

Installing and Configuring an Oracle Solaris Cluster 4.4 Environment



E69313-04
July 2023



Installing and Configuring an Oracle Solaris Cluster 4.4 Environment,

E69313-04

Copyright © 2000, 2023, Oracle and/or its affiliates.

Primary Author: Cathleen Reiher, Sharon Veach, Kristi Caruso

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, software documentation, data (as defined in the Federal Acquisition Regulation), or related documentation that is delivered to the U.S. Government or anyone licensing it on behalf of the U.S. Government, then the following notice is applicable:

U.S. GOVERNMENT END USERS: Oracle programs (including any operating system, integrated software, any programs embedded, installed, or activated on delivered hardware, and modifications of such programs) and Oracle computer documentation or other Oracle data delivered to or accessed by U.S. Government end users are "commercial computer software," "commercial computer software documentation," or "limited rights data" pursuant to the applicable Federal Acquisition Regulation and agency-specific supplemental regulations. As such, the use, reproduction, duplication, release, display, disclosure, modification, preparation of derivative works, and/or adaptation of i) Oracle programs (including any operating system, integrated software, any programs embedded, installed, or activated on delivered hardware, and modifications of such programs), ii) Oracle computer documentation and/or iii) other Oracle data, is subject to the rights and limitations specified in the license contained in the applicable contract. The terms governing the U.S. Government's use of Oracle cloud services are defined by the applicable contract for such services. 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®, Java, and MySQL are registered trademarks of Oracle and/or its affiliates. Other names may be trademarks of their respective owners.

Intel and Intel Inside 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, Epyc, and the AMD 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 about 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 unless otherwise set forth in an applicable agreement between you and Oracle. 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, except as set forth in an applicable agreement between you and Oracle.

Copyright © 2000, 2023, Oracle et/ou ses affiliés.

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 stipulation expresse de votre contrat de licence ou de la loi, vous ne pouvez pas copier, reproduire, traduire, diffuser, modifier, accorder de licence, 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, la documentation du logiciel, les données (telles que définies dans la réglementation "Federal Acquisition Regulation") ou la documentation qui l'accompagne sont livrés sous licence au Gouvernement des États-Unis, ou à quiconque qui aurait souscrit la licence de ce logiciel pour le compte du Gouvernement des États-Unis, la notice suivante s'applique :

UTILISATEURS DE FIN DU GOUVERNEMENT É.-U. : programmes Oracle (y compris tout système d'exploitation, logiciel intégré, tout programme intégré, installé ou activé sur le matériel livré et les modifications de tels programmes) et documentation sur l'ordinateur d'Oracle ou autres logiciels Oracle Les données fournies aux utilisateurs finaux du gouvernement des États-Unis ou auxquelles ils ont accès sont des "logiciels informatiques commerciaux", des "documents sur les logiciels informatiques commerciaux" ou des "données relatives aux droits limités" conformément au règlement fédéral sur l'acquisition applicable et aux règlements supplémentaires propres à l'organisme. À ce titre, l'utilisation, la reproduction, la duplication, la publication, l'affichage, la divulgation, la modification, la préparation des œuvres dérivées et/ou l'adaptation des i) programmes Oracle (y compris tout système d'exploitation, logiciel intégré, tout programme intégré, installé, ou activé sur le matériel livré et les modifications de ces programmes), ii) la documentation informatique d'Oracle et/ou iii) d'autres données d'Oracle, sont assujetties aux droits et aux limitations spécifiés dans la licence contenue dans le contrat applicable. Les conditions régissant l'utilisation par le gouvernement des États-Unis des services en nuage d'Oracle sont définies par le contrat applicable à ces services. Aucun autre droit n'est accordé au gouvernement américain.

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 un risque de 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 des applications dangereuses.

Oracle®, Java, et MySQL sont des marques déposées d'Oracle Corporation et/ou de ses affiliés. Tout autre nom mentionné peut être une marque appartenant à un autre propriétaire qu'Oracle.

Intel et Intel Inside 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, Epyc, et le logo AMD sont des marques ou des marques déposées d'Advanced Micro Devices. UNIX est une marque déposée de 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 et excluent toute garantie expresse ou implicite quant aux contenus, produits ou services émanant de tiers, sauf mention contraire stipulée dans un contrat entre vous et Oracle. 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, sauf mention contraire stipulée dans un contrat entre vous et Oracle.

Contents

Using This Documentation

Product Documentation Library	ix
Feedback	ix

1 Planning the Oracle Solaris Cluster Configuration

Finding Oracle Solaris Cluster Installation Tasks	1-1
Planning the Oracle Solaris OS	1-2
Guidelines for Selecting Your Oracle Solaris Installation Method	1-2
Oracle Solaris OS Feature Requirements and Restrictions	1-2
System Disk Partitions	1-3
Volume Management Software	1-4
Planning the Oracle Solaris Cluster Environment	1-5
Oracle Solaris Cluster Software Version	1-5
Memory Requirements	1-5
Licensing	1-6
Software Updates	1-6
Disaster Recovery Framework	1-6
Public-Network IP Addresses	1-6
Console-Access Devices	1-7
Public Network Configuration	1-7
Quorum Server Configuration	1-9
NFS Guidelines	1-9
Service Restrictions	1-10
Network Time Protocol (NTP)	1-11
Oracle Solaris Cluster Configurable Components	1-12
Global-Cluster Name	1-12
Global-Cluster Node Names and Node IDs	1-12
Private Network Configuration	1-13
Private Hostnames	1-15
Cluster Interconnect	1-15
Global Fencing	1-17
Quorum Devices	1-18

SPARC: Guidelines for Oracle VM Server for SPARC Logical Domains as Cluster Nodes	1-19
Zone Clusters	1-21
Global-Cluster Requirements and Guidelines	1-21
Zone-Cluster Requirements and Guidelines	1-22
Guidelines for Trusted Extensions in a Zone Cluster	1-24
Oracle Solaris Cluster Manager	1-24
Planning Global Devices, Device Groups, and Cluster File Systems	1-25
Planning Global Devices	1-25
Planning Device Groups	1-25
Planning Cluster File Systems	1-26
Choosing Options for ZFS Cluster File Systems	1-27
Property Settings for ZFS File System Datasets and ZFS Storage Pools	1-27
Encrypting ZFS Cluster File Systems	1-29
Choosing Mount Options for UFS Cluster File Systems	1-30
Mount Information for Cluster File Systems	1-31
Planning Volume Management	1-32
Guidelines for Volume Manager Software	1-32
Guidelines for Solaris Volume Manager Software	1-33
UFS Cluster File System Logging	1-33
Mirroring Guidelines	1-34
Guidelines for Mirroring Multihost Disks	1-34
Guidelines for Mirroring the ZFS Root Pool	1-34

2 Installing Software on Global-Cluster Nodes

Overview of Installing the Software	2-1
Installation Considerations	2-2
Installing the Software	2-2
How to Prepare for Cluster Software Installation	2-3
How to Install Oracle Solaris Software	2-4
How to Install pconsole Software on an Administrative Console	2-8
How to Install and Configure Oracle Solaris Cluster Quorum Server Software	2-10
How to Configure Internal Disk Mirroring	2-13
How to Install Oracle VM Server for SPARC Software and Create Domains (SPARC Only)	2-13
How to Install Oracle Solaris Cluster Software Using Centralized Installer (clinstall)	2-14
How to Install Oracle Solaris Cluster Software (pkg)	2-19
How to Install Oracle HSM Software	2-24
How to Set Up the Root Environment	2-24

3 Establishing the Global Cluster

Overview of Establishing a New Cluster or Cluster Node	3-1
Establishing a New Global Cluster or New Global-Cluster Node	3-3
Configuring Oracle Solaris Cluster Software on All Nodes (scinstall)	3-4
How to Configure Oracle Solaris Cluster Software on All Nodes (scinstall)	3-7
How to Configure Oracle Solaris Cluster Software on All Nodes (XML)	3-12
Establishing a New Cluster With the Automated Installer	3-19
How to Install and Configure Oracle Solaris and Oracle Solaris Cluster Software (IPS Repositories)	3-26
How to Install and Configure Oracle Solaris and Oracle Solaris Cluster Software (Unified Archives)	3-32
How to Replicate a Cluster from the Unified Archives	3-37
Establishing a New Logical-Domain Cluster by Deploying the Oracle VM Server for SPARC Template for Oracle Solaris Cluster	3-42
How to Deploy the Oracle VM Server for SPARC Template for Oracle Solaris Cluster to Configure a Logical Domain Cluster	3-42
How to Prepare the Cluster for Additional Global-Cluster Nodes	3-47
How to Change the Private Network Configuration When Adding Nodes or Private Networks	3-49
Configuring Oracle Solaris Cluster Software on Additional Global-Cluster Nodes (scinstall)	3-53
How to Configure Oracle Solaris Cluster Software on Additional Global-Cluster Nodes (scinstall)	3-54
How to Configure Oracle Solaris Cluster Software on Additional Global-Cluster Nodes (XML File)	3-59
How to Add a New Logical-Domain Node to an Existing Cluster by Using the Oracle VM Server for SPARC Template for Oracle Solaris Cluster	3-63
How to Update Quorum Devices After Adding a Node to a Global Cluster	3-64
How to Configure Quorum Devices	3-66
How to Verify the Quorum Configuration and Installation Mode	3-71
How to Change Private Hostnames	3-72
Configuring Network Time Protocol (NTP)	3-73
How to Use Your Own /etc/inet/ntp.conf File	3-73
How to Install NTP After Adding a Node to a Single-Node Cluster	3-74
How to Update NTP After Changing a Private Hostname	3-75
Configuring Precision Time Protocol (PTP) Instead of NTP	3-75
How to Configure PTP (Oracle Solaris PTP)	3-76
How to Configure PTP (Third Party PTP)	3-76
Securing the Interconnect for Oracle Solaris Cluster With IPsec	3-76
How to Protect Oracle Solaris Cluster Interconnect with IPsec	3-77
How to Validate the Cluster	3-78

4 Configuring Solaris Volume Manager Software

Configuring Solaris Volume Manager Software	4-1
How to Install Solaris Volume Manager	4-1
How to Create State Database Replicas	4-2
Creating Disk Sets in a Cluster	4-2
How to Create a Disk Set	4-3
Adding Drives to a Disk Set	4-5
How to Add Drives to a Disk Set	4-5
How to Repartition Drives in a Disk Set	4-6
How to Create an md.tab File	4-7
How to Activate Volumes	4-8
Configuring Dual-String Mediators	4-9
Requirements for Dual-String Mediators	4-10
How to Add Mediator Hosts	4-10
How to Check For and Fix Bad Mediator Data	4-11

5 Creating a Cluster File System

Creating Cluster File Systems	5-1
How to Create UFS Cluster File Systems	5-1
Creating ZFS Cluster File Systems	5-4
Creating ZFS Cluster File Systems With and Without Using HAStoragePlus	5-4
Configuring Encryption for a Globally Mounted ZFS File System	5-4
Creating Oracle ACFS File Systems	5-8
Sample Configurations of an Oracle ACFS File System	5-9
How to Register and Configure the Framework Resource Groups	5-13
How to Create an Oracle ACFS File System	5-16
How to Register and Configure the Scalable Device-Group Resource Group	5-17
How to Register and Configure the Oracle ASM Resource Group	5-18
How to Register and Configure the Oracle ACFS Proxy Resource Group	5-20
How to Create an Oracle Grid Infrastructure Resource for Interoperation With Oracle Solaris Cluster	5-22

6 Creating Zone Clusters

Overview of Creating and Configuring a Zone Cluster	6-1
Creating and Configuring a Zone Cluster	6-2
Creating a Zone Cluster	6-2
How to Install and Configure Trusted Extensions	6-3

How to Create a Zone Cluster (clsetup)	6-5
How to Create a solaris10 Brand Zone Cluster (CLI)	6-17
How to Configure a Zone Cluster to Use Trusted Extensions	6-22
How to Configure a Zone Cluster to Be Immutable	6-23
Importing Zones into Zone Clusters	6-24
Requirements and Restrictions	6-25
How to Import an Installed Zone into an existing Zone Cluster	6-25
How to Import an Installed Zone into a New Zone Cluster	6-26
Adding File Systems to a Zone Cluster	6-28
How to Add a Highly Available Local File System to a Zone Cluster (clsetup)	6-28
How to Add a Highly Available Local File System to a Zone Cluster (CLI)	6-30
How to Add a ZFS Storage Pool to a Zone Cluster (clsetup)	6-32
How to Add a ZFS Storage Pool to a Zone Cluster (CLI)	6-33
How to Add a Cluster File System to a Zone Cluster (clsetup)	6-34
How to Add a ZFS Cluster File System to a Zone Cluster (CLI)	6-37
How to Add a UFS Cluster File System to a Zone Cluster (CLI)	6-39
How to Add a Oracle HSM Shared File System to a Zone Cluster (CLI)	6-41
How to Add an Oracle ACFS File System to a Zone Cluster (CLI)	6-42
Adding Local File Systems to a Specific Zone-Cluster Node	6-43
How to Add a Local File System to a Specific Zone-Cluster Node (CLI)	6-43
How to Add a Local ZFS Storage Pool to a Specific Zone-Cluster Node (CLI)	6-45
Adding Storage Devices to a Zone Cluster	6-46
How to Add a Global Storage Device to a Zone Cluster (clsetup)	6-46
How to Add a DID Device to a Zone Cluster (CLI)	6-48
How to Add a Disk Set to a Zone Cluster (CLI)	6-49
How to Add a Raw-Disk Device to a Specific Zone-Cluster Node (CLI)	6-51

7 Uninstalling Software From the Cluster

Uninstalling the Software	7-1
How to Unconfigure Oracle Solaris Cluster Software to Correct Installation Problems	7-1
How to Uninstall Oracle Solaris Cluster Software From a Cluster Node	7-4
How to Unconfigure a Zone Cluster	7-6
How to Uninstall Oracle Solaris Cluster Quorum Server Software	7-8
How to Uninstall Cluster Software Packages from an Automated Installer Server	7-8
How to Uninstall a Specific Oracle Solaris Cluster Package	7-9

Index

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.

Product Documentation Library

Documentation and resources for this product and related products are available at http://docs.oracle.com/cd/E69294_01.

Feedback

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

1

Planning the Oracle Solaris Cluster Configuration

This chapter provides planning information and guidelines specific to an Oracle Solaris Cluster 4.4 configuration.

This chapter contains the following overview information:

- [Finding Oracle Solaris Cluster Installation Tasks](#)
- [Planning the Oracle Solaris OS](#)
- [Planning the Oracle Solaris Cluster Environment](#)
- [Planning Global Devices, Device Groups, and Cluster File Systems](#)
- [Planning Volume Management](#)

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.	Managing Hardware With Oracle Solaris Cluster 4.4 Documentation that shipped with your server and storage devices
Plan global-cluster software installation.	Planning the Oracle Solaris Cluster Configuration
Install software packages. Optionally, install and configure Oracle HSM software.	Installing the Software Oracle HSM documentation
Establish a new global cluster or a new global-cluster node.	Establishing a New Global Cluster or New Global-Cluster Node
Configure Solaris Volume Manager software.	Configuring Solaris Volume Manager Software Solaris Volume Manager Administration Guide
Configure cluster file systems, if used.	How to Create UFS Cluster File Systems
<i>(Optional)</i> Create zone clusters.	Creating and Configuring a Zone Cluster
Plan, install, and configure resource groups and data services. Create highly available local file systems, if used.	Planning and Administering Data Services for Oracle Solaris Cluster 4.4
Develop custom data services.	Developing Data Services

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](#)
- [Oracle Solaris OS Feature Requirements and Restrictions](#)
- [System Disk Partitions](#)

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 Oracle Solaris 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 \(IPS Repositories\)](#) 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:

- **VNIC** – When using VNICs for zone clusters, the name of the VNIC must be less than 16 characters long. Otherwise, you might not be able to choose the VNIC during system configuration.
- **Minimum Oracle Solaris Package** – Oracle Solaris Cluster software requires at least the `solaris-minimal-server` group package of Oracle Solaris software. To use the GUI form of the `clinstall` utility to perform a centralized installation of the Oracle Solaris Cluster software, the `solaris-desktop` group package must be installed.
- **Oracle Solaris Operating System** – Oracle Solaris Cluster 4.4 core software and Oracle Solaris Cluster 4.4 Quorum Server software require a minimum of Oracle Solaris 11.4 software.
- **Required Java Version** – Java 1.8 is required by Oracle Solaris 11.4 software and Oracle Solaris Cluster 4.4 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(8)` man page for more information.
- **Packet Filter (PF) feature** – When IP network multipathing (IPMP) is used for public network, Oracle Solaris Cluster relies on IPMP for public network monitoring. Any PF configuration must be made in accordance with IPMP configuration guidelines and restrictions concerning PF.
- **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 `fssnap` command is not supported on local file systems under the control of HAStoragePlus.
- **hostmodel Parameter for IPv4 or IPv6** - On a global cluster and in zone cluster nodes, set IP strict multihoming to `weak` or `src-priority`. See [hostmodel Parameter \(IPv4 or IPv6\) in Oracle Solaris 11.4 Tunable Parameters Reference Manual](#).
For example, the following commands set IPv4 and IPv6 strict multihoming to `weak`:

```
# ipadm set-prop -p hostmodel=weak ipv4
# ipadm set-prop -p hostmodel=weak ipv6
```

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 Mbytes 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.

 **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.

- **Solaris volume manager** – Create a 20- Mbytes partition on slice 6 for Solaris 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](#)
- [Memory Requirements](#)
- [Licensing](#)
- [Software Updates](#)
- [Disaster Recovery Framework](#)
- [Public-Network IP Addresses](#)
- [Console-Access Devices](#)
- [Public Network Configuration](#)
- [Quorum Server Configuration](#)
- [NFS Guidelines](#)
- [Service Restrictions](#)
- [Network Time Protocol \(NTP\)](#)
- [Oracle Solaris Cluster Configurable Components](#)
- [SPARC: Guidelines for Oracle VM Server for SPARC Logical Domains as Cluster Nodes](#)
- [Zone Clusters](#)
- [Oracle Solaris Cluster Manager](#)

For detailed information about Oracle Solaris Cluster components, see the [Concepts for Oracle Solaris Cluster 4.4](#).

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 4.4 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.

Disaster Recovery Framework

If a zone cluster will be configured in an Oracle Solaris Cluster Disaster Recovery Framework (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 Disaster Recovery Framework configuration's partner cluster.
- Each zone-cluster 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 browser interface to administer Disaster Recovery Framework components, all cluster nodes must have the same root password. For more information about Oracle Solaris Cluster Manager, see [Chapter 12, Using the Oracle Solaris Cluster Manager Browser Interface in *Administering an Oracle Solaris Cluster 4.4 Configuration*](#).

Public-Network IP Addresses

For information about the use of public networks by the cluster, see [Public Network Adapters in *Concepts for Oracle Solaris Cluster 4.4*](#).

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
(Optional) Non-global zones	1 IP address per subnet
Console-access device	1 IP address
Logical addresses	1 IP address per logical host resource, per subnet
Disaster Recovery Framework	1 IP address per global cluster/zone cluster, where the framework runs

For more information about planning IP addresses, see [Planning for Network Deployment in Oracle Solaris 11.4](#).

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 [Concepts for Oracle Solaris Cluster 4.4](#).

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 [Concepts for Oracle Solaris Cluster 4.4](#).
- **IPv4** – Oracle Solaris Cluster software supports IPv4 addresses on the public network.
- **IPv6** – The Disaster Recovery Framework does not support IPv6 addresses on the public network. 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 a list of data services that support IPv6 addresses on the public network.
- **Public Network Management** – Each public-network adapter that is used for data-service traffic must belong to a PNM object that includes IPMP groups, link aggregations, and VNICs that are directly backed by link aggregations. If a public-network adapter is not used for data-service traffic, you do not have to configure it in a PNM object.

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.

 **Note:**

IPMP groups are created only on unused physical adapters.

You can 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 11.4](#). To modify IPMP groups after cluster installation, follow the guidelines in [How to Administer IP Network Multipathing Groups in a Cluster in Administering an Oracle Solaris Cluster 4.4 Configuration](#) and procedures in [Chapter 3, Administering IPMP in Administering TCP/IP Networks, IPMP, and IP Tunnels in Oracle Solaris 11.4](#). For more information on link aggregations, see [Chapter 2, Configuring High Availability by Using Link Aggregations in Managing Network Datalinks in Oracle Solaris 11.4](#).

- **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 [Concepts for Oracle Solaris Cluster 4.4](#).

