems, 176
- physical memory, 195
- physical property, 195, 195
- planning, 32
- scope property, 194
- semaphore IDs, 194
- setting the Global_zone resource-type property, 34
- shared memory, 194
- shared memory IDs, 194
- software updates, 203
- solaris brand options, 197
- solaris10 brand
 - supported versions, 197
 - zone root password, 194
- solaris10 brand options, 197
- swap memory, 195
- swap property, 195
- sysid configuration, 198
- unconfiguring, 231
- zonepath property, 193
- zone root password
 - solaris10 brand zone clusters, 194
- zonepath property, 193
- zoneshare Trusted Extension script, 189