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 4.4 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** – A cluster node cannot be an NFS client of an Oracle Solaris Cluster HA for NFS exported file system when the underlying file system is UFS. However, you can perform this type of cross mounting when the underlying file system is ZFS.

This cross-mounting restriction applies to a client node and an NFS server in the same cluster. Instead, use the cluster file system to share files between cluster nodes.

A cluster node can be an NFS client of a qualified NAS device.

- **HA-NFS as cluster-shared storage** – You cannot use HA-NFS as a cluster-shared storage method for running HA data services on cluster nodes that mount the exported NFS file systems. You cannot use this method because there is no mechanism to release NFS file locks and perform NFS fencing against a possibly faulted NFS client cluster node. As a result, the data service might not be able to start on a new node.

This restriction applies when the data services are on the same or a different cluster than the HA-NFS cluster.

This restriction only applies when the HA-NFS-exported file systems are used as cluster-shared storage. So, you can use this configuration for user home directories, offline copy of data, and so on.

If you must use NFS as a shared-storage method for HA data services, you must use a qualified NAS device.

- **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(8)` 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.
- **IP Forwarding** – While Solaris Cluster does not support Solaris cluster nodes as routers (gateways), forwarding can be selectively enabled for network interfaces other than private interconnect and `clprivnet` interfaces.
 - Do not enable forwarding for private interconnect interfaces or `clprivnet` interfaces. Solaris Cluster software will automatically turn off IP forwarding for private interconnect interfaces and `clprivnet` interfaces upon booting of cluster nodes.
 - Do not use IP forwarding within the cluster host for an external client to connect to an application instance using the Shared Address resource.
- **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 [Concepts for Oracle Solaris Cluster 4.4](#) for further information about cluster time. For more information about NTP, see the [ntpd\(8\)](#) man page.

Oracle Solaris Cluster Configurable Components

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

- [Global-Cluster Name](#)
- [Global-Cluster Node Names and Node IDs](#)
- [Private Network Configuration](#)
- [Private Hostnames](#)
- [Cluster Interconnect](#)
- [Global Fencing](#)
- [Quorum Devices](#)

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](#).

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\(5\)](#) 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](#).

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, 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 *Administering an Oracle Solaris Cluster 4.4 Configuration*](#). 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 11.4](#) 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(8CL)` 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 Concepts for Oracle Solaris Cluster 4.4](#).



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](#).
- 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 *Managing Hardware With Oracle Solaris Cluster 4.4*](#). For general information about the cluster interconnect, see [Cluster Interconnect in *Concepts for Oracle Solaris Cluster 4.4*](#).

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 *Managing Hardware With Oracle Solaris Cluster 4.4*](#). For information about creating and administering VLANs, see the `dladm(8)` man page and [Chapter 4, Configuring Virtual Networks by Using Virtual Local Area Networks in *Managing Network Datalinks in Oracle Solaris 11.4*](#).

- **SPARC: Oracle VM Server for SPARC guest domains** – For Oracle VM Server for SPARC guest domains that are configured as cluster nodes, 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 `switch N`, 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 Concepts for Oracle Solaris Cluster 4.4](#). For more information about setting the fencing protocol of individual shared disks, see the [cldevice\(8CL\)](#) man page. For more information about the global fencing setting, see the [cluster\(8CL\)](#) 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 Concepts for Oracle Solaris Cluster 4.4](#).

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 Concepts for Oracle Solaris Cluster 4.4](#).

SPARC: Guidelines for Oracle VM Server for SPARC Logical Domains as Cluster Nodes

Consider the following points when you create an Oracle VM Server for SPARC logical domain or guest domain on a physically clustered machine that is SPARC hypervisor capable, for use as a cluster node:

- **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 cluster node on the same physical machine unless you also disable fencing for that device from the guest domain nodes that you export the LUN to. Otherwise, the SCSI reservation on the device can be modified by one of the clusters and exclude nodes of other clusters. When that happens any other guest domain cluster nodes that subsequently try to access the device will panic. Such a situation could be triggered for example when one node of a cluster is halted. When you export a storage LUN to more than one guest domain cluster

node on the same physical machine, login to each guest domain cluster node that you have exported the LUN to and use `cldev` command to disable fencing of the device.

For example:

```
# cldev set -p default_fencing=nofencing d2"
```

- **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.
- **Live migration restriction** – The live migration and cold migration features of Oracle VM Server for SPARC are not supported for logical domains that are configured to run as cluster nodes.

However, guest domains that are configured to be managed by the HA for Oracle VM Server for SPARC data service do support cold migration and live migration. But no other type of logical domain in this data service configuration is supported to use cold migration or live migration.

- **mpgroup boot drive support** – Oracle Solaris Cluster supports `mpgroups` for boot drives with fencing disabled using the `cldev` command. You must disable fencing before configuring the `vdisk` for the boot drive in an `mpgroup`.

For example:

```
# cldev set -p default_fencing=nofencing d2
```

- **Virtual disk multipathing** – With the exception of boot drives, 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.

For more information about Oracle VM Server for SPARC, see the [Oracle VM Server for SPARC 3.6 Administration Guide](#).

Zone Clusters

A zone cluster is a cluster of Oracle Solaris non-global zones. You can use the `clsetup` utility or the Oracle Solaris Cluster Manager browser interface 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(8CL)` man page. For more information about Oracle Solaris Cluster Manager, see [Chapter 12, Using the Oracle Solaris Cluster Manager Browser Interface in Administering an Oracle Solaris Cluster 4.4 Configuration](#).

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](#)
- [Zone-Cluster Requirements and Guidelines](#)
- [Guidelines for Trusted Extensions in a Zone Cluster](#)

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*)
- **Solaris Cluster Manager browser** – All Global-cluster nodes must have the same root password if you want to use the Oracle Solaris Cluster Manager browser interface to administer Oracle Solaris Cluster components.

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. If the Disaster Recovery Framework will be enabled in the zone cluster, the cluster name must be suitable as a hostname, and must be in the naming system.
- **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](#).
- **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.
- **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
 - * Oracle HSM stand-alone file system
 - * Oracle HSM 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
 - * ZFS cluster file system
 - * Oracle HSM stand-alone file system
 - * Oracle HSM 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.

Oracle Solaris Cluster Manager

The Oracle Solaris Cluster manager can be installed as a standalone package without installing the Oracle Solaris Cluster software. The manager can then be used to access Oracle Solaris cluster nodes in global clusters on zone clusters to configure

and manage these clusters. For more information, see [Chapter 12, Using the Oracle Solaris Cluster Manager Browser Interface in *Administering an Oracle Solaris Cluster 4.4 Configuration*](#).

Planning Global Devices, Device Groups, and Cluster File Systems

This section provides the following information:

- [Planning Global Devices](#)
- [Planning Device Groups](#)
- [Planning Cluster File Systems](#)
- [Choosing Options for ZFS Cluster File Systems](#)
- [Choosing Mount Options for UFS Cluster File Systems](#)
- [Mount Information for Cluster File Systems](#)

Planning Global Devices

For information about the purpose and function of global devices, see [Global Devices in *Concepts for Oracle Solaris Cluster 4.4*](#).

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 *Concepts for Oracle Solaris Cluster 4.4*](#).

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](#) for additional guidelines. See [Configuring Solaris Volume Manager Software](#) 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 Concepts for Oracle Solaris Cluster 4.4](#).

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 Planning and Administering Data Services for Oracle Solaris Cluster 4.4](#).

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 ZFS or UFS cluster file systems directly in a zone cluster. However, you can configure them in the global cluster and loopback-mounted into the zone cluster; or you can use highly available local file systems. For further information, see [Adding File Systems to a Zone Cluster](#)
- **Zone file systems** – Zones cannot be installed on a global ZFS file system. Configuring the zone root path of a zone under a file system of a zpool for a globally mounted ZFS file system is not supported.
- **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.
- **Using `chmod` to change `setuid` permissions** – The `chmod` command might fail to change `setuid` permissions on a file in a cluster file system. If the `chmod` command is run on a non-global zone and the non-global zone is not on the PxFS primary server, the `chmod` command fails to change the `setuid` permission.

Use one of these methods to successfully change `setuid` permissions:

- Perform the operation on any global-cluster node that accesses the cluster file system.
- Perform the operation on any non-global zone that runs on the PxFS primary node that has a loopback mount to the cluster file system.
- Switch the PxFS primary to the global-cluster node where the non-global zone that encountered the error is running.

Choosing Options for ZFS Cluster File Systems

This section describes the ZFS properties that control ZFS file system datasets and ZFS storage pools (zpools), and encrypted ZFS cluster file systems.

- [Property Settings for ZFS File System Datasets and ZFS Storage Pools](#)
- [Encrypting ZFS Cluster File Systems](#)

Property Settings for ZFS File System Datasets and ZFS Storage Pools

This section describes requirements and restrictions for property settings of ZFS file system datasets and ZFS storage pools (zpools).

 **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 *Planning and Administering Data Services for Oracle Solaris Cluster 4.4*](#).

The following is an Oracle Solaris ZFS pool property.

```
clustered=on | off
```

For ZFS pools, the `clustered` property determines the accessibility of the datasets of the pool. If the pool is imported with the `clustered` property set to `on`, then all the file system datasets in that pool will be mounted globally and available from all the cluster nodes. Normally, a ZFS pool that hosts the ZFS cluster file systems is managed by a device group which automatically sets the `clustered` property to `on` at pool import time. Hence, you do not have to set this manually.

For more information about the `clustered` property, see the [`zpool\(8\)`](#) man page.

The following ZFS dataset properties are not allowed to be set to new values while the file system is globally mounted. To change any of these properties, unmount the ZFS dataset or mount it locally. For more information on these and other `zfs` properties, see the [`zfs\(8\)`](#) man page.

- `atime`
- `devices`
- `exec`
- `readonly`
- `rstchown`
- `setuid`
- `xattr`
- `sync`
- `canmount`
- `mountpoint`
- `zoned`

A ZFS file system must have its `zoned` property set to `off` for a global mount to succeed.

For example, to change the `mountpoint` property of dataset `globalzpool/fs1` where `globalzpool` has `clustered` property set to `on`, complete the following steps.

1. Unmount the dataset.

```
# zfs unmount globalzpool/fs1
```

2. Update the `mountpoint` property.

```
# zfs set mountpoint=/global/gzpool/fs1 globalzpool/fs1
```

3. Re-mount the dataset.

```
# zfs mount globalzpool/fs1
```

Sharing of a ZFS file system using the `zfs share` and `zfs unshare` commands with the `share.nfs` or `share.smb` ZFS property is not currently supported for globally mounted ZFS file systems. Instead, configure legacy sharing using the Oracle Solaris Cluster HA for Network File System (NFS) data service. For more information, see the [Oracle Solaris Cluster Data Service for NFS Guide](#).

Encrypting ZFS Cluster File Systems

You can encrypt ZFS cluster file systems in Oracle Solaris Cluster 4.4 and use such file systems in the global cluster and in zone clusters by using the virtual loopback file system, `lofs`. See the `lofs(4FS)` man page.

Requirements for ZFS Cluster File System Encryption

To use this feature, you must run at least the following software versions:

- Oracle Solaris 11.4 SRU 33
- Oracle Solaris Cluster 4.4 SRU 10

You can encrypt a ZFS cluster file system upon its creation. To ensure that the encryption key is available during Oracle Solaris Cluster failovers and switchovers, you must specify the cluster as the location of the encryption key in the `keysource` property value.

`keysource` ZFS Encryption Property

For ZFS cluster file systems, you must specify the value of the `keysource` property value. The value includes the format (`hex`) and location of the wrapping key (`cluster:keyname`) as follows:

```
keysource=hex,cluster:keyname
```

Where `hex`, `cluster` keywords indicate that the encryption key is a hexadecimal number (`hex`) and that ZFS must contact Oracle Solaris Cluster for the wrapping key (`cluster`). `keyname` specifies the name of the unique Oracle Solaris Cluster encryption key. You use the `cldevicegroup key-create` command to create the encryption key associated with `keyname`. For more information about the `keysource` property, see [Encrypting ZFS File Systems in Managing ZFS File Systems in Oracle Solaris 11.4](#).

Managing Encryption Keys

Each encrypted cluster file system requires a wrapping key, which persistently stores the encrypted data-encryption keys that ZFS manages. Oracle Solaris Cluster stores the wrapping key in the Cluster Configuration Repository (CCR). By using the CCR, the wrapping key is always available during a failover or switchover.

Maintaining Keys

Back up the encryption key information to ensure that encrypted cluster file systems are accessible. Backing up this information is important in cases where the keys in the CCR have been deleted inadvertently. If one or more keys becomes unavailable, any encrypted cluster file systems that use that key become inaccessible.

Keys might become unavailable if you perform any of the following actions:

- Delete a key
- Roll back to an Oracle Solaris Cluster release that does not support encryption
- Roll back to an Oracle Solaris Cluster release in which no keys are configured

To record encryption key information, write down the key name and its associated encryption key value or save this information to a file that only you can read.

For example, the following commands save the key name and its associated encryption key value to a file called `~/osck` that only you can read:

```
phys-schost# cldevicegroup key-list -v >~/osck
phys-schost# chmod 400 ~/osck
```

**Note:**

For more information about configuring encryption, see [Configuring Encryption for a Globally Mounted ZFS File System](#).

Choosing Mount Options for UFS Cluster File Systems

This section describes requirements and restrictions for mount options 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 Planning and Administering Data Services for Oracle Solaris Cluster 4.4](#).

Follow the guidelines in the following list of mount options to determine which 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\(8\)](#) 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** –By default, ZFS datasets are mounted under the zpool's root file system. For example, a ZFS pool named `pool1` will have its file system datasets mounted on `/pool1`, `/pool1/fs1`, `/pool1/fs2`, and so on.

In earlier releases of Oracle Solaris Cluster that did not support cluster file systems on ZFS, there was a convention of mounting cluster file systems under the `/global` directory to more easily distinguish them from the local file systems. With ZFS, it is possible to specify non-default mount points by executing the `zfs set` command on the ZFS dataset. However, the mountpoint of a ZFS file system dataset can be changed only while the dataset is not mounted globally. The ZFS command must be executed on the node where the ZFS pool is imported. To avoid dealing with non-default mountpoints, it is recommended to accept the default ZFS mounting behavior. For more information, see the `zfs(8)` man page.

- **Nesting mount points** – ZFS datasets within a given ZFS storage pool may be nested. For example, a pool named `globalpool` might have file system datasets named `globalpool/fs1`, `globalpool/fs2`, `globalpool/fs2/fs3`, and so on. These file system datasets will be mounted on the corresponding paths under the pool's root file system `/globalpool`.

Do not nest the mount points for non-ZFS 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.

An 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 Oracle HSM shared file systems even if the two file system devices have the same physical host connectivity.

Mounting a local filesystem or an HA local filesystem on a global filesystem is not supported. Use symbolic links from the global filesystem to the non-global filesystem instead.

- `forcedirectio` (UFS only) – 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](#)
- [Guidelines for Solaris Volume Manager Software](#)
- [UFS Cluster File System Logging](#)
- [Mirroring Guidelines](#)

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](#) 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 Concepts for Oracle Solaris Cluster 4.4](#) and [Device Groups in Concepts for Oracle Solaris Cluster 4.4](#).

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](#) 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](#) 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 [Planning and Administering Data Services for Oracle Solaris Cluster 4.4](#) 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\(4D\)](#) 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\(8\)](#) 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](#)
- [Guidelines for Mirroring the ZFS Root Pool](#)

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 Concepts for Oracle Solaris Cluster 4.4](#).

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 11.4](#) for instructions about how to mirror the ZFS root pool. Also see [Chapter 6, Managing the ZFS Root Pool in Managing ZFS File Systems in Oracle Solaris 11.4](#) 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.

2

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](#)
- [Installing the Software](#)

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
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
(Optional) Install <code>pconsole</code> software on an administrative console.	How to Install pconsole Software on an Administrative Console
(Optional) Install and configure a quorum server.	How to Install and Configure Oracle Solaris Cluster Quorum Server Software
(Optional) Configure internal disk mirroring.	How to Configure Internal Disk Mirroring
(Optional) Install Oracle VM Server for SPARC software and create domains.	How to Install Oracle VM Server for SPARC Software and Create Domains
Install Oracle Solaris Cluster software and any data services that you will use.	How to Install Oracle Solaris Cluster Software (pkg)
(Optional) Install Oracle HSM software.	How to Install Oracle HSM Software
Set up directory paths.	How to Set Up the Root Environment
(Optional) Configure the Packet Filter feature of Oracle Solaris.	How to Configure Packet Filter

▲ Caution:

The `/var/cluster.sav/` directory is created and used during current and future Oracle Solaris Cluster installation, update, and uninstallation operations. This directory and its contents are for Oracle Solaris Cluster internal use only, and must not be deleted by the user. User removal of this directory, or any of its contents, would pose a high risk of compromising the cluster during any future installation, update, or uninstallation of Oracle Solaris Cluster software.

Installation Considerations

The following table describes the file systems that work with the Oracle Solaris Cluster 4.4 software on SPARC and x86 platforms.

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

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](#)

- [How to Install Oracle Solaris Software](#)
- [How to Install pconsole Software on an Administrative Console](#)
- [How to Install and Configure Oracle Solaris Cluster Quorum Server Software](#)
- [How to Configure Internal Disk Mirroring](#)
- [How to Install Oracle VM Server for SPARC Software and Create Domains](#)
- [How to Install Oracle Solaris Cluster Software Using Centralized Installer \(clinstall\)](#)
- [How to Install Oracle Solaris Cluster Software \(pkg\)](#)
- [How to Install Oracle HSM Software](#)
- [How to Set Up the Root Environment](#)
- [How to Configure Packet Filter](#)

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 Concepts for Oracle Solaris Cluster 4.4](#) for information about physical and virtual machines that are supported as cluster nodes.
 - See [Oracle Solaris Cluster 4 Compatibility Guide \(http://www.oracle.com/technetwork/server-storage/solaris-cluster/overview/solariscluster4-compatibilityguide-1429037.pdf\)](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 4.4 Release Notes](#) – Restrictions, bug workarounds, and other late-breaking information.
 - [Concepts for Oracle Solaris Cluster 4.4](#) - 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.
 - [Planning and Administering Data Services for Oracle Solaris Cluster 4.4](#) – 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
 - Oracle HSM software
 - Third-party applications
- 4. Plan your cluster configuration.**

Use the planning guidelines in [Planning the Oracle Solaris Cluster Configuration](#) and in the [Planning and Administering Data Services for Oracle Solaris Cluster 4.4](#) 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 [Updating Your Oracle Solaris Cluster 4.4 Environment](#).

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](#).
- 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](#).
- Otherwise, choose the Oracle Solaris installation procedure to use.
 - To configure Oracle Solaris Cluster software by using the `scinstall(8)` utility, go to [How to Install Oracle Solaris Software](#) 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 \(IPS Repositories\)](#).

How to Install Oracle Solaris Software

Perform the following tasks:

- Ensure that the hardware setup is complete and that connections are verified before you install Oracle Solaris software. See the [Managing Hardware With Oracle Solaris Cluster 4.4](#) 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](#) 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](#) for planning guidelines. See your Oracle Solaris system administrator documentation for information about using Oracle Solaris naming services.

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

- (Optional) An administrative console that you will install with `pconsole` software. For more information, see [How to Install pconsole Software on an Administrative Console](#).
- (Optional) A quorum server. For more information, see [How to Install and Configure Oracle Solaris Cluster Quorum Server Software](#).
- 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 \(IPS Repositories\)](#).

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](#) for information about required root-disk partitioning and other Oracle Solaris Cluster installation requirements.

1. **Connect to the consoles of each node.**
2. **Install the Oracle Solaris OS.**

Follow installation instructions in [Automatically Installing Oracle Solaris 11.4 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](#).
- 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 browser interface to administer components of the disaster recovery framework, formerly called Geographic Edition. For more information about Oracle Solaris Cluster Manager, see [Chapter 12, Using the Oracle Solaris Cluster Manager Browser Interface in Administering an Oracle Solaris Cluster 4.4 Configuration](#).

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 [Updating Systems and Adding Software in Oracle Solaris 11.4](#).

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 11.4](#) 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 [Updating Your Oracle Solaris Cluster 4.4 Environment](#).

7. **(x86 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 11.4 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. (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 11.4](#) 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. (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(8)` 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 I/O Multipathing in Oracle Solaris 11.4](#) and the `stmsboot(8)` 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](#).

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

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](#).

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](#).

Alternatively, to use the Oracle VM Server for SPARC template for Oracle Solaris Cluster to install and configure a cluster of guest domains or I/O domains, go to [Establishing a New Logical-Domain Cluster by Deploying the Oracle VM Server for SPARC Template for Oracle](#)

[Solaris Cluster](#). However, you cannot use this template to install and configure control domains as cluster nodes. Instead, follow software installation and configuration procedures as for a physical machine.

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 \(pkg\)](#) or [How to Install Oracle Solaris Cluster Software Using Centralized Installer \(clinstall\)](#).
- 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 \(IPS Repositories\)](#).

See Also

See the [Administering an Oracle Solaris Cluster 4.4 Configuration](#) for procedures to perform dynamic reconfiguration tasks in an Oracle Solaris Cluster configuration.

How to Install `pconsole` Software on an Administrative Console

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

Note:

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

This procedure describes how to install the 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 4.4 software as an 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  ha-cluster-
```

```

repository
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 *Updating Systems and Adding Software in Oracle Solaris 11.4*](#).

 **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 *Updating Systems and Adding Software in Oracle Solaris 11.4*](#).

3. Install the terminal/pconsole package.

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

4. Install the Oracle Solaris Cluster man page packages.

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

Package Name	Description
ha-cluster/system/manual	Oracle Solaris Cluster framework man pages
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 Disaster Recovery Framework 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. `/opt/SUNWcluster/bin/` directory**
 - 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. `startingpconsole` utility**
 - a. `startingpconsole` utility**
 - b. `startingpconsole` 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 Administering an Oracle Solaris Cluster 4.4 Configuration](#) and [How to Connect Securely to Cluster Consoles in Administering an Oracle Solaris Cluster 4.4 Configuration](#) 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](#).

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](#).

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](#).

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 \(pkg\)](#) or [How to Install Oracle Solaris Cluster Software Using Centralized Installer \(clinstall\)](#).
- 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 \(IPS Repositories\)](#).

How to Install and Configure Oracle Solaris Cluster Quorum Server Software

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.

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

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  ha-cluster-
repository
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 *Updating Systems and Adding Software in Oracle Solaris 11.4*](#).

 **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 *Updating Systems and Adding Software in Oracle Solaris 11.4*](#).

3. Install the Quorum Server group package.

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

4. 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. 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. Enable the quorum server SMF service.

```
quorumserver# svcadm enable quorumserver
quorumserver# /usr/cluster/bin/clquorumserver show
```

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-110n` 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](#).

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](#).

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 \(pkg\)](#).
- 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 \(IPS Repositories\)](#).

How to Configure Internal Disk Mirroring

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

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 *Managing Hardware With Oracle Solaris Cluster 4.4*](#).

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\(8\)](#) 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](#).

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 \(pkg\)](#).
- 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 \(IPS Repositories\)](#).

How to Install Oracle VM Server for SPARC Software and Create Domains (SPARC Only)

Perform the following tasks:

- Ensure that the machine is SPARC hypervisor capable.

- Have available [Oracle VM Server for SPARC 3.6 Administration Guide](#) and [Oracle VM Server for SPARC 3.6 Release Notes](#).
- Read the requirements and guidelines in [SPARC: Guidelines for Oracle VM Server for SPARC Logical Domains as Cluster Nodes](#).

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

Alternatively, to use the Oracle VM Server for SPARC template for Oracle Solaris Cluster to install and configure a cluster of guest domains or I/O domains, go to [Establishing a New Logical-Domain Cluster by Deploying the Oracle VM Server for SPARC Template for Oracle Solaris Cluster](#). However, you cannot use this template to install and configure control domains as cluster nodes. Instead, follow software installation and configuration procedures as for a physical machine.

1. **Assume the root role on the machine.**
2. **Install Oracle VM Server for SPARC software and configure domains by following the procedures in the correct version of Oracle VM Server for SPARC 3.6 Administration Guide.**

Observe the following special instructions:

- Adhere to the Oracle Solaris Cluster guidelines for creating guest domains or I/O domains for use as cluster nodes.
- 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](#).

Otherwise, install the Oracle Solaris Cluster software packages. Go to [How to Install Oracle Solaris Cluster Software \(pkg\)](#) or [How to Install Oracle Solaris Cluster Software Using Centralized Installer \(clinstall\)](#).

How to Install Oracle Solaris Cluster Software Using Centralized Installer (clinstall)

Perform the following tasks:

- Verify installation prerequisites.
 - All nodes to install have the same hardware architecture.
 - All nodes are installed with the same version of Oracle Solaris software.
 - 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.

 **Note:**

If you will run the `clinstall` utility in GUI mode, you must install the Oracle Solaris `solaris-desktop` group package on all nodes you want to install.

- Each node name has a valid format and can be resolved by the name server.
- All nodes are online.
- If you install from an IPS repository, the repository is already configured with the Oracle Solaris Cluster packages you want to install.
- If you install from an ISO image, a directory where you will extract the image file is created on a shared file system that is accessible by all nodes that you will install.
- Determine which Oracle Solaris Cluster framework group package or packages you will install. The following matrix lists the framework group packages and which high-level Oracle Solaris Cluster components each package provides.

 **Note:**

You cannot add or remove individual packages that are part of the `ha-cluster-minimal` group package except by complete reinstallation or uninstallation. However, you can add or remove other, optional packages without removing the `ha-cluster-minimal` group package.

Feature	ha-cluster-full	ha-cluster-framework-full	ha-cluster-data-services-full	ha-cluster-geo-full	ha-cluster/system/manager	ha-cluster-framework-minimal	*data-service
Framework	X	X	X	X		X	
Agents	X		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				
Graphical User Interface	X				X		
Disaster Recovery Framework	X			X			

This procedure uses the `clinstall` utility to install Oracle Solaris Cluster framework software in the global zone of one or more potential cluster nodes. You can also use this procedure to install additional software packages to a node that already has Oracle Solaris Cluster framework software installed.

 **Note:**

To instead use the Oracle Solaris `pkg` command to install Oracle Solaris Cluster group packages and individual packages, go to [How to Install Oracle Solaris Cluster Software \(pkg\)](#).

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 -s svc:/network/rpc/bind setprop config/local_only = false
# 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. Determine which cluster node to use to issue installation commands.

This system, called the *control node*, will run the `clinstall` utility and issue the commands to install software on all specified nodes.

 **Tip:**

If you are installing nodes for a new cluster, when you establish the installed nodes as a cluster, issue the `scinstall` command from this same control node. Otherwise, if you choose a different node to issue the `scinstall` command, you would first have to run the `clauth` command again on the other nodes, to authorize the new control node.

4. Verify that the control node can reach each of the other cluster nodes to install.
5. Set the location of the Oracle Solaris Cluster 4.4 package repository on all nodes.

```
# pkg set-publisher -p file:///mnt/repo
```

- a. Before installing the Oracle Solaris Cluster 4.4 software on a host upgraded from Oracle Solaris 11.3 to 11.4, ensure that consolidation/cacao/cacao-incorporation has been upgraded to Oracle Solaris 11.4 version. For example:

```
# pkg list consolidation/cacao/cacao-incorporation
NAME (PUBLISHER) VERSION IFO
consolidation/cacao/cacao-incorporation 11.4-11.4.0.0.1.0 i--
```

- b. If the package version is still from Oracle Solaris 11.3, upgrade the cacao package on all nodes.

```
# pkg update consolidation/cacao/cacao-incorporation
```

- c. Use pkg list to confirm the package was updated.

```
# pkg list consolidation/cacao/cacao-incorporation
```

6. Install ha-cluster/system/preinstall on all nodes, including the control node.

The ha-cluster/system/preinstall package delivers required software for cinstall.

```
# pkg install ha-cluster/system/preinstall
# pkg list ha-cluster/system/preinstall
NAME (PUBLISHER)          VERSION          IFO
ha-cluster/system/preinstall 4.4-0.x.0       i--
```

The ha-cluster/system/preinstall package is included in the Oracle Solaris Cluster 4.4 repository. See [How to Install Oracle Solaris Cluster Software \(pkg\)](#).

7. Authorize acceptance of cluster installation commands by the control node.

Perform this task on all cluster nodes to install.

 **Note:**

The clauth command does not need to be run on the control node, but doing so causes no problems for clauth operations or for cluster installation operations.

```
phys-schost# clauth enable -n control-node
```

 **Note:**

To use the des (Diffie-Hellman) authentication protocol instead of the sys (unix) protocol, include -p des in the command.

```
phys-schost# clauth enable -p des -n control-node
```

For information about setting up DES authentication, see Chapter 5, Authenticating NFS Connections in Oracle Solaris in *Managing Authentication in Oracle Solaris 11.4*.

8. From the control node, launch the cinstall installer.

- To launch the graphical user interface (GUI), use the `-g` option.

```
phys-control# cinstall -g
```

- To launch the interactive text-based utility, use the `-t` option.

```
phys-control# cinstall -t
```

9. Follow the prompts in the `cinstall` installer to provide the following information.

Some GUI or text screens provide a Help option to access additional information about that screen.

Data to Provide	Guidelines	Your Values
Installation mode	<p>Select Initial Installation if cluster framework software is not yet installed on the nodes.</p> <p>Select Additional Cluster Package Installation if cluster framework software is installed and you want to add other component software:</p> <ul style="list-style-type: none"> Data services, data service man pages Generic Data Service Oracle Solaris Cluster Manager browser interface Disaster Recovery Framework (formerly called Geographic Edition) 	
Package source type	Select whether to use an IPS repository or an ISO image file.	
Cluster nodes to install	<p>For GUI mode, type a comma-separated list of the names of all nodes to install. Include the name of the control node, if you want to install the software on it. For example:</p> <pre>pnode1, pnode2, pcontrolnode</pre> <p>For text mode, specify one node name at a time.</p>	
IPS repository (if used)	Provide the fully-qualified path or URI to the IPS repository.	
ISO image (if used)	<p>Provide the fully-qualified path and name of the ISO image file.</p> <p>Provide the fully-qualified path to the directory where you want to extract the ISO image file. The directory must reside on a shared file system that is accessible to all nodes to be installed.</p>	
Package selection	For initial installation, see the table in How to Install Oracle Solaris Cluster Software (pkg) for the high-level contents of each cluster framework package group.	

Data to Provide	Guidelines	Your Values
Data services (if adding after framework installation)	Select one data service package at a time from the displayed pick list. Repeat for each additional package you want to install.	

10. In the Operation Summary panel, review your selection.

Click Back if needed to change any choices, then click Install to begin installation. The utility displays progress messages during installation processing.

11. In the Completion Summary panel, review installation status and perform final tasks.

- If using text mode, press any key to return to the Main Menu, then perform one of the following actions:
 - To install additional cluster components, select Additional cluster package installation.
Follow the prompts to install additional cluster packages.
 - To exit the installer, select Quit.
In text mode, you also have the option to set the source of the Oracle Solaris Cluster software as the `ha-cluster` publisher on all nodes, where it is not already set.
- If using GUI mode, perform the following actions:
 - a. Choose whether to launch the Oracle Solaris Cluster Manager browser interface.
Deselect the checkbox at the bottom of the panel if you do *not* want the browser interface started when you exit the installer.
 - b. Click Finish to exit the installer.
 - c. If you want to install additional components, restart the installer, and select Additional Cluster Package Installation.

Next Steps

If you want to install the Oracle HSM file system software, follow the procedures for initial installation. See [How to Install Oracle HSM Software](#).

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

How to Install Oracle Solaris Cluster Software (pkg)

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](#) for more information about installing the Oracle Solaris software to meet Oracle Solaris Cluster software requirements.

 **Note:**

If you want to use the Oracle Solaris Cluster Manager browser interface to administer Oracle Solaris Cluster components, ensure that all cluster nodes have the same root password.

- Choose which Oracle Solaris Cluster software packages to install.

The following table lists the primary group packages for the Oracle Solaris Cluster 4.4 software and the principal features that each group package contains. You must install at least the `ha-cluster-framework-minimal` group package.

 **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](#) for a procedure 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.

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
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			
Graphical User Interface	X				X	
Disaster Recovery Framework	X			X		

This procedure uses the Oracle Solaris `pkg` command to install Oracle Solaris Cluster framework package groups and individual packages.

 **Note:**

To use the interactive `clinstall` utility instead, go to [How to Install Oracle Solaris Cluster Software Using Centralized Installer \(clinstall\)](#).

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, control domains, or guest domains.

- To install data services.
- 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. CR 6558275 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 -s svc:/network/rpc/bind setprop config/local_only = false
# 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. Per Venkat, deleted obsolete step for NWAM.**

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. 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 4.4 repository.

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

```
# pkg set-publisher -p \  
-k /var/pkg/ssl/Oracle_Solaris_Cluster_4.4.key.pem \  
-c /var/pkg/ssl/Oracle_Solaris_Cluster_4.4.certificate.pem  
  
-k /var/pkg/ssl/Oracle_Solaris_Cluster_4.4.key.pem  
Specifies the full path to the downloaded SSL key file.  
  
-c /var/pkg/ssl/Oracle_Solaris_Cluster_4.4.certificate.pem  
Specifies the full path to the downloaded certificate file.
```

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 4.4 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 4.4 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 4.4 package repository.

```
# pkg set-publisher -p file:///mnt/repo
```

Before installing the Oracle Solaris Cluster 4.4 software on a host upgraded from Oracle Solaris 11.3 to 11.4, ensure that `consolidation/cacao/cacao-incorporation` has been upgraded to Oracle Solaris 11.4 version. For example:

```
# pkg list consolidation/cacao/cacao-incorporation
NAME (PUBLISHER) VERSION          IFO
consolidation/cacao/cacao-incorporation 11.4-11.4.0.0.0.1.0      i--
```

If the package version is still from Oracle Solaris 11.3, upgrade the `cacao` package on all nodes.

```
# pkg update consolidation/cacao/cacao-incorporation
```

Use `pkg list` to confirm the package was updated.

```
# pkg list consolidation/cacao/cacao-incorporation
```

5. 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  ha-cluster-repository
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 Updating Systems and Adding Software in Oracle Solaris 11.4](#).

 **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 Updating Systems and Adding Software in Oracle Solaris 11.4](#).

6. Install the Oracle Solaris Cluster 4.4 software.

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

7. Verify that the package installed successfully.

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

Package installation succeeded if the state is `Installed`.

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

For instructions on updating your software, see [Updating Your Oracle Solaris Cluster 4.4 Environment](#).

Next Steps

If you want to install the Oracle HSM file system software, follow the procedures for initial installation. See [How to Install Oracle HSM Software](#).

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

How to Install Oracle HSM 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 \(pkg\)](#).
2. **Assume the `root` role on a cluster node.**
3. **Install Oracle HSM file system software.**

Follow procedures for initial installation in your Oracle HSM documentation.

Next Steps

Set up the `root` user environment. Go to [How to Set Up the Root Environment](#).

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 Work Environment in Managing User Accounts and User Environments in Oracle Solaris 11.4](#) 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. **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 PF feature of Oracle Solaris, go to [How to Configure Packet Filter](#).

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

How to Configure Packet Filter

Read the guidelines and restrictions to follow when you configure the Packet Filter (PF) feature in a cluster. See the "Packet Filter (PF) Feature" bullet item in [Oracle Solaris OS Feature Requirements and Restrictions](#).

Perform this procedure to configure the Packet Filter (PF) feature of Oracle Solaris software on the global cluster.



Note:

Only use PF with failover data services. The use of PF with scalable data services is not supported.

For more information about the PF feature, see [Oracle Solaris Firewall in *Securing the Network in Oracle Solaris 11.4*](#).

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

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

- In the `pf.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 PF configuration will interfere with cluster membership and infrastructure operations.

For example, suppose the following rules are currently used:

```
# Default block { tcp, udp } unless some later rule overrides
block return in proto { tcp, udp } from any to any

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

# Allow traffic on localhost
pass in quick to localhost
pass out quick from localhost
```

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

```
# clintr status
=== Cluster Transport Paths ===

Endpoint1      Endpoint2      Status
-----
node1:net1     node2:net1     Path online
node1:net2     node2:net2     Path online
# ipadm show-addr | egrep "net1|net2|clprivnet"
net1/?         static  ok         172.16.0.65/26
net2/?         static  ok         172.16.0.129/26
clprivnet0/?   static  ok         172.16.2.1/24
```

The first interconnect `net1` adapter is on subnet `172.16.0.64/26`

The second interconnect `net2` adapter is on subnet `172.16.0.128/26`.

The private network interface `clprivnet0` is on subnet `172.16.2.0/24`.

The PF rules corresponding to networks derived are:

```
# Unblock cluster traffic on 172.16.0.64/26 subnet (physical
interconnect)
pass in quick proto { tcp, udp } from 172.16.0.64/26 to any flags any
pass out quick proto { tcp, udp } from 172.16.0.64/26 to any flags any

# Unblock cluster traffic on 172.16.0.128/26 subnet (physical
interconnect)
pass in quick proto { tcp, udp } from 172.16.0.128/26 to any flags any
pass out quick proto { tcp, udp } from 172.16.0.128/26 to any flags any

# Unblock cluster traffic on 172.16.2.0/24 (clprivnet0 subnet)
pass in quick proto { tcp, udp } from 172.16.2.0/24 to any flags any
pass out quick proto { tcp, udp } from 172.16.2.0/24 to any flags 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 flags any
pass in quick on net2 all flags any
pass in quick on clprivnet0 all flags any
```

- 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 PF'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.
- All cluster specific filtering rules must be added prior to configuring Oracle Solaris Cluster if PF has already been enabled on the cluster nodes.
- Open all ports that are used by your Oracle Solaris Cluster configuration. For instance, if you run the Oracle Solaris Cluster disaster recovery framework, ensure that you open all ports that are used by the framework. See the ports

described in [Configuring Firewalls in *Installing and Configuring the Disaster Recovery Framework for Oracle Solaris Cluster 4.4*](#).

- When you enable a firewall inside of an exclusive IP zone cluster, perform this procedure inside of each zone cluster node.
- A shared IP zone cluster shares the private interconnect with the global zone. Add the `clprivnet` private network interface that is used in the shared IP zone cluster node to the packet filter rules on the global zone.

For more information about PF rules, see the [`pf.conf\(7\)`](#) man page.

Example 2-1 Using an Exclusive IP Zone Cluster

This example shows how to configure an exclusive IP zone cluster.

```
# clnode status -m
--- Node Public Network Status ---

Node Name          PNM Object Name   Status   Adapter          Status
-----
node1              sc_ipmp0          Online   scld02zc2pub1   Online
node2              sc_ipmp0          Online   scld02zc2pub1   Online

# clintr status
=== Cluster Transport Paths ===

Endpoint1          Endpoint2          Status
-----
node2:scl02zc2priv2  node1:scl02zc2priv2  Path online
node2:scl02zc2priv1  node1:scl02zc2priv1  Path online

# ipadm show-addr | egrep "scl02zc2priv1|scl02zc2priv2|clprivnet2"
scl02zc2priv1/?    static   ok           172.18.4.66/26
scl02zc2priv2/?    static   ok           172.18.4.130/26
clprivnet2/?       static   ok           172.18.4.2/26

# grep -v ^# /etc/firewall/pf.conf | grep -v ^$
set reassemble yes no-df
set skip on lo0
ext_if="sc_ipmp0"
client_out="{22, 2084, 5201, 111, 8059, 8060, 8061, 8062, 6499, 11161, 11162, 11163,
11164, 11165}"
pass in quick on scl02zc2priv1 all flags any
pass in quick on scl02zc2priv2 all flags any
pass in quick on clprivnet2 all flags any
block in log quick on egress proto tcp to port { 22 }
block return log all
pass in log proto tcp from any to any port 22 <> 23
pass out log proto tcp from any to any
pass in log proto udp from any to any
pass out inet proto icmp all icmp-type echoreq keep state
pass in log proto icmp from any to any
pass out on $ext_if proto udp all
pass out
#
```

Next Steps

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

3

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](#). 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](#)
- [Establishing a New Global Cluster or New Global-Cluster Node](#)

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.

- [Task Map: Establish a New Global Cluster](#)
- [Task Map: Add a Node to an Existing Global Cluster](#)

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

Method	Instructions
Use one of the following methods to establish a new global cluster:	
<ul style="list-style-type: none">• Use the <code>scinstall</code> utility to establish the cluster.	Configuring Oracle Solaris Cluster Software on All Nodes (scinstall)
<ul style="list-style-type: none">• Use an XML configuration file to establish the cluster.	How to Configure Oracle Solaris Cluster Software on All Nodes (XML)
<ul style="list-style-type: none">• Set up an Automated Installer (AI) install server. Then use the <code>scinstall AI</code> option to install the software on each node and establish the cluster.	Establishing a New Cluster With the Automated Installer
Assign quorum votes and remove the cluster from installation mode, if this operation was not already performed.	How to Configure Quorum Devices
Validate the quorum configuration.	How to Verify the Quorum Configuration and Installation Mode
<i>(Optional)</i> Change a node's private hostname.	How to Change Private Hostnames
Create or modify the NTP configuration file, if not already configured.	Configuring Network Time Protocol (NTP)

Table 3-1 (Cont.) Task Map: Establish a New Global Cluster

Method	Instructions
If using a volume manager, install the volume management software.	Configuring Solaris Volume Manager Software
Create cluster file systems or highly available local file systems as needed.	Creating a Cluster File System or Enabling Highly Available Local File Systems in <i>Planning and Administering Data Services for Oracle Solaris Cluster 4.4</i>
Install third-party applications, register resource types, set up resource groups, and configure data services.	Planning and Administering Data Services for Oracle Solaris Cluster 4.4 Documentation that is supplied with the application software
(Optional) Configure IP Security Architecture (IPsec) on the Cluster Private Interconnect.	Securing the Interconnect for Oracle Solaris Cluster With IPsec
Validate the cluster.	How to Validate the Cluster
Take a baseline recording of the finished cluster configuration.	How to Record Diagnostic Data of the Cluster Configuration

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
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
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 (scinstall)
<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)
Update the quorum configuration information.	How to Update Quorum Devices After Adding a Node to a Global Cluster
Validate the quorum configuration.	How to Verify the Quorum Configuration and Installation Mode
(Optional) Change a node's private hostname.	How to Change Private Hostnames
Modify the NTP configuration.	Configuring Network Time Protocol (NTP)
If using a volume manager, install the volume management software.	Configuring Solaris Volume Manager Software
Create cluster file systems or highly available local file systems as needed.	Creating a Cluster File System or Enabling Highly Available Local File Systems in <i>Planning and Administering Data Services for Oracle Solaris Cluster 4.4</i>

Table 3-2 (Cont.) Task Map: Add a Node to an Existing Global Cluster

Method	Instructions
Install third-party applications, register resource types, set up resource groups, and configure data services.	Planning and Administering Data Services for Oracle Solaris Cluster 4.4 Documentation that is supplied with the application software
(Optional) Configure IP Security Architecture (IPsec) on the Cluster Private Interconnect.	Securing the Interconnect for Oracle Solaris Cluster With IPsec
Validate the cluster.	How to Validate the Cluster
Take a baseline recording of the finished cluster configuration.	How to Record Diagnostic Data of the Cluster Configuration

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, or control 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](#).

This section contains the following information and procedures:

- [Configuring Oracle Solaris Cluster Software on All Nodes \(scinstall\)](#)
- [How to Configure Oracle Solaris Cluster Software on All Nodes \(XML\)](#)
- [Establishing a New Cluster With the Automated Installer](#)
- [Establishing a New Logical-Domain Cluster by Deploying the Oracle VM Server for SPARC Template for Oracle Solaris Cluster](#)
- [How to Prepare the Cluster for Additional Global-Cluster Nodes](#)
- [How to Change the Private Network Configuration When Adding Nodes or Private Networks](#)
- [Configuring Oracle Solaris Cluster Software on Additional Global-Cluster Nodes \(scinstall\)](#)
- [How to Configure Oracle Solaris Cluster Software on Additional Global-Cluster Nodes \(XML File\)](#)
- [How to Update Quorum Devices After Adding a Node to a Global Cluster](#)
- [How to Configure Quorum Devices](#)
- [How to Verify the Quorum Configuration and Installation Mode](#)
- [How to Change Private Hostnames](#)
- [Configuring Network Time Protocol \(NTP\)](#)
- [Configuring Precision Time Protocol \(PTP\) Instead of NTP](#)
- [Securing the Interconnect for Oracle Solaris Cluster With IPsec](#)
- [How to Validate the Cluster](#)

- [How to Record Diagnostic Data of the Cluster Configuration](#)

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?	<i>First:</i> <i>Second:</i>
Quorum Configuration <i>(two-node cluster only)</i>	Do you want to disable automatic quorum device selection? <i>(Answer Yes if any shared storage is not qualified to be a quorum device or if you want to configure a quorum server as a quorum device.)</i>	Yes No
Check	Do you want to interrupt cluster creation for <code>cluster check</code> errors?	Yes No

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

 **Note:**

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

Component	Description/Example	Answer
Cluster Name	What is the name of the cluster that you want to establish?	
Cluster Nodes	List the names of the other cluster nodes planned for the initial cluster configuration. <i>(For a single-node cluster, press Control-D alone.)</i>	
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: <code>switch1</code> and <code>switch2</code>	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: <code>switch1</code> and <code>switch2</code>	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:	

Component	Description/Example	Answer
	Transport adapter name:	First: Second:
	Will this be a dedicated cluster transport adapter? (Answer No if using tagged VLAN adapters.)	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 (a switch or another adapter)? – 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 (multiple-node cluster only)	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? (Choose from the values calculated by scinstall or supply your own.)	____.____.____.____
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?	____ exclusive-IP zone clusters ____ shared-IP zone clusters
Global Fencing	Do you want to disable global fencing? (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.)	First: Yes No Second: Yes No
Quorum Configuration (two-node cluster only)	Do you want to disable automatic quorum device selection? (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.)	First: Yes No Second: Yes No
Check (multiple-node cluster only)	Do you want to interrupt cluster creation for cluster check errors?	Yes No
(single-node cluster only)	Do you want to run the cluster check utility to validate the cluster?	Yes No
Automatic Reboot (single-node cluster only)	Do you want scinstall to automatically reboot the node after installation?	Yes No

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

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](#) for more information about installing the Oracle Solaris software to meet Oracle Solaris Cluster software requirements.

- 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](#).
- Ensure that Oracle Solaris Cluster software packages and updates are installed on each node. See [How to Install Oracle Solaris Cluster Software \(pkg\)](#).
- 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 \(scinstall\)](#).

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(8)` man page.

Note:

Alternatively, you can use the Oracle Solaris Cluster Manager browser interface to perform this task. The cluster configuration wizard can be launched only on a cluster that has not been configured already. You cannot launch it on a node that is already a part of a cluster. The `clauth` command controls the network access policies for machines that are to be configured as nodes of a new-cluster. For more information about the `clauth` command, see the `clauth(8CL)` man page. To use the Oracle Solaris Cluster Manager to configure the cluster, run steps 1 through 5 in this section then from a host that runs the Oracle Solaris Cluster Manager log in to the control node and configure the cluster. For Oracle Solaris Cluster Manager log in instructions, see [How to Access Oracle Solaris Cluster Manager in Administering an Oracle Solaris Cluster 4.4 Configuration](#). After logging in, if the node is not configured as part of the cluster, the wizard will display a screen with the Configure button. Click Configure to launch the cluster configuration wizard.

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

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

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

1. If you 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 restart rpc/bind
```

4. Prepare public-network interfaces.

a. Ensure that static IP addresses for each public-network interface exist.

```
# ipadm
```

If static IP addresses for each public-network interface do not exist, create them.

```
# 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 Managing Network Components in Oracle Solaris 11.4](#).

b. 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 11.4](#).

c. Create trunk and DLMP link aggregations and VNICs that are directly backed by link aggregations for public-network interfaces.

For more information, see [Chapter 2, Configuring High Availability by Using Link Aggregations in Managing Network Datalinks in Oracle Solaris 11.4](#).

5. Authorize acceptance of cluster configuration commands by the control node.

a. Determine which system to use to issue the cluster creation command.

This system is the control node.

b. On all systems that you will configure in the cluster, other than the control node, authorize acceptance of commands from the control node.

```
phys-schost# clauth enable -n control-node
```

If you want to use the des (Diffie-Hellman) authentication protocol instead of the `sys (unix)` protocol, include `-p des` option in the command.

```
phys-schost# clauth enable -p des -n control-node
```

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

```
phys-schost# scinstall
```

7. 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.

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

The Typical or Custom Mode menu is displayed.

**9. 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.

**10. 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.

**11. 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
```

**12. 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\(8CL\)](#) man page.

**13. Take the cluster out of `installmode`.**

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

**14. 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
...
```

**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. 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 for Autofs Maps in Managing Network File Systems in Oracle Solaris 11.4](#) 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 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](#) 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](#).
- If you installed a multiple-node cluster and declined automatic quorum configuration, perform post installation setup. Go to [How to Configure Quorum Devices](#).

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

Otherwise, go to [How to Verify the Quorum Configuration and Installation Mode](#).

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

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](#) for more information about installing the Oracle Solaris software to meet Oracle Solaris Cluster software requirements.

- 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](#).

- 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 4.4 software and updates are installed on each node that you will configure. See [How to Install Oracle Solaris Cluster Software \(pkg\)](#).

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 the Oracle Solaris Cluster software.

This procedure configures the following cluster components:

- Cluster name
  - Cluster node membership
  - Cluster interconnect
- 1. Ensure that the Oracle Solaris Cluster 4.4 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 4.4 software, instead perform upgrade procedures in [Updating Your Oracle Solaris Cluster 4.4 Environment](#).

- 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 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 4.4 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\(8CL\)](#) 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\(7CL\)](#) 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\(7CL\)](#) 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.

### Tip:

If you used the `clinstall` utility to install cluster software on the cluster nodes, you can skip the next step if you issue the `scinstall` command from the same control node. The `clauth` authorizations you made before running the `clinstall` utility stay in force until you reboot the nodes into the cluster at the end of this procedure.

## 8. Authorize acceptance of cluster configuration commands by the control node.

### a. Determine which system to use to issue the cluster creation command.

This system is the control node.

### b. On all systems that you will configure in the cluster, other than the control node, authorize acceptance of commands from the control node.

```
phys-schost# clauth enable -n control-node
```

If you want to use the `des` (Diffie-Hellman) authentication protocol instead of the `sys` (unix) protocol, include `-p des` option in the command.

```
phys-schost# clauth enable -p des -n control-node
```

## 9. 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.

## 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(8CL)` man page.

**12. Perform any necessary updates to the Oracle Solaris Cluster software.**

For instructions on updating your software, see [Updating Your Oracle Solaris Cluster 4.4 Environment](#).

**13. 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.****14. 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 for Autofs Maps in Managing Network File Systems in Oracle Solaris 11.4](#) for more information about modifying the automounter map.

**15. 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](#).

**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](#).

- 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 *Managing Network-Attached Storage Devices in an Oracle Solaris Cluster 4.4 Environment*.
  - 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.

- 16. Remove the cluster from installation mode.

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

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

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

- 18. 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 Oracle Solaris 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\(8CL\)](#)

For Solaris Volume Manager, first create the disk sets that you specify in the cluster configuration XML file.

- Resource Group Manager components

- Resources: `clresource(8CL)`
- Shared address resources: `clressharedaddress(8CL)`
- Logical hostname resources: `clreslogicalhostname(8CL)`
- Resource types: `clresourcetype(8CL)`
- Resource groups: `clresourcegroup(8CL)`

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(8CL)`

You must first set up the NAS device as described in the device's documentation.

- SNMP hosts: `clsnmphost(8CL)`

The `clsnmphost create -i` command requires that you specify a user password file with the `-f` option.

- SNMP users: `clsnmpuser(8CL)`

- Thresholds for monitoring system resources on cluster objects:  
`cltelemetryattribute(8CL)`

### 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 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](#).

## Establishing a New 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

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 Unified 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 \(Unified Archives\)](#), [How to Replicate a Cluster from the Unified Archives](#), [How to Restore a Node from the Unified Archive in \*Administering an Oracle Solaris Cluster 4.4 Configuration\*](#), and [How to Install a Zone Cluster from the Unified Archive in \*Administering an Oracle Solaris Cluster 4.4 Configuration\*](#).

These nodes can be physical machines or supported (SPARC only) Oracle VM Server for SPARC logical 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>What is the full path name of the Automated Installer boot image ISO file? |        |

| Component                                                                    | Description/Example                                                                                                                                                                                                                                                                                                                              | Answer            |
|------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------|
|                                                                              | <p>If you plan to use a repository to get the AI boot image, you will need the following information:</p> <p>What is the publisher for the boot image <code>install-image/solaris-auto-install</code> package?</p> <p>What is the repository of the publisher?</p> <p>What is the architecture of the cluster nodes?</p>                         |                   |
|                                                                              | <p>For repositories that use HTTPS:</p> <p>What is the full path of the certificate file for the repository?</p> <p>What is the full path of the private key file for the repository?</p> <p>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.</p> |                   |
| Unified Archives                                                             | <p>If you plan to use the Unified Archives to install, you will need the following information:</p> <p>What is the location of the unified archive?</p>                                                                                                                                                                                          |                   |
| Custom Automated Installer User root Password                                | What is the password for the root account of the cluster nodes?                                                                                                                                                                                                                                                                                  |                   |
| Custom Automated Installer Repositories (when not using the Unified Archive) | What is the repository of publisher <code>solaris</code> ?                                                                                                                                                                                                                                                                                       |                   |
|                                                                              | What is the repository of publisher <code>ha-cluster</code> ?                                                                                                                                                                                                                                                                                    |                   |
|                                                                              | <p>For repositories that use HTTPS:</p> <p>What is the full path of the certificate file for the repository?</p> <p>What is the full path of the private key file for the repository?</p> <p>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.</p> |                   |
|                                                                              | 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 Cables (VLAN cannot be used with cluster AI.) | First node name:                                                                                                                                                                                                                                                                                                                                 |                   |
|                                                                              | Transport adapter names:                                                                                                                                                                                                                                                                                                                         | First:<br>Second: |
| <i>Specify for each additional node</i>                                      | Node name:                                                                                                                                                                                                                                                                                                                                       |                   |
|                                                                              | Transport adapter names:                                                                                                                                                                                                                                                                                                                         | First:<br>Second: |

| Component                                       | Description/Example                                                                                                                                                                                    | Answer                              |
|-------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------------------------|
| Quorum Configuration<br>(two-node cluster only) | Do you want to disable automatic quorum device selection?<br>(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.) | 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 Boot Image ISO source                             | If you plan to use a downloaded AI ISO image file, you will need the following information:<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>What is the publisher for the boot image <code>install-image/solaris-auto-install</code> package?<br>What is the repository of the publisher?<br>What is the architecture of the cluster nodes?                         |        |
|                                                                              | For repositories that use HTTPS:<br>What is the full path of the certificate file for the repository?<br>What is the full path of the private key file for the repository?<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. |        |
| Unified Archives                                                             | If you plan to use the Unified Archives to install, you will need the following information:<br>What is the location of the unified archive?                                                                                                                                                                                  |        |
| Custom Automated Installer User root Password                                | What is the password for the root account of the cluster nodes?                                                                                                                                                                                                                                                               |        |
| Custom Automated Installer Repositories (when not using 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>What is the full path of the certificate file for the repository?<br>What is the full path of the private key file for the repository?<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. |        |

| Component                                                                                                              | Description/Example                                                                                                                                    | Answer                                                                               |
|------------------------------------------------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------|--------------------------------------------------------------------------------------|
|                                                                                                                        | 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.                                                                                 |                                                                                      |
| 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?                       | ____ 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>(VLAN cannot be used with cluster AI.) | First 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                                                  |

| Component                                                        | Description/Example                                                                                                                                                                                        | Answer                              |
|------------------------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------------------------|
|                                                                  | – If no, what is the name of the port that you want to use?                                                                                                                                                | First:<br>Second:                   |
| Specify for each additional node<br>(multiple-node cluster only) | Node name:                                                                                                                                                                                                 |                                     |
|                                                                  | Transport adapter name:                                                                                                                                                                                    | First:<br>Second:                   |
|                                                                  | Where does each transport adapter connect to (a switch or another adapter)?<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>(two-node cluster only)                  | Do you want to disable automatic quorum device selection? <i>(Answer Yes if any shared storage is not qualified to be a quorum device or if you want to configure a quorum server 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, control 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 [Managing Hardware With Oracle Solaris Cluster 4.4](#) 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 [Configuring an AI Server in Automatically Installing Oracle Solaris 11.4 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](#) 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](#).
- 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 4.4 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       |
| Disaster Recovery Framework |                           |                               |                              | X                   |         |

- Have available your completed Typical Mode or Custom Mode installation worksheet. See [Establishing a New Cluster With the Automated Installer](#).

## 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(8)` 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 4, Setting Up the AI Server in Automatically Installing Oracle Solaris 11.4 Systems](#) and [Working With DHCP in Oracle Solaris 11.4](#).

### 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
```

#### Tip:

If you used the `clinstall` utility to install cluster software on the cluster nodes, you can skip the next step if you issue the `scinstall` command from the same control node. The `clauth` authorizations you made before running the `clinstall` utility stay in force until you reboot the nodes into the cluster at the end of this procedure.

### 4. Authorize acceptance of cluster configuration commands by the control node.

#### a. Determine which system to use to issue the cluster creation command.

This system is the control node.

- b. On all systems that you will configure in the cluster, other than the control node, authorize acceptance of commands from the control node.**

```
phys-schost# clauth enable -n control-node
```

If you want to use the des (Diffie-Hellman) authentication protocol instead of the sys (unix) protocol, include `-p des` option in the command.

```
phys-schost# clauth enable -p des -n control-node
```

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

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

The `scinstall` Main Menu is displayed.

- 6. 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:
```

- 7. Follow the menu prompts to supply your answers from the configuration planning worksheet.**
- 8. 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.

- 9. (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/autoscinstall.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.example.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 [Configuring an AI Server in Automatically Installing Oracle Solaris 11.4 Systems](#) and the `ai_manifest(5)` man page.

**c. 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.

**10. 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.

**11. 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](#), then install and configure the cluster by following instructions in [How to Install Oracle Solaris Cluster Software \(pkg\)](#).

- **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.

- 12. 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 for Autofs Maps in Managing Network File Systems in Oracle Solaris 11.4](#) for more information about modifying the automounter map.

- 13. (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 \(kmdb\) Enabled in Booting and Shutting Down Oracle Solaris 11.4 Systems](#).

- 14. 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](#).

**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 [About Run Level Booting in Booting and Shutting Down Oracle Solaris 11.4 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.

**15. 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(8CL)` man page.

**16. 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.

- 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.

17. 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
...
```

18. 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](#)
- [How to Install Oracle VM Server for SPARC Software and Create Domains](#)
- [How to Install Oracle HSM Software](#)
- [How to Set Up the Root Environment](#)
- [How to Configure Packet Filter](#)

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](#).
- If you installed a multiple-node cluster and declined automatic quorum configuration, perform post installation setup. Go to [How to Configure Quorum Devices](#).

- 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](#).
- 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](#).
- 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](#).
- If you installed a single-node cluster, cluster establishment is complete. Go to [Creating Cluster File Systems](#) 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 9](#) 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\(8\)](#) 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 4, Setting Up the AI Server in Automatically Installing Oracle Solaris 11.4 Systems](#) and [Working With DHCP in Oracle Solaris 11.4](#).

**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/autoscinstall.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 [Configuring an AI Server in Automatically Installing Oracle Solaris 11.4 Systems](#) and the `ai_manifest(5)` 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.

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](#), then install and configure the cluster by following instructions in [How to Install Oracle Solaris Cluster Software \(pkg\)](#).

- 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\(8CL\)](#) 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\(8\)](#) 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 4, Setting Up the AI Server in Automatically Installing Oracle Solaris 11.4 Systems](#) and [Working With DHCP in Oracle Solaris 11.4](#).

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/autoscinstall.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 [Configuring an AI Server in Automatically Installing Oracle Solaris 11.4 Systems](#) and the `ai_manifest(5)` 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.

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](#), then install and configure the cluster by following instructions in [How to Install Oracle Solaris Cluster Software \(pkg\)](#).

- 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(8CL)` 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 Disaster Recovery Framework, follow the procedures in [Chapter 5, Administering Cluster Partnerships in Administering the Disaster Recovery Framework for Oracle Solaris Cluster 4.4](#) 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.

Establishing a New Logical-Domain Cluster by Deploying the Oracle VM Server for SPARC Template for Oracle Solaris Cluster

This section provides procedures to use the Oracle VM Server for SPARC template for Oracle Solaris Cluster, to configure only guest domains or I/O domains as global cluster nodes. The new cluster can have a mix of guest domains and I/O domains.

Note:

This template is not valid to use with control domains. To configure control domains as cluster nodes instead, follow procedures as for physical machines to install software and establish the cluster. See [Finding Oracle Solaris Cluster Installation Tasks](#).

To use the template to add a logical domain to a cluster, go to [How to Add a New Logical-Domain Node to an Existing Cluster by Using the Oracle VM Server for SPARC Template for Oracle Solaris Cluster](#).

How to Deploy the Oracle VM Server for SPARC Template for Oracle Solaris Cluster to Configure a Logical Domain Cluster

- Ensure that the `ovmtutils` package is installed in the control domain. You can use the following command to verify whether the package has been installed.

```
# pkg info ovmtutils
```

- Ensure that the Oracle VM Server for SPARC template file is accessible from the control domains.
- Ensure that the Oracle VM Server for SPARC services have been defined:

- **Virtual disk service** - The `ovmtutils` create and configure various aspects of the Oracle VM Server for SPARC environment during deployment, but will require some services to be present. Additionally, some services are required by subsequent tasks. The following is an example command to create a disk service and is run from the control domain:

```
# /usr/sbin/ldm add-vds primary-vds0 primary
```

- **Virtual console concentrator service** - The following is an example command to create a console concentrator service and is run from the control domain:

```
# /usr/sbin/ldm add-vcc port-range=5000-5100 primary-vcc0 primary
```

- Have available the following information:

- * `com.oracle.hacluster.config.cluster_name` – Specifies the cluster name.
- * `com.oracle.hacluster.config.node_list` – Comma-separated list of hostnames of the logical domains to form the cluster. The first node in the list is the first one to be added to the cluster and serves as the sponsor node for the rest of the nodes. It is required that all domains deployed by using the template to have the exact same list, as the order of the list matters: the first host name is the sponsor node.
- * `com.oracle.hacluster.config.interconnect` – Comma-separated list of the interconnect adapters, or `pkeys` if you are using InfiniBand partitions.
- * `com.oracle.hacluster.config.private_netaddr` – (Optional) Specify a private network address that is compatible with netmask 255.255.240.0. The default address is 172.16.0.0. When using InfiniBand, the default private network address can be used as the `pkeys` are different among the clusters.

You can use the same cluster configuration property file for all the new domains.

Type all commands in this procedure from the control domain. The template file is located at `/export/ovmt/ovm_osc5_sparc.ova`.

Perform this procedure to create a cluster of guest domains or of I/O domains.

Note:

This procedure cannot be used for the following tasks:

- Create a cluster of control domains.
- Add nodes to an existing cluster of logical domains.

Instead, follow procedures as for physical machines to perform these tasks. See [Finding Oracle Solaris Cluster Installation Tasks](#).

1. Remove the target logical domain if it already exists.

```
# ovmtdeploy -U newdomain
```

2. Create the working directory.

If the working directory already exists, make sure that the directory is empty.

```
# mkdir -p /domains/newdomain
# ls -l /domains/newdomain
total 0
```

3. List the contents of the template without deploying the template.

```
# ovmtdeploy -n -l -d newdomain /export/ovmt/ovm_osc5_sparc.ova
```

```
Oracle VM for SPARC Deployment Utility
ovmtdeploy Version 3.4.0.0.11
Copyright (c) 2014, 2015, Oracle and/or its affiliates. All rights reserved.
```

```
STAGE 1 - EXAMINING SYSTEM AND ENVIRONMENT
-----
```

```
Checking user privilege
Performing platform & prerequisite checks
Checking for required services
```

```

Named resourced available
STAGE 2 - ANALYZING ARCHIVE & RESOURCE REQUIREMENTS
-----
Checking .ova format and contents
Validating archive configuration
Listing archive configuration
Assembly
-----
Assembly name: ovm_osc5_sparc.ovf
Global settings:
References: system -> System.img.gz
Disks: system -> system
Networks: primary-vsw0
Virtual machine 1
-----
Name: newdomain
Description: Oracle Solaris Cluster 4.4 with 2 vCPUs, 4G memory, 1 disk
image(s)
vcpu Quantity: 2
Memory Quantity: 4G
Disk image 1: ovf:/disk/system -> system
Network adapter 1: Ethernet_adapter_0 -> primary-vsw0
Oracle Solaris Cluster 4.4
name
Solaris 11 System
computer-name
ifname
time-zone
keyboard
language
Solaris 11 Root Account
root-password
Solaris 11 User Account
name.0
real-name.0
password.0
Solaris 11 Network
ipaddr.0
netmask
gateway.0
dns-servers.0
dns-search-domains.0
name-service
domain-name
nis-servers
ldap-profile
ldap-servers
ldap-search-base
ldap-proxy-bind-distinguished-name
ldap-proxy-bind-password
Oracle Solaris Cluster
cluster_name
node_list
interconnect
private_netaddr

```

4. **Prepare the system configuration property files which are required for configuring each domain.**

Use the template Oracle Solaris system configuration file to compose your own file. The template for the Oracle Solaris system configuration property file is available at `/opt/ovmtutils/share/props/solaris.properties`.

The system configuration property file is different for each node. A name service must be provided in the Oracle Solaris property file, so that the nodes can resolve the remote sponsor node name when they join the cluster.

5. Prepare the cluster configuration property files which is required to add each domain to form the cluster.

The cluster configuration file includes the following Oracle Solaris Cluster properties:

- `com.oracle.hacluster.config.cluster_name` – Specifies the cluster name.
- `com.oracle.hacluster.config.node_list` – Comma-separated list of hostnames of the logical domains to form the cluster. The first node in the list is the first one to be added to the cluster and serves as the sponsor node for the rest of the nodes. It is required that all domains deployed by using the template to have the exact same list, as the order of the list matters: the first host name is the sponsor node.
- `com.oracle.hacluster.config.interconnect` – Comma-separated list of the interconnect adapters, or `pkeys` if you are using InfiniBand partitions.
- `com.oracle.hacluster.config.private_netaddr` – (Optional) Specify a private network address that is compatible with netmask 255.255.240.0. The default address is 172.16.0.0. When using InfiniBand, the default private network address can be used as the `pkeys` are different among the clusters.

You can use the same cluster configuration property file for all the new domains.

6. In the control domains, type the `ovmtdeploy` command to deploy the new domains.

You can use different options in the following scenarios:

- If the control domain is in the `vanilla` state and does not have the switches created yet, use the `-e` option to specify the adapters for creating vswitches or vnets.
- If the switches are already created in the control domain, you can use the order in the template as shown in the output of `ovmtdeploy -n -l`, or use the `-t` option to specify the order for using existing vswitches for each of the vnets.
- Specify the SR-IOV virtual functions by using the `-I` option.

The following example deploys a new domain with the specified switches/adapters and disks. The first disk specified by the `-v` option is the local root disk for the new domain, and the following two disks are shared disks:

```
# /opt/ovmtutils/bin/ovmtdeploy -d newdomain -o /domains/newdomain \
-k -s -c 8 -t primary-vsw0,priv-vsw1,priv-vsw2 -e net0,net1,net2 \
-v /dev/rdisk/c0tNd0s2,/dev/rdisk/c0tNd0s2,/dev/rdisk/c0tNd0s2 \
/export/ovmt/ovm_osc5_sparc.o
```

The following example uses SR-IOV virtual functions for deploying a new domain:

```
# /opt/ovmtutils/bin/ovmtdeploy -d newdomain -o /domains/newdomain -k -s -c 8 \
-I /SYS/PCI-EM0/IOVIB.PF0.VF0,/SYS/PCI-EM4/IOVIB.PF0.VF0 \
-e net0 \
-v /dev/rdisk/c0tNd0s2,/dev/rdisk/c0tNd0s2,/dev/rdisk/c0tNd0s2 \
/export/ovmt/ovm_osc5_sparc.o
```

The `-v` option specifies a comma-separated list of target devices. For cluster, you can specify an Oracle Solaris raw whole disk device, for example, `/dev/rdskc3t3d0s2` or an iSCSI device, such as `/dev/rdskc0t600144F00021283C1D7A53609BE10001d0s2`. A target device on a slice and a target device on a block device are not supported. Specify the `root zpool` disk as the very first one. If you specify multiple disks including local disks and shared devices, specify them in the same order for all the domains. For more information, see the `ovmtdeploy(1M)` man page.

7. Configure the new domain to form the cluster.

In all the control domains, use the `ovmtconfig` command to configure the new domains with the system and Oracle Solaris Cluster configuration property files created in Step 4 and Step 5. The `ovmtconfig` command will also boot the domain to complete the remaining configuration operations done by the software. In this process the domain will be rebooted twice, and the last reboot will bring it into cluster mode.

Use the `-P` option to specify the system and Oracle Solaris Cluster configuration property files, or use the `-p` option to specify an individual property which overrides the same property specified inside the property file. Information about the use of other options can be found in the `ovmtconfig(8)` man page.

```
# ovmtconfig -d newdomain -s -v \  
-P /export/ovmt/properties/system_node10VM.props,/export/ovmt/properties/  
cluster_newdomain.props  
# ldm ls
```

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

8. Identify the console port number of the domain and then connect to the console of that domain.

```
# ldm ls newdomain  
# telnet 0 console-port-number-of-newdomain
```

a. When all the domains join the cluster, log in to the domain and check the cluster configuration and status.

```
# cluster show  
# cluster status
```

b. Use the `pkg info` command to confirm whether the cluster packages are installed.

c. Use the `cluster check` command to verify the cluster configuration.

d. Check whether any SMF services are in maintenance mode.

```
# svcs -xv
```

e. Check the public network configuration.

```
# ipmpstat -g
```

9. If the `svc:/system/cluster/sc-ovm-config:default` SMF service is failed and is in maintenance mode, check the deploy log file at `/var/cluster/logs/install` for the detailed list of errors.

10. Request and download your own key and certificate files.

The `solaris` and `ha-cluster` publishers that are set in the deployed domain do not work until you perform this step.

a. Unset the `solaris` and `ha-cluster` publishers.

```
# pkg unset-publisher solaris
# pkg unset-publisher ha-cluster
```

b. Go to `https://pkg-register.oracle.com`.

c. Choose Oracle Solaris Cluster software.

d. Accept the license.

e. 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.

f. Download the key and certificate files and install them as described in the returned certification page.

g. Configure the `ha-cluster` publisher with the downloaded SSL keys and set the location of the Oracle Solaris Cluster 4.4 repository.

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

```
# pkg set-publisher -p \  
-k /var/pkg/ssl/Oracle_Solaris_Cluster_4.4.key.pem \  
-c /var/pkg/ssl/Oracle_Solaris_Cluster_4.4.certificate.pem \  
-k /var/pkg/ssl/Oracle_Solaris_Cluster_4.4.key.pem
```

Specifies the full path to the downloaded SSL key file.

```
-c /var/pkg/ssl/Oracle_Solaris_Cluster_4.4.certificate.pem
```

Specifies the full path to the downloaded certificate file.

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

How to Prepare the Cluster for Additional Global-Cluster Nodes

Perform the following tasks:

- Ensure that all necessary hardware is installed.
 - Ensure that the host adapter is installed on the new node. See the [Managing Hardware With Oracle Solaris Cluster 4.4](#).
 - Verify that any existing cluster interconnects can support the new node. See the [Managing Hardware With Oracle Solaris Cluster 4.4](#).
 - Ensure that any additional storage is installed.

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

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.
 - i. Choose the **Add a Transport Switch** menu item and configure a minimum of two switches (one per adapter pair).
 - ii. Then choose the **Add a Transport Adapter to a Node** menu item and ensure a minimum of two adapters are configured.
 - iii. Then choose the **Add a Transport Cable** menu item to connect each adapter to each switch (`netX <-> switch 1; netY <-> switch 2`).
- c. When finished, quit the `clsetup` utility.
- d. 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 \(scinstall\)](#).
 - 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](#). 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 \(scinstall\)](#) or [How to Configure Oracle Solaris Cluster Software on Additional Global-Cluster Nodes \(XML File\)](#).

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 `num_zoneclusters`, do not perform this procedure. Instead, you can run the following command from a global-cluster node that is in cluster mode to specify the expected number of zone clusters:

```
phys-schost# cluster set-netprops -p num_zoneclusters=N <cluster-name>
```

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 [About Run Level Booting in *Booting and Shutting Down Oracle Solaris 11.4 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.

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 [About Run Level Booting in *Booting and Shutting Down Oracle Solaris 11.4 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 \(scinstall\)](#)
- [How to Install and Configure Oracle Solaris and Oracle Solaris Cluster Software \(IPS Repositories\)](#)
- [How to Configure Oracle Solaris Cluster Software on Additional Global-Cluster Nodes \(XML File\)](#)

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](#) and [Planning the Oracle Solaris Cluster Environment](#) 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> |

| Component | Description/Example | Answer |
|---------------------------------------|--|---|
| 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 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](#) for more information about installing the Oracle Solaris software to meet Oracle Solaris Cluster software requirements.

- 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](#).
- Ensure that Oracle Solaris Cluster software packages and updates are installed on the node. See [How to Install Oracle Solaris Cluster Software \(pkg\)](#).
- Ensure that the cluster is prepared for the addition of the new node. See [How to Prepare the Cluster for Additional Global-Cluster Nodes](#).
- Have available your completed Typical Mode or Custom Mode installation worksheet. See [Configuring 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 \(IPS Repositories\)](#).

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(8)` 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.

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 restart rpc/bind
```

3. Prepare public-network interfaces.

- a. Ensure that static IP addresses for each public-network interface exist. If static IP addresses for each public-network interface do not exist, create them.

```
# 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 Managing Network Components in Oracle Solaris 11.4*](#).

- b. 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 11.4*](#).

- c. Create trunk and DLMP link aggregations and VNICs that are directly backed by link aggregations for public-network interfaces.

For more information, see [Chapter 2, Configuring High Availability by Using Link Aggregations in *Managing Network Datalinks in Oracle Solaris 11.4*](#).

4. Authorize acceptance of cluster configuration commands by the control node.

- a. Determine which system to use to issue the cluster creation command.

This system is the control node.

- b. On all systems that you will configure in the cluster, other than the control node, authorize acceptance of commands from the control node.

```
phys-schost# clauth enable -n control-node
```

If you want to use the des (Diffie-Hellman) authentication protocol instead of the `sys (unix)` protocol, include `-p des` option in the command.

```
phys-schost# clauth enable -p des -n control-node
```

5. Start the `scinstall` utility.

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

The `scinstall` Main Menu is displayed.

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 Add This Machine as a Node in an Existing Cluster and press the Return key.**
8. **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.

9. **Repeat this procedure on any other node to add to the cluster until all additional nodes are fully configured.**
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 an active cluster member, prevent any other nodes from joining the cluster.**

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

Alternatively, you can use the `clsetup` utility. See [How to Add a Node to an Existing Cluster or Zone Cluster in Administering an Oracle Solaris Cluster 4.4 Configuration](#) for procedures.

12. **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(8CL)` man page.

13. **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.

14. Verify that all necessary software updates are installed.

```
phys-schost# pkg list
```

15. 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.added + at the end of the following command for bug 21848444

```
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
...
```

16. 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 for Autofs Maps in Managing Network File Systems in Oracle Solaris 11.4](#) 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.
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 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](#).

Otherwise, go to [How to Verify the Quorum Configuration and Installation Mode](#).

How to Configure Oracle Solaris Cluster Software on Additional Global-Cluster Nodes (XML File)

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](#) for more information about installing the Oracle Solaris software to meet Oracle Solaris Cluster software requirements.

- 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](#).
- Ensure that Oracle Solaris Cluster software packages and any necessary updates are installed on the node. See [How to Install Oracle Solaris Cluster Software \(pkg\)](#).
- Ensure that the cluster is prepared for the addition of the new node. See [How to Prepare the Cluster for Additional Global-Cluster Nodes](#).

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 4.4 software.

This procedure configures the following cluster components on the new node:

- Cluster node membership
 - Cluster interconnect
 - Global devices
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 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 [About Run Level Booting in *Booting and Shutting Down Oracle Solaris 11.4 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 4.4 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\(8CL\)](#) 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 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\(7CL\)](#) 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\(7CL\)](#) 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. 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 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](#).

Otherwise, go to [How to Verify the Quorum Configuration and Installation Mode](#).

How to Add a New Logical-Domain Node to an Existing Cluster by Using the Oracle VM Server for SPARC Template for Oracle Solaris Cluster

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](#) for more information about installing the Oracle Solaris software to meet Oracle Solaris Cluster software requirements.

- 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](#).
- Ensure that the cluster is prepared for the addition of the new node. See [How to Prepare the Cluster for Additional Global-Cluster Nodes](#).

Perform this procedure to add a logical domain as a node to an existing cluster by using the Oracle VM Server for SPARC template for Oracle Solaris Cluster.

1. **Use the `ovmtdeploy` command to create the new logical domain and add it to the cluster as a new node.**

For more information, see [Establishing a New Logical-Domain Cluster by Deploying the Oracle VM Server for SPARC Template for Oracle Solaris Cluster](#).

2. **Use the `ovmtconfig` command to configure the new domain with the system and cluster properties.**

For more information, see [Establishing a New Logical-Domain Cluster by Deploying the Oracle VM Server for SPARC Template for Oracle Solaris Cluster](#).

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](#).

Otherwise, go to [How to Verify the Quorum Configuration and Installation Mode](#).

How to Update Quorum Devices After Adding a Node to a Global Cluster

Ensure that you have completed installation of the Oracle Solaris Cluster software on the added node.

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.

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. 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. 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/clt1d0 |
| d3 | phys-schost-1:/dev/rdisk/clt1d0 |
| ... | |

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/clt1d0
d3              phys-schost-1:/dev/rdisk/clt1d0
...
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](#).

How to Configure Quorum Devices

- **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](#).
 - 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 [Managing Network-Attached Storage Devices in an Oracle Solaris Cluster 4.4 Environment](#) and your device documentation for requirements and installation procedures for NAS hardware and software.

 **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](#).

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.

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
ipmp0/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/clt1d0 |
| d3 | phys-schost-1:/dev/rdisk/clt1d0 |
| ... | |

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](#) 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/rdsk/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](#).

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

How to Modify a Quorum Device Node List in *Administering an Oracle Solaris Cluster 4.4 Configuration*.

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. **cluster command verifying installation mode `/usr/cluster/bin/cluster` command verifying installation mode verifying installation mode installation mode verifying global-cluster nodes verifying installation mode** 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](#).
- If you want to install or modify the NTP configuration file, go to [Configuring Network Time Protocol \(NTP\)](#).
- If you want to install a volume manager, go to [Configuring Solaris Volume Manager Software](#).
- If you want to create cluster file systems, go to [How to Create UFS Cluster File Systems](#).
- 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 [Planning and Administering Data Services for Oracle Solaris Cluster 4.4](#).
- When your cluster is fully configured, validate the configuration. Go to [How to Validate the Cluster](#).
- 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](#).

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 Administering an Oracle Solaris Cluster 4.4 Configuration](#).

How to Change Private Hostnames

Perform this task if you do not want to use the default private host names, `clusternode $nodeID$ -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 host name.**
Repeat for each private host name to change.
6. **Verify the new private host names.**

```
phys-schost# cluster 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](#).

Configuring Network Time Protocol (NTP)

This section contains the following procedures:

- [How to Use Your Own /etc/inet/ntp.conf File](#)
- [How to Install NTP After Adding a Node to a Single-Node Cluster](#)
- [How to Update NTP After Changing a Private Hostname](#)

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](#).

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 create cluster file systems, go to [How to Create UFS Cluster File Systems](#).
- 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 *Planning and Administering Data Services for Oracle Solaris Cluster 4.4*.
- When your cluster is fully configured, validate the configuration. Go to [How to Validate the Cluster](#).
- 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](#).

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 create cluster file systems, go to [How to Create UFS Cluster File Systems](#).
- 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 *Planning and Administering Data Services for Oracle Solaris Cluster 4.4*.
- When your cluster is fully configured, validate the configuration. Go to [How to Validate the Cluster](#).
- 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](#).

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 host name.**
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 create cluster file systems, go to [How to Create UFS Cluster File Systems](#).
- 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 *Planning and Administering Data Services for Oracle Solaris Cluster 4.4*.
- When your cluster is fully configured, validate the configuration. Go to [How to Validate the Cluster](#).
- 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](#).

Configuring Precision Time Protocol (PTP) Instead of NTP

This section contains information about using PTP instead of NTP for clock synchronization.

- [How to Configure PTP \(Oracle Solaris PTP\)](#)

- [How to Configure PTP \(Third Party PTP\)](#)

How to Configure PTP (Oracle Solaris PTP)

Use the following procedure to configure PTP, if Oracle Solaris PTP software is being used.

Note:

If PTP Hardware assistance is enabled, the following restrictions will apply to the configuration:

- Vnics can not be configured over the data links assigned to the PTP service.
- Data link assigned to PTP service cannot be used by the Oracle Solaris Cluster public network to configure Logical Hostname resource as it requires IPMP group configured over the data link.
- Network adapter used by Solaris PTP can not be used by Solaris Cluster for hosting Logical Hostname or Shared Address resources.

1. **Assume the `root` role on a cluster node.**
2. **On each node of the cluster, disable the NTP SMF service.**

```
phys-schost# svcadm disable ntp
```

3. **On each node of the cluster, configure and enable the PTP SMF service.**

For configuration directions, see [Oracle Solaris Cluster 4.3.9.2.0 Support Repository Updates \(SRU 9.2\) Readme \(Doc ID 2335972.1\)](#) and [Managing the Precision Time Protocol in *Managing Clock Synchronization in Oracle Solaris 11.4*](#).

How to Configure PTP (Third Party PTP)

1. **Assume the `root` role on a cluster node.**
2. **On each node of the cluster, disable the NTP SMF service.**

```
phys-schost# svcadm disable ntp
```

3. **On each node of the cluster, ensure that the cluster software will not re-enable NTP service upon reboot of the node, which can interfere with PTP.**

```
phys-schost# /usr/bin/touch /etc/cluster/notntp
```

Securing the Interconnect for Oracle Solaris Cluster With IPsec

You can optionally configure IPsec to protect traffic on the interconnect, which carries traffic both for internal cluster operations as well as user data for global filesystems.

The following requirements have to be met while configuring IPsec and IKE for Oracle Solaris Cluster interconnect:

- Configure IKEv2 to automatically manage keys for IPsec on the interconnect.

 **Note:**

You can also configure IKE for traffic on the public network.

- Use the cluster private network address and prefix the same in IKE and IPsec configuration files. The `cluster` command provides this information. For more information, see the `cluster(8CL)` man page.

The `network/prefix` identifier allows the same IKE rule and IPsec policy to be added identically to all cluster nodes irrespective of the local address assignment on the interconnect.

- IKEv2 rules pertaining to the interconnect must be tagged with "cluster_interconnect" attribute and with value "yes". This ensures proper handling of cluster interconnect traffic during system shutdown.
- IKEv2 and IPsec configuration files on all cluster nodes must contain the same IKEv2 rule and IPsec policy that pertain to the interconnect. When new nodes are added to the cluster, you must ensure that it contains the same IKEv2 rules and IPsec policies as the other nodes in the existing cluster.

 **Note:**

IKEv2 service should be enabled on the private network either in non-cluster mode or simultaneously on all the nodes of the cluster when in cluster mode. Enabling IPsec on a subset of nodes can lead to disconnects in the cluster private network paths and nodes could go down due to loss of quorum.

How to Protect Oracle Solaris Cluster Interconnect with IPsec

1. Determine the cluster private network and its' prefix.

```
# /usr/cluster/bin/cluster show-netprops

=== Private Network ===

private_netaddr:          172.16.0.0
private_netmask:         255.255.240.0
```

The netmask of `255.255.240.0` indicates a prefix length of 20 bits. Thus the private network prefix to be used for the following IKE and IPsec configuration files is `172.16.0.0/20`.

2. Add IKEv2 rules for the interconnect.

For more information, see `ikev2.config(5)` man page.

```
# pfedit /etc/inet/ike/ikev2.config

ikesa_xform { dh_group 21 auth_alg sha512 encr_alg aes }
## Rules for cluster interconnect
{
    label "cluster-foobar"
    auth_method preshared
    local_addr 172.16.0.0/20
    remote_addr 172.16.0.0/20
```

```
cluster_interconnect yes
}
```

 **Note:**

You can configure IKEv2 to authenticate by using methods other than preshared keys. For more information, Consult <reference XXX> for more details.

3. Configure IKEv2 with preshared keys.

For more information, see the `ikev2.preshared(5)` man page.

```
# pfedit /etc/inet/ike/ikev2.preshared

## label must match the rule that uses this key
{
  label "cluster-foobar"
  key 0001020304050607
}
```

 **Note:**

The above sample key is an example. You must create an unique key for each cluster.

4. Add IPsec policies for interconnect.

For more information, see the `ipsecconf(8)` man page.

```
# pfedit /etc/inet/ipsecinit.conf

{ laddr 172.16.0.0/20 raddr 172.16.0.0/20 } ipsec
{ encr_algs aes encr_auth_algs sha256 }
```

How to Validate the Cluster

Ensure that you have completed the installation and configuration of all hardware and software components in the cluster, including firmware and software updates.

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\(8CL\)](#) man page.

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 Administering an Oracle Solaris Cluster 4.4 Configuration](#).

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](#).

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.**

```
# pkg install support/explorer
```

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(8)` man page in the `/opt/SUNWexplo/man/man8` 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.

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](#). See your Solaris Volume Manager documentation for additional details.

This chapter contains the following sections:

- [Configuring Solaris Volume Manager Software](#)
- [Creating Disk Sets in a Cluster](#)
- [Configuring Dual-String Mediators](#)

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 |
| Install the Solaris Volume Manager software. | How to Install Solaris Volume Manager |
| Create state database replicas on the local disks. | How to Create State Database Replicas |

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 [Creating Package Repositories in Oracle Solaris 11.4](#).

3. **Install the Solaris Volume Manager package.**

```
# pkg install storage/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 (*cNtXdYsZ*), 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(8)` 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      /dev/dsk/c0t0d0s7
a          u              16                8192
a          u              16                8192
a          u              16                8192
```

Next Steps

Go to [Creating Disk Sets in a Cluster](#) 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 |
| Add drives to the disk sets. | How to Add Drives to a Disk Set |
| (Optional) Repartition drives in a disk set to allocate space to different slices. | How to Repartition Drives in a Disk Set |
| List DID pseudo-driver mappings and define volumes in the <code>/etc/lvm/md.tab</code> files. | How to Create an md.tab File |
| Initialize the <code>md.tab</code> files. | How to Activate Volumes |

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](#) 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\(8CL\)](#) 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](#).

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 Administering an Oracle Solaris Cluster 4.4 Configuration](#) 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 Administering an Oracle Solaris Cluster 4.4 Configuration](#).

- 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\(8CL\)](#) 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](#).

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

Ensure that the disk set has been created. For instructions, see [How to Create a Disk Set](#).

1. **Assume the root role.**
2. **List the DID mappings.**

```
phys-schost# cldevice show
```

- 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/d N`, 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/clt1d0
Full Device Path:               phys-schost-2:/dev/rdisk/clt1d0
...
```

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 (`c N t X d Y`) 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](#).

Otherwise, go to [How to Create an md.tab File](#) 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(8)` 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(8)` man page for details.

Next Steps

Define volumes by using an `md.tab` file. Go to [How to Create an md.tab File](#).

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 (`c N t X d Y`). The DID device name takes the form `/dev/did/rdisk/d N`.

```
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 [Sample md.tab File](#) 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.
2. The second line defines volume `d10`, the first submirror of `d0`, as a one-way stripe.
3. The third line defines volume `d20`, the second submirror of `d0`, as a one-way stripe.

```
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
```

Next Steps

Activate the volumes that are defined in the `md.tab` files. Go to [How to Activate Volumes](#).

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\(8\)](#) 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\(8\)](#) man page for more information.

8. 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 Administering an Oracle Solaris Cluster 4.4 Configuration](#).

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](#).

Otherwise, go to [How to Create UFS Cluster File Systems](#) 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
How to Add Mediator Hosts |
| Check the status of mediator data and, if necessary, fix bad mediator data. | How to Check For and Fix Bad Mediator Data |

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\(4D\)](#) 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](#).

How to Check For and Fix Bad Mediator Data

Ensure that you have added mediator hosts as described in [How to Add Mediator Hosts](#).

Perform this procedure to repair bad mediator data.

1. Display the status of the mediator data.

```
phys-schost# medstat -s setname
```

-s *setname*

Specifies the disk set name.

See the [medstat\(8\)](#) 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\(4D\)](#) 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 UFS Cluster File Systems](#).
- 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 [Planning and Administering Data Services for Oracle Solaris Cluster 4.4](#).

5

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 guide for that data service. For general information about creating a highly available local file system, see [Enabling Highly Available Local File Systems in *Planning and Administering Data Services for Oracle Solaris Cluster 4.4*](#).

The following sections are in this chapter:

- [Creating Cluster File Systems](#)
- [Creating Oracle ACFS File Systems](#)

Creating Cluster File Systems

This section provides information to create the UFS and ZFS cluster filesystems that support data services.

Perform these 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.

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](#).
- Ensure that you established the new cluster or cluster node as described in [Establishing a New Global Cluster or New Global-Cluster Node](#).
- If you are using a volume manager, ensure that volume management software is installed and configured.
- Determine the options to use for each ZFS pool that you want to create. See [Choosing Options for ZFS Cluster File Systems](#).
- 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](#).

How to Create UFS Cluster File Systems

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 <code>d1</code> within the <code>nfs</code> disk set |
| None | <code>/dev/global/rdisk/d1s3</code> | Raw disk device <code>d1s3</code> |

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\(8CL\)](#) 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\(8\)](#) man page or [mount\(8\)](#) 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 [Planning and Administering Data Services for Oracle Solaris Cluster 4.4](#).

Creating ZFS Cluster File Systems

This section describes how to create a ZFS cluster file system with and without using HAStoragePlus as well as how to configure encryption for a globally mounted ZFS file system.

- [Creating ZFS Cluster File Systems With and Without Using HAStoragePlus](#)
- [Configuring Encryption for a Globally Mounted ZFS File System](#)

Creating ZFS Cluster File Systems With and Without Using HAStoragePlus

For information about how to configure cluster file systems by using a ZFS storage pool with HAStoragePlus, see [How to Set Up an HAStoragePlus Resource for a zpool for Globally Mounted ZFS File Systems in *Planning and Administering Data Services for Oracle Solaris Cluster 4.4*](#). For information about how to use a ZFS storage pool to set up a cluster file system without using HAStoragePlus, see [How to Configure a zpool for Globally Mounted ZFS File Systems Without HAStoragePlus in *Administering an Oracle Solaris Cluster 4.4 Configuration*](#).

Configuring Encryption for a Globally Mounted ZFS File System

This section provides procedures that describe how to configure a globally mounted ZFS file system to use encryption.

For information about software requirements, see [Encrypting ZFS Cluster File Systems](#).

How to Create an Encrypted ZFS Cluster File System on a New Zpool

1. Create an encryption key.

Use the `cldevicegroup key-create` command to create an encryption key and specify its name.

```
phys-schost# cldevicegroup key-create -k keyname
```

2. Create an encrypted zpool.

Use the `zpool` command to create an encrypted zpool by specifying the `encryption` and `keysource` property values, the name of the zpool, and the name of the virtual disk device. The key name portion of the `keysource` property value specifies the encryption key that you created in the previous step.

```
phys-schost# zpool create -O encryption=on -O keysource=hex,cluster:keyname  
zpoolname virtual-disk-device
```

3. Place the encrypted zpool in a global cluster device group.

Use the `cldevicegroup create` command to add the encrypted zpool to a global cluster device group by specifying `poolaccess=global` and `import-at-boot=true` properties, the names of one or more cluster nodes, the `zpool` device type, and the name of the zpool.

```
phys-schost# cldevicegroup create -p poolaccess=global -p import-at-  
boot=true \  
-n node1,node2 -t zpool zpoolname
```

4. Bring the cluster device group online.

Use the `cldevicegroup online` command to activate the specified cluster device group.

```
phys-schost# cldevicegroup online zpoolname
```

Example 5-2 Creating an Encrypted Cluster File System on a New Zpool

This example shows how to create an encrypted cluster file system in a new zpool.

The first command creates a new encryption key called `gpool-key`:

```
phys-schost# cldevicegroup key-create -k gpool-key
```

The following command creates an encrypted zpool called `gpool` on the `c1t0d1` device. The command also specifies `gpool-key` as the `keyname` value of `keysource` property:

```
phys-schost# zpool create -O encryption=on -O keysource=hex,cluster:gpool-key gpool  
c1t0d1
```

The following command creates a device group for the `gpool` zpool that has the following characteristics:

- Specifies that the zpool has global access
- Imports the device group at boot time
- Specifies that the `node1` and `node2` nodes, in that order, are available to take over as the primary node for the device group
- Identifies the device type as `zpool`

```
phys-schost# cldevicegroup create -p poolaccess=global -p import-at-boot=true \  
-n node1,node2 -t zpool gpool
```

The final command imports the zpool on the preferred node and then mounts the encrypted cluster file system on both the `node1` and `node2` nodes:

```
phys-schost# cldevicegroup online gpool
```

How to Create an Encrypted ZFS Cluster File System on an Existing Unencrypted Zpool

Use this procedure when a device group already exists for the zpool for globally mounted ZFS file systems. Perform these steps from the node that is the primary for this device group or from the node where the zpool is imported. The zpool is imported to the node where the device group is primary.

1. Create an encryption key.

Use the `cldevicegroup key-create` command to create an encryption key and specify its name.

```
phys-schost# cldevicegroup key-create -k keyname
```

2. Create an encrypted cluster file system.

Run the `zfs create` command in the node on which the zpool for the globally mounted file system is imported to create an encrypted cluster file system by specifying the `keysource` value, setting the `encryption` property to `on`, and specifying the name of the zpool.

```
phys-schost# zfs create -o encryption=on -o keysource=hex,cluster:keyname  
zpoolname/fsname
```

Example 5-3 Creating an Encrypted Cluster File System on an Existing Unencrypted Zpool

The following command creates a new encryption key called `fs1-key`:

```
phys-schost# cldevicegroup key-create -k fs1-key
```

The following command, which must be run from the node where the zpool for the globally mounted ZFS file system is imported, uses the `fs1-key` encryption key to create an encrypted cluster file system called `gpool/fs1`.

```
phys-schost# zfs create -o encryption=on -o keysource=hex,cluster:fs1-key  
gpool/fs1
```

How to Create an Encrypted ZFS Cluster File System on an Existing Encrypted Zpool

Note that the encrypted child file system that you create inherits its encryption key information from the zpool.

- **Create a global encrypted ZFS file system.**

```
phys-schost# zfs create zpoolname/fsname
```

The following example command creates an encrypted cluster file system called `gpool/fs2`:

```
phys-schost# zfs create gpool/fs2
```

How to Remove an Encrypted ZFS Cluster File System and Its Encryption Key

1. **Identify the encryption key that is associated with the file system that you want to remove.**

```
phys-schost# zfs get -r keysource zpoolname/fsname
```

The following example command shows the `keysource` property value of the `gpool/fs1` file system:

```
phys-schost# zfs get -r keysource gpool/fs1  
NAME PROPERTY VALUE SOURCE  
gpool keysource hex,cluster:gpool-key local
```

2. **Remove the file system.**

Use the `zfs destroy` or the `zfs -r` command to remove all the child file systems of the specified file system.

```
phys-schost# zfs destroy zpoolname/fsname
```

The following example command removes the `gpool/fs1` file system:

```
phys-schost# zfs destroy gpool/fs1
```

3. **Remove the encryption key.**

Use the `cldevicegroup key-remove` command to remove the specified encryption key. Specify the key name of the encryption key that you identified in Step 1.

```
phys-schost# cldevicegroup key-remove -k key-name
```

The following example command removes the `gpool-key` encryption key:

```
phys-schost# cldevicegroup key-remove -k gpool-key
WARNING: Please ensure keyname 'gpool-key' is not being used in any offlined
Device Group.
Are you sure you want to remove keyname 'gpool-key' (y/n) [n]? y
```

How to List the Configured Keys in Oracle Solaris Cluster

- **Generate a list of key names that are configured in Oracle Solaris Cluster.**

```
phys-schost# cldevicegroup key-list
```

How to Back Up an Encryption Key

1. **Retrieve information about the key that you want to back up.**

```
phys-schost# cldevicegroup key-list -v
```

The following example command shows key information for the `gpool-key` encryption key:

```
phys-schost# cldevicegroup key-list -v
Keyname    Key
-----    ---
gpool-key  007781ebafcbd7fa94de683ae65f1e71293d48e5e32ba6bd2751a2e9299707d7
```

2. **Record the encryption key information.**

You can write down the key name and the key or save this information to a file that only you can read.

The following command examples save the key name and its associated encryption key value to a file called `~/osck` that only you can read:

```
phys-schost# cldevicegroup key-list -v >~/osck
phys-schost# chmod 400 ~/osck
```

How to Restore an Encryption Key

To restore a deleted encryption key, you must recreate the key by using the same key name and value as the key that you deleted. Ensure that you record or back up the key name and its value before you delete the key. For information about how to back up a key, see [How to Back Up an Encryption Key](#).

While deleting an encryption key renders associated shares inaccessible, you can make such shares accessible again by recreating the encryption key.

- **Restore a backed up key.**

```
phys-schost# cldevicegroup key-create -k keyname -c key
```

The following example command uses the `cldevicegroup key-create` command to recreate the `gpool-key` by specifying the `-c` option and the encryption key value:

```
phys-schost# cldevicegroup key-create -k gpool-key \
-c 007781ebafcbd7fa94de683ae65f1e71293d48e5e32ba6bd2751a2e9299707d7
```

How to Replace a ZFS Cluster File System Encryption Key

Use the `zfs key -c` command to replace an encrypted file system's wrapping key.

For more information about ZFS encryption, see [Encrypting ZFS File Systems in Managing ZFS File Systems in Oracle Solaris 11.4](#).

1. Obtain the list of encryption keys.

```
phys-schost# cldevicegroup key-list -v
```

The following example shows the existing encryption keys:

```
phys-schost# cldevicegroup key-list -v
Keyname      Key
-----      ---
DEC-2020-key
b268ae3f9beebbc12296aa4f472328b5deca0b33786eb92d25ccd4d2f465fc1c
```

2. Create a new encryption key by specifying a different key name.

```
phys-schost# cldevicegroup key-create -k new-keyname
```

The following example command creates a new encryption key called JAN-2021-key:

```
phys-schost# cldevicegroup key-create -k JAN-2021-key
```

3. Replace the encryption key for the existing global ZFS file system with the new key.

```
phys-schost# zfs key -c -o keysource=hex,cluster:new-keyname zpool
```

The following example command replaces the existing encryption key of the `gpool/fs1` file system with the JAN-2021-key key:

```
phys-schost# zfs key -c -o keysource=hex,cluster:JAN-2021-key gpool/fs1
```

4. Remove the obsolete encryption key.

```
phys-schost# cldevicegroup key-remove -k key-name
```

The following example command removes the obsolete encryption key, DEC-2020-key:

```
phys-schost# cldevicegroup key-remove -k DEC-2020-key
WARNING: Please ensure keyname 'DEC-2020-key' is not being used in any
offlined Device Group.
Are you sure you want to remove keyname 'DEC-2020-key' (y/n) [n]? y
```

5. Verify that the device group uses the new encryption key.

```
phys-schost# cldevicegroup key-list -v
```

The following example shows the new JAN-2021-key key and not the DEC-2020-key key that you removed:

```
phys-schost# cldevicegroup key-list -v
Keyname      Key
-----      ---
JAN-2021-key
052b4ce65bbda7e2402fbc965fde122aae0ce1d86581688e983f28aecc749de
```

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.

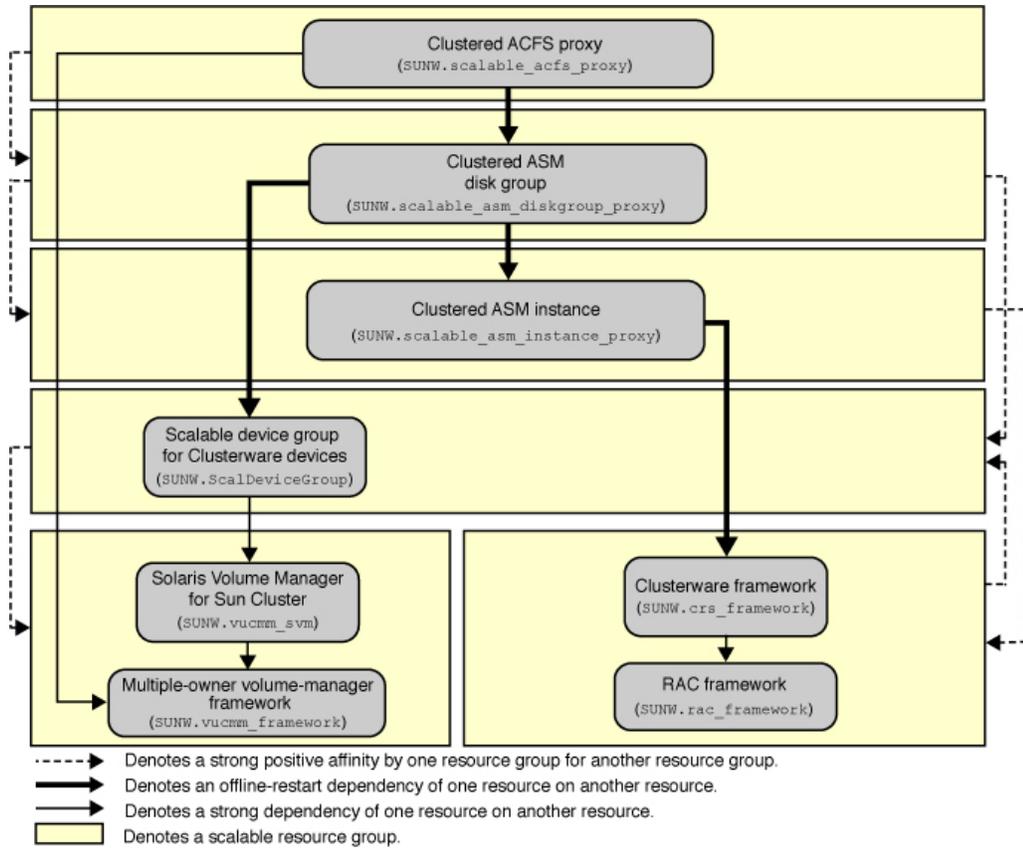
Table 5-1 Task Map: Creating Oracle ACFS File Systems

| Task | Instructions |
|--|---|
| Review Oracle ACFS configuration diagrams. | Sample Configurations of an Oracle ACFS File System |
| Register and configure the framework resource groups. | How to Register and Configure the Framework Resource Groups |
| Create an Oracle ACFS file system. | How to Create an Oracle ACFS File System |
| Register and configure the scalable device-group resource group. | How to Register and Configure the Scalable Device-Group Resource Group |
| Register and configure the Oracle ASM resource groups. | How to Register and Configure the Oracle ASM Resource Group |
| Configure interoperability between Oracle Grid Infrastructure and Oracle Solaris Cluster | How to Create an Oracle Grid Infrastructure Resource for Interoperation With Oracle Solaris Cluster |
| Register and configure the Oracle ACFS proxy resource group. | How to Register and Configure the Oracle ACFS Proxy Resource Group |

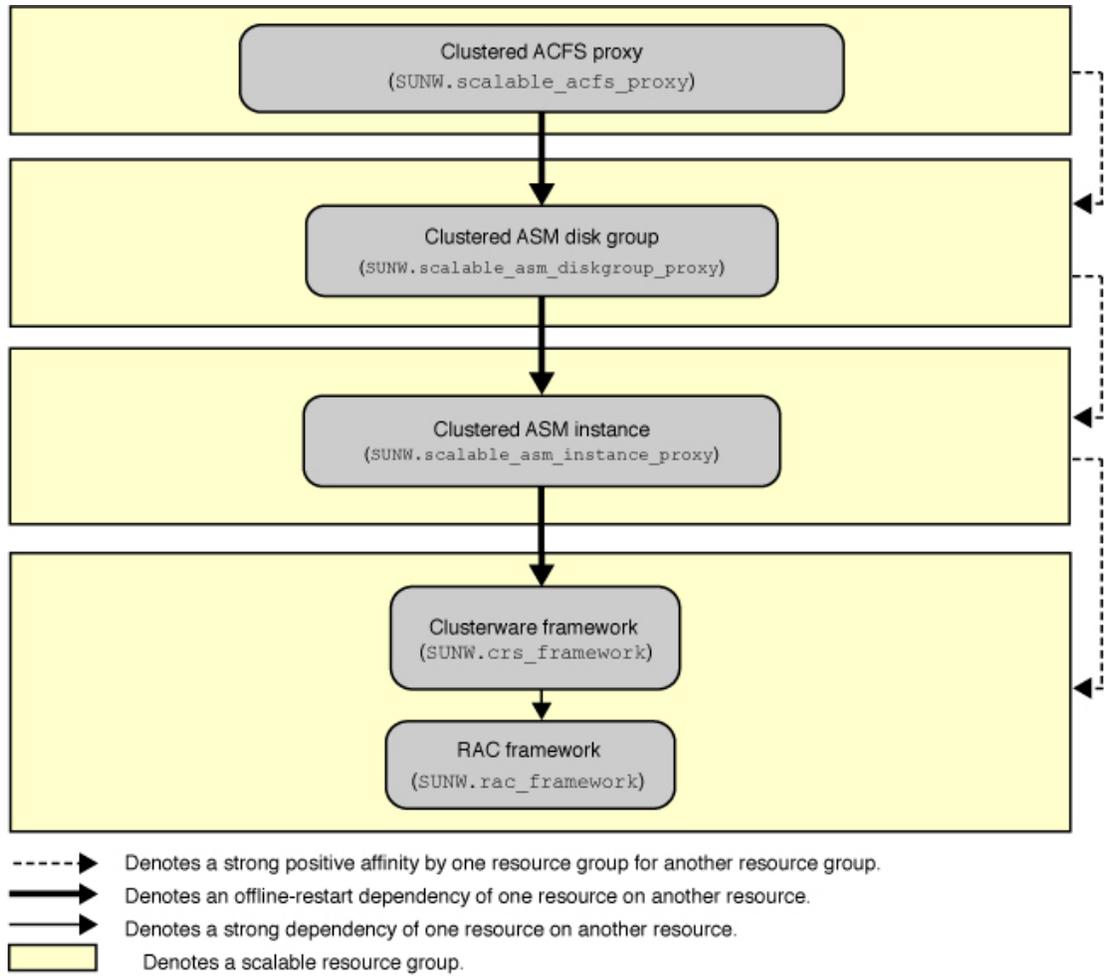
Sample Configurations of an Oracle ACFS File System

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.

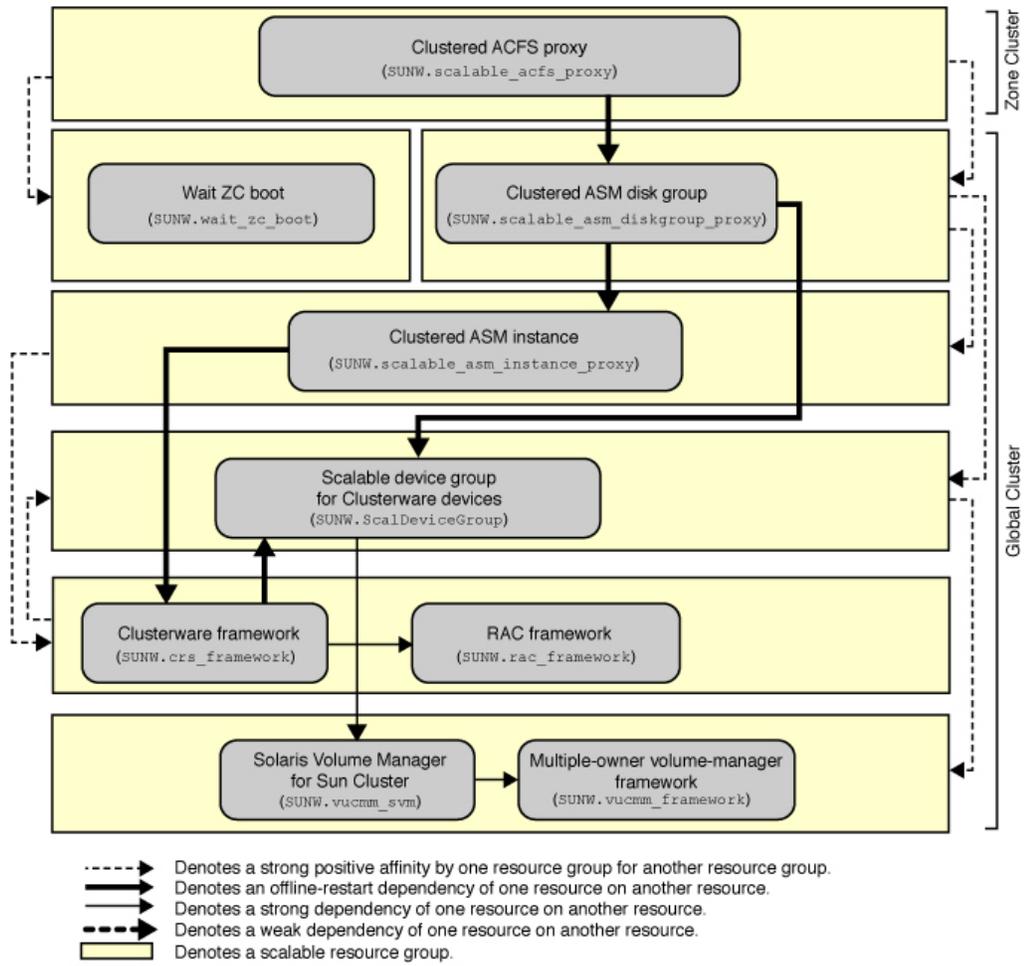
Configuration of Oracle ACFS in the Global Zone With Solaris Volume Manager



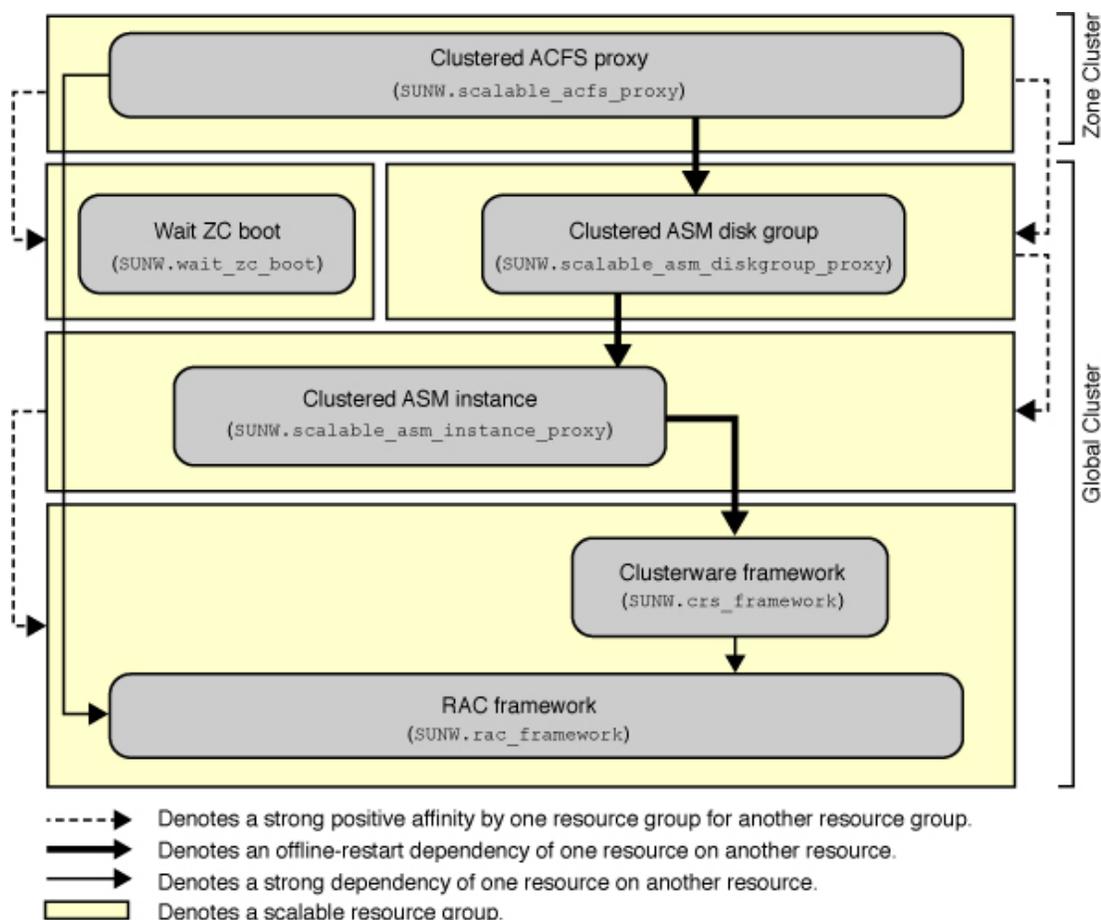
Configuration of Oracle ACFS in the Global Zone With Hardware RAID



Configuration of Oracle ACFS in a Zone Cluster with Solaris Volume Manager



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](#).

How to Create an Oracle ACFS File System

- Ensure that the resource framework groups are configured. See [How to Register and Configure the Framework Resource Groups](#).
- 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.

Use this procedure to create an Oracle ACFS file system. Perform all steps from the one node of the global cluster.

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](#).

Otherwise, go to [How to Register and Configure the Oracle ASM Resource Group](#).

How to Register and Configure the Scalable Device-Group Resource Group

Ensure that the Oracle ACFS file system is created. See [How to Create an Oracle ACFS File System](#).

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](#).

- 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`.

```
# 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](#).

How to Register and Configure the Oracle ASM Resource Group

- Ensure that the framework resource groups are created. See [How to Register and Configure the Framework Resource Groups](#).
- 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](#).

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.

- 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=+ASM*n*

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
```

- Bring online the `asm-dg-rg` resource group in a managed state on a cluster node.

```
# clresourcegroup online -eM asm-dg-rg
```

- 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](#).

How to Register and Configure the Oracle ACFS Proxy Resource Group

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](#).

Use this procedure to register and configure the Oracle ACFS proxy resource group. Perform all steps from one node of the global cluster.

- Assume the `root` role or become an administrator that provides `solaris.cluster.admin` and `solaris.cluster.modify` authorizations.
- Register the `SUNW.scalable_acfs_proxy` resource type.
 - If the file system is for use in the global cluster, use the following command:

```
# clresourcetype register SUNW.scalable_acfs_proxy
```
 - If the file system is for use by a zone cluster, use the following command:

```
# clresourcetype register -Z zonecluster SUNW.scalable_acfs_proxy
```
- 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 resource group in a managed state on a cluster node.

If the resource group is in a global cluster:

```
# clresourcegroup online -eM acfs-rg
```

If the resource group is in a zone cluster:

```
# clresourcegroup online -eM -Z zonecluster
acfs-rg
```

6. Verify the Oracle ACFS configuration.

For a global cluster:

```
# clresource status +
```

For a zone cluster:

```
# clresource status -Z zonecluster +
```

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](#).
- 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 [Planning and Administering Data Services for Oracle Solaris Cluster 4.4](#).
- To create an Oracle Grid Infrastructure resource, go to [How to Create an Oracle Grid Infrastructure Resource for Interoperation With Oracle Solaris Cluster](#).

How to Create an Oracle Grid Infrastructure Resource for Interoperation With Oracle Solaris Cluster

Ensure that you have registered and configured the Oracle ASM resource groups. See [How to Register and Configure the Oracle ASM Resource Group](#).

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.

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,pggrp: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,pggrp: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
```

6

Creating Zone Clusters

This chapter provides the following information to create and configure a zone cluster:

- [Overview of Creating and Configuring a Zone Cluster](#)
- [Creating and Configuring a Zone Cluster](#)

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 |
| If using Trusted Extensions, install and configure the Trusted Extensions software. | How to Install and Configure Trusted Extensions |
| Create the zone cluster. | How to Create a Zone Cluster (clsetup) |
| <i>(Optional)</i> Create a <code>solaris10</code> brand zone cluster. | How to Create a solaris10 Brand Zone Cluster (CLI) |
| <i>(Optional)</i> Add a file system for use by multiple nodes of the zone cluster. | Adding File Systems to a Zone Cluster |
| <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 |
| <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 |
| <i>(Optional)</i> Migrate an <code>solaris10</code> system to a node of a <code>solaris10</code> brand zone cluster. | Creating the Image for Directly Migrating Oracle Solaris 10 Systems Into Zones in <i>Creating and Using Oracle Solaris 10 Zones</i> |

 **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.

Table 6-1 (Cont.) Task Map: Create a Zone Cluster

| Task | Instructions |
|--|--|
| <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 . |
| <i>(Optional)</i> Validate the zone cluster. | How to Validate a Basic Cluster Configuration in Administering an Oracle Solaris Cluster 4.4 Configuration |

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](#)
- [Importing Zones into Zone Clusters](#)
- [Adding File Systems to a Zone Cluster](#)
- [Adding Local File Systems to a Specific Zone-Cluster Node](#)
- [Adding Storage Devices to a Zone Cluster](#)

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(8CL)` man page for more information.

You can also configure the Packet Filter (PF) feature in a zone cluster. See [How to Configure Packet Filter](#).

Note:

You cannot change the zone cluster name after the zone cluster is created. Also, once the zone cluster is configured, switching ip-type between exclusive and shared is not supported.

 **Note:**

When a zone is booted into a zone cluster, do not use the `zoneadm` command to halt or reboot the zone. Instead, use the `clzonecluster halt` and `clzonecluster reboot` commands. The `zoneadm halt` and `zoneadm reboot` commands do not synchronize the shutdown of objects that Oracle Solaris Cluster manages. So, these halting actions might cause applications in the zone or the zone to enter a faulted state.

This section contains the following procedures:

- [How to Install and Configure Trusted Extensions](#)
- [How to Create a Zone Cluster \(csetup\)](#)
- [How to Create a solaris10 Brand Zone Cluster \(CLI\)](#)
- [How to Configure a Zone Cluster to Use Trusted Extensions](#)
- [How to Configure a Zone Cluster to Be Immutable](#)

How to Install and Configure Trusted Extensions

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](#) 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 6, 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](#).

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](#).

Perform this procedure on each node in the global cluster.

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 /bin/true /usr/lib/zones/zoneshare
phys-schost# ln -s /bin/true /usr/lib/zones/zoneunshare
```

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. 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 using `ssh`.

```
phys-schost# svcadm enable ssh
```

- c. Modify the `/etc/pam.conf` file.

- i. Save a copy of the login file, then open the original file.

```
# cd /etc/pam.d
# cp login login.orig
# pfedit login
```

- ii. 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. Make the global-cluster node an LDAP client.

See [Chapter 6, Configuring LDAP for 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](#).

How to Create a Zone Cluster (`clsetup`)

- Create a global cluster. See [Establishing the Global Cluster](#).
- Read the guidelines and requirements for creating a zone cluster. See [Zone Clusters](#).
- If you plan to use a zone cluster configuration profile when creating a `solaris` or labeled brand zone cluster, ensure that the file is created and the file name has the `.xml` extension. See the Example section of the `clzonecluster(8CL)` man page for an example of the profile contents.
- 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](#).
- 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 *Administering an Oracle Solaris Cluster 4.4 Configuration*](#).
- 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 [Configurable Resource Types and Global Properties in *Oracle Solaris Zones Configuration Resources*](#). 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 Disaster Recovery Framework configuration. Otherwise, this requirement is optional. For more information about this Disaster Recovery Framework requirement, see [Disaster Recovery Framework](#).

- If applicable, the name of the public network management object that each zone-cluster node uses to connect to the public network. For a `solaris10` branded exclusive-IP zone cluster, you can only use an IPMP group as the public network management object.

 **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 Oracle Solaris Cluster Manager to create a zone cluster. For the browser interface log-in instructions, see [How to Access Oracle Solaris Cluster Manager in Administering an Oracle Solaris Cluster 4.4 Configuration](#).

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 Administering an Oracle Solaris Cluster 4.4 Configuration](#) and the `clzonecluster(8CL)` man page.

 **Note:**

You cannot change the zone cluster name after the zone cluster is created.

 **Note:**

When a zone is booted into a zone cluster, do not use the `zoneadm` command to halt or reboot the zone. Instead, use the `clzonecluster halt` and `clzonecluster reboot` commands. The `zoneadm halt` and `zoneadm reboot` commands do not synchronize the shutdown of objects that Oracle Solaris Cluster manages. So, these halting actions might cause applications in the zone or the zone to enter a faulted state.

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> . |

| Property | Description |
|--------------------------|--|
| brand= <i>brand-type</i> | <p>Specifies the solaris, solaris10, or labeled zones brand used in the zone cluster.</p> <div style="border-left: 2px solid #0070C0; border-right: 2px solid #0070C0; border-bottom: 2px solid #0070C0; padding: 10px; background-color: #E6F2FF;"> <p> Note:</p> <p>To use Trusted Extensions, you must use only the labeled brand. To create an exclusive-IP solaris10 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> </div> |
| ip-type= <i>value</i> | <p>Specifies the type of network IP address used by the zone cluster. Valid ip-type values are shared and exclusive. 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(8CL) man page.</p> |

| Property | Description |
|---|--|
| <p><code>enable_priv_net=value</code></p> | <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 (solaris brand only)</code> – 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> <p>If <code>net2</code> and <code>net3</code> are IB type network interfaces, the two IB partitions created for the zone cluster will have the names <code>net2_zone1_ibp0</code> and <code>net3_zone1_ibp0</code>.</p> |
| <p><code>enable_scalable_svc=value</code></p> | <p>The default value for this property is set to <code>false</code>.</p> <p>This property can be set only when the zone cluster is not in a running state. You will need to halt all the zone cluster nodes before you can change the property settings.</p> <p>If a scalable service is to be hosted on a zone cluster, this property must be set to <code>true</code> before the scalable service resource group and resources are created for the zone cluster.</p> <p>If this property has to be set to <code>false</code>, any scalable service already configured for a zone cluster must be removed before setting the property to <code>false</code>.</p> |

8. For a solaris10 brand zone cluster, enter a zone root password.

A `root` account password is required for a `solaris10` brand zone.

9. Choose the Zone System Resource Control property that you want to change.

You can set the following properties:

| Property | Description |
|-----------------------------------|---|
| <code>max-lwps=value</code> | Specifies the maximum number of lightweight processes (LWPs) simultaneously available to this zone cluster. |
| <code>max-shm-memory=value</code> | Specifies the maximum amount of shared memory in GBytes allowed for this zone cluster. |
| <code>max-shm-ids=value</code> | Specifies the maximum number of shared memory IDs allowed for this zone cluster. |
| <code>max-msg-ids=value</code> | Specifies the maximum number of message queue IDs allowed for this zone cluster. |
| <code>max-sem-ids=value</code> | Specifies the maximum number of semaphore IDs allowed for this zone cluster. |
| <code>cpu-shares=value</code> | Specifies the number of Fair Share Scheduler (FSS) shares to allocate to this zone cluster. |

10. Choose the Zone CPU Resource Control property that you want to change.

You can set the following properties:

| Property | Description |
|-------------------------------|---|
| <code>scope=scope-type</code> | Specifies whether the <code>ncpus</code> property used in a zone cluster is <code>dedicated-cpu</code> or <code>capped-cpu</code> . |
| <code>ncpus=value</code> | Specifies the limit for the scope type. <ul style="list-style-type: none"> If the <code>scope</code> property is set to <code>dedicated-cpu</code>, the <code>ncpus</code> 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(8) and poolcfg(8) man pages for more information on resource pools. If the <code>scope</code> property is set to <code>capped-cpu</code>, the <code>ncpus</code> 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 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(8) and poolcfg(8) man pages for more information on resource pools. |

11. 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 Oracle Solaris Cluster Manager to view the capped-cpu memory configuration of a zone cluster, as well as the dedicated-CPU configuration. For the browser interface log-in instructions, see [How to Access Oracle Solaris Cluster Manager in Administering an Oracle Solaris Cluster 4.4 Configuration](#).

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>192.0.2.1</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_ipmp0</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(8) man page for more information about the <code>defrouter</code> property. |
| <code>allowed-address=ipaddresses</code> | Specifies the IP addresses the exclusive-IP zone can use for the physical address. If not specified, then the exclusive-IP zone can use any IP address on the associated physical interface for the net resource. See the zonecfg(8) man page for more information about the <code>allowed-address</code> property. |
| <code>configure-allowed-address=true false</code> | If <code>configure-allowed-address</code> is set to <code>true</code> , the addresses specified by <code>allowed-address</code> are automatically configured on the interface each time the zone boots. When it is set to <code>false</code> , the <code>allowed-address</code> will not be configured on zone boot. By default, <code>configure-allowed-address</code> is set to <code>true</code> when an <code>allowed-address</code> is specified. When global net resources are present, this must be set to <code>false</code> . See the zonecfg(8) man page for more information about the <code>configure-allowed-address</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=192.0.2.1
set physical=net0
end
end
add net
set address=192.0.2.2
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 configuration 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` BUI.

```

phys-schost# txzonemgr

```

Select the global zone, then select the item, Configure per-zone name service.

19. If you typed No in Step 14, then 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 |
|------------------------------------|--|
| <code>-c config-profile.xml</code> | Includes system configuration information. The <code>-c config-profile.xml</code> option provides a configuration profile for all non-global zones of the zone cluster. All profiles must have a <code>.xml</code> extension. The contents of the file is derived from the <code>profile.xml</code> created by the <code>sysconfig(8) create-profile</code> subcommand. See the Example section of the clzonecluster(8CL) man page for an example of the profile contents. |
| <code>-M manifest.xml</code> | 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 <code>-M</code> 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 when using the `clzonecluster install` and the `clzonecluster install-cluster` commands.

When using the `clzonecluster install` command, use either the `-a` option or the `-d` option to install the `solaris10` image.

When using the `clzonecluster install-cluster` command, you can use the `-d`, `-s`, and `-p` options in the same command, to install cluster core packages, Geographic Edition software, and agents that are supported in the zone cluster, as well as patches.

 **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 |
|--|--|
| <pre>-a absolute_path_to_archive</pre> | <p>Specifies the absolute path to a solaris10 system archive to be used as the source image. The archive has to be accessible from all the nodes where the zone cluster is configured.</p> <pre># clzonecluster install [-n nodename[,...]] \ -a absolute_path_to_archive \ zone-cluster-name</pre> |
| <pre>-d absolute_directory_path</pre> | <p>Specifies the full directory path to the root directory of an installed solaris10 non-global zone. The path should be accessible on all the physical nodes of the cluster where the zone cluster will be installed.</p> <pre># clzonecluster install \ [-n nodename[,...]] \ -d absolute_directory_path zone-cluster-name</pre> |
| <pre>-d dvd-image- directory zone- cluster-name -p patchdir=patchdir[,pa tchlistfile=patchlist file] -s {all software- component</pre> | <div data-bbox="781 852 1365 1346" style="border: 1px solid #0070C0; padding: 10px; background-color: #E6F2FF;"> <p> Note:</p> <p>Oracle Solaris Cluster patch 145333-15 for SPARC and 145334-15 for x86 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. 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 <code>-p</code> option to install the patch.</p> </div> <p>The <code>-d</code> option specifies the full path to a DVD image directory for an Oracle Solaris Cluster release that supports the solaris10 brand zones. The cluster software DVD directory must be accessible from the global zone of the node where you run the command.</p> <p>In the <code>-p</code> option, <code>patchdir</code> specifies the directory of Oracle Solaris Cluster patches, and <code>patchlistfile</code> is a file that contains the list of patches in the <code>patchdir</code> directory to install. The <code>patchdir</code> directory is required, and must be accessible from inside the solaris10 brand zone on all nodes of the zone cluster. 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> <p>The <code>-s</code> option specifies the cluster software components that include the geo edition and data services, in addition to the core packages.</p> |
| | <pre># clzonecluster install-cluster \ -d dvd-image-directory \ [-p patchdir=patchdir[,patchlistfile=filename] \</pre> |

For more information, see the `clzonecluster(8CL)` man page.

20. If in Step 18, you did not use the `-c config-profile.xml` option when you installed the zone cluster, perform `sysid` configuration.

If in [Step 18](#), you *did* use the `-c config-profile.xml` option when you installed the zone cluster, you do not need to perform `sysid` configuration. Proceed to [Step 21](#).

 **Note:**

In the following steps, the non-global zone `zcnode` 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 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. Repeat Step b through Step 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 zcnode
```

- 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 b through Step 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, and add the underlying public network interface to the IPMP group. Since the underlying interface might have addresses associated with it, you must move the associated addresses to the IPMP group.

In each of the nodes of the zone cluster, configure the IPMP group and add an underlying public network interface to it. Delete any address that is already associated with the underlying interface as shown in the output of the `ipadm show-addr` command, and create it back on the IPMP interface.

```
zcnode# ipadm create-ipmp -i interface sc_ipmp0
zcnode# ipadm show-addr interface
zcnode# ipadm delete-addr interface/name
zcnode# ipadm create-addr -T static -a IPaddress/prefix sc_ipmp0/name
```

Note:

If the zone cluster's public networking interface is created over a global zone link aggregation or a global zone VNIC that is directly backed by a link aggregation, you do not need to create IPMP groups over it.

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 guide. See [Oracle Solaris Cluster 3.3 Documentation](#).

To complete Trusted Extensions configuration, go to [How to Configure a Zone Cluster to Use Trusted Extensions](#).

Otherwise, add file systems or storage devices to the zone cluster. See the following sections:

- [Adding File Systems to a Zone Cluster](#)
- [Adding Local File Systems to a Specific Zone-Cluster Node](#)
- [Adding Storage Devices to a Zone Cluster](#)

How to Create a `solaris10` Brand Zone Cluster (CLI)

Perform the following tasks:

- Ensure that all requirements in [Planning the Oracle Solaris Cluster Environment](#) are met.
- Select a zone image to migrate or install. The archive types that are supported for installing a zone cluster are the following:
 - Native brand zone on an Oracle Solaris10 system.
 - Cluster brand zone on an Oracle Solaris Cluster node with proper patch level, archive derived from a physical system installed with Oracle Solaris 10 software.
 - `solaris10` brand zone archive derived from an installed `solaris10` brand zone.
 - An Oracle Solaris 10 physical system.
 - An Oracle Solaris 10 physical cluster node.

For more information about `solaris10` brand zones, see [Creating and Using Oracle Solaris 10 Zones](#).

The `solaris10` brand zone cluster provides a virtualized Oracle Solaris 10 cluster environment in an Oracle Solaris 11 configuration. You can use the `solaris10` brand zone cluster to run or migrate cluster applications that are deployed on the Oracle Solaris 10 operating system, without any modification to the application.

1. Assume the `root` role on an active member node of a global cluster.

Perform all steps of this procedure from a node of the global cluster.

2. Create an archive and store it in a shared location.

```
# flarcreate -S -n s10-system -L cpio /net/mysharehost/share/s10-system.flar
```

This archiver format is NOT VALID for flash installation of ZFS root pool.

```
This format is useful for installing the system image into a zone.
Reissue command without -L option to produce an archive for root pool install.
Full Flash
Checking integrity...
Integrity OK.
Running precreation scripts...
Precreation scripts done.
Creating the archive...
6917057 blocks
Archive creation complete.
Running postcreation scripts...
Postcreation scripts done.

Running pre-exit scripts...
Pre-exit scripts done.
```

For more information about creating archives, see [Chapter 2, Assessing an Oracle Solaris 10 System and Creating an Archive in *Creating and Using Oracle Solaris 10 Zones*](#).

3. Configure the zone cluster.

Create and configure the zone cluster on the global cluster, as shown in the following example.

 **Note:**

The main difference between the `solaris` and `solaris10` brand zone cluster is setting the brand to `solaris10` and adding the `sysid` configuration.

```
# clnode status

=== Cluster Nodes ===

--- Node Status ---

Node Name                Status
-----                -
phys-host-1              Online
phys-host-2              Online

# cat ./s10-zc.config

create -b

set zonepath=/zones/s10-zc

set brand=solaris10

set autoboot=true

set bootargs="-m verbose"

add attr

set name=cluster

set type=boolean

set value=true

end

add node

set physical-host=phys-host-1

set hostname=zc-host-1

add net
```

```
set address=10.134.90.196/24

set physical=sc_ipmp0

end

end

add node

set physical-host=phys-host-2

set hostname=zc-host-2

add net

set address=10.134.90.197/24

set physical=sc_ipmp0

end

end

add sysid

set root_password=N4l3cWQb/s9zY

set name_service="DNS{domain_name=mydomain.com
name_server=13.35.24.52,13.35.29.41,14.13.8.13 search=mydomain.com}"

set nfs4_domain=dynamic

set security_policy=NONE

set system_locale=C

set terminal=vt100

set timezone=US/Pacific

end
```

In the above configuration, the `root_password` mentioned is `solaris`.

```
# clzonecluster configure -f ./s10-zc.config s10-zc

# clzonecluster verify s10-zc

# clzonecluster status s10-zc

=== Zone Clusters ===

--- Zone Cluster Status ---

Name          Brand          Node Name          Zone Host Name      Status
Zone Status
----          -
-----          -
-----          -

s10-zc        solaris10      phys-host-1        zc-host-1          offline
```

Configured

```

phys-host-2      zc-host-2
offline          Configured

```

4. Install the zone image for the zone cluster.

Use the zone image obtained in Step 3.

```
# clzonecluster install -a /net/mysharehost/share/s10-system.flar s10-zc
```

5. Install the cluster software.

Perform this step **only** if the archive does **not** contain cluster software in the image.

a. Boot the zone cluster into Offline/Running mode.

```
# clzonecluster boot -o s10-zc
```

b. Access the zone on all nodes of zone cluster and make sure that system configuration is complete.

```
# zlogin -C s10-zc
```

If the configuration is not complete, finish any pending system configuration.

c. From the global zone, check the zone cluster status.

```
# clzonecluster status s10-zc
```

```
=== Zone Clusters ===
```

```
--- Zone Cluster Status ---
```

| Name | Brand | Node Name | Zone Host Name |
|---------|-------------|-------------|----------------|
| Status | Zone Status | | |
| ---- | ----- | ----- | ----- |
| ----- | ----- | | |
| s10-zc | solaris10 | phys-host-1 | zc-host-1 |
| offline | Running | | |
| | | phys-host-2 | zc-host-2 |
| offline | Running | | |

d. Install the zone cluster software.

```
# clzonecluster install-cluster -d /net/mysharehost.com/osc-dir/ \
```

```
-p patchdir=/net/mysharehost/osc-dir,patchlistfile=plist-sparc \
```

```
-s all s10-zc
```

```
-p patchdir
```

Specifies the location of the patches to be installed along with the cluster software.

```
patchlistfile
```

Specifies the file that contains the list of patches to be installed inside the zone cluster along with the cluster software.

In this example, the contents of the file `plist-sparc` are as follows:

```
# cat /net/mysharehost/osc-dir/plist-sparc
```

```
145333-15
```

Note - Both the patchdir and patchlistfile locations must be accessible to all nodes of the cluster.

```
-s
```

Specifies the agent packages that should be installed along with core cluster software. In this example, all is specified to install all the agent packages.

6. Boot the zone cluster.

a. Reboot the zone cluster to boot the zone into Online/Running mode.

You might have to wait for some time to get the status to Online/Running.

```
# clzonecluster reboot s10-zc
```

b. From the global zone, check the zone cluster status.

The status of zone cluster will now be in Online/Running mode.

```
# clzonecluster status s10-zc
```

```
=== Zone Clusters ===
```

```
--- Zone Cluster Status ---
```

| Name | Brand | Node Name | Zone Host Name | Status |
|--------|-----------|-------------|----------------|--------|
| ---- | ----- | ----- | ----- | ----- |
| s10-zc | solaris10 | phys-host-1 | zc-host-1 | online |
| | | phys-host-2 | zc-host-2 | online |

7. Log into the zone.

```
# zlogin s10-zc
```

```
[Connected to zone 's10-zc' pts/2]
```

```
Last login: Mon Nov 5 21:20:31 on pts/2
```

8. Verify the status of the zone.

```
# /usr/cluster/bin/clnode status
```

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

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

| Node Name | Status |
|-----------|--------|
| zc-host-1 | Online |
| zc-host-2 | Online |

Next Steps

The solaris10 brand zone cluster configuration is now complete. You can now install and bring up any Oracle Solaris 10 applications and make them highly available by creating the necessary resources and resource groups.

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/nodeidN
```

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 ffffffff broadcast 10.11.166.255
groupname sc_ipmp0
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_ipmp0
ether 0:14:4f:24:74:d8
ce0:3: flags=1000843<UP,BROADCAST,RUNNING,MULTICAST,IPv4> mtu 1500 index
4
zone zcl
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 zcl
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 host name 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 Parameter \(IPv4 or IPv6\) in *Oracle Solaris 11.4 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](#)
- [Adding Local File Systems to a Specific Zone-Cluster Node](#)
- [Adding Storage Devices to a Zone Cluster](#)

See Also

If you want to update the software on a zone cluster, follow procedures in [Chapter 3, Installing and Updating Software Packages in *Updating Systems and Adding Software in Oracle Solaris 11.4*](#).

How to Configure a Zone Cluster to Be Immutable

Oracle Solaris Cluster 4.4 supports the use of immutable zone clusters, which are zone clusters with read-only roots. A read-only zone cluster can be configured by setting the `file-mac-profile` property. Using a read-only zone cluster root expands the secure runtime boundary.

Several values for `file-mac-profile` restrict access to all or part of the runtime environment from inside the zone. All of the profiles except `none` will cause the `/var/pkg` directory and its contents to be read-only from inside the zone. With zone clusters, the following options are available:

- `fixed-configuration` – Permits updates to `/var/*` directories, with the exception of directories that contain system configuration components. IPS packages, including new packages, cannot be installed. Persistently enabled SMF services are fixed. SMF manifests cannot be added from the default locations. Logging and auditing configuration files can be local. `syslog` and audit configuration are fixed.
- `flexible-configuration` – Permits modification of files in `/etc/*` directories, changes to `root`'s home directory, and updates to `/var/*` directories. IPS packages, including new packages, cannot be installed. Persistently enabled SMF services are fixed. SMF manifests cannot be added from the default locations. Logging and auditing configuration files can be local. `syslog` and audit configuration can be changed.
- `none` – This option is equivalent to an unset MWAC security policy.

By default, the `file-mac-profile` property is not set and the zone has a writable root dataset.

1. To change the `file-mac-profile` property, choose one of the following options:

- To change the zone cluster to a `fixed-configuration`, type:

```
root@dev1403:~# clzc configure zc1
clzc:zc1> info file-mac-profile
file-mac-profile:none
clzc:zc1> set file-mac-profile=fixed-configuration
clzc:zc1> info file-mac-profile
file-mac-profile:fixed-configuration
```

- To change the zone cluster to a `flexible-configuration`, type:

```
root@dev1403:~# clzc configure zc1
clzc:zc1> info file-mac-profile
file-mac-profile:none
clzc:zc1> set file-mac-profile=flexible-configuration
clzc:zc1> info file-mac-profile
file-mac-profile:flexible-configuration
```

- To change the zone cluster to a `none`, type:

```
clzc:zc1> set file-mac-profile=none
clzc:zc1> info file-mac-profile
file-mac-profile:none
```

2. To perform administration on immutable zone clusters, see [Administering an Immutable Zone Cluster by Making It Writable in Administering an Oracle Solaris Cluster 4.4 Configuration](#).

Importing Zones into Zone Clusters

This section provides procedures on how to import installed zones into an existing zone cluster or into a new zone cluster.

You can use the `import-zone` command to import the zones into an existing zone cluster or a new zone cluster. The `import-zone` command can be run in the

interactive mode as well as the non-interactive mode. For more information about the `import-zone` command, see the [clzonecluster\(8CL\)](#) man page.

Requirements and Restrictions

Consider the following points when you plan the importing of installed zones into an existing zone cluster or into a new zone cluster:

- When you import an installed zone into a new zone cluster, zone properties such as `zonepath`, `ip-type`, and `brand` must be defined in the zone cluster configuration before you run the `import-zone` command.
- The `import-zone` command can be executed only in the node scope.
- The physical host must be specified in the node scope before running the `import-zone` command.
- The `local:rad` SMF service, `svc:/system/rad:local` must be online on the physical host hosting the zone to be imported.
- The zone to be imported must be in an `installed` state.
- You cannot import an Oracle Solaris Zone that has zone properties such as `zonepath`, `ip-type`, and `brand` that are different from the zone cluster, unless you move the zone to the zone-cluster path using the Oracle Solaris `zoneadm move` command. See [Moving a Non-Global Zone](#).
- The zone to be imported cannot be on a shared storage.
- The zone to be imported cannot be configured as a failover highly available zone.
- The zone to be imported cannot be part of another zone cluster configuration.
- The Oracle Solaris Zone that is imported will be renamed to the zone cluster's name. You will be prompted to accept the renaming of the zone.
- When you run the `import-zone` command in the non-interactive mode, you will need to set an additional flag in order to rename the zone. For example:

```
add node
set physical-host=phost_name
import-zone -y zonename=zone_name
set hostname=host
end
```

How to Import an Installed Zone into an existing Zone Cluster

1. **Assume the `root` role on a node of the global cluster.**

You perform all steps of the procedure from a node of the global cluster.

2. **Verify that the zone that you want to import is in an installed state.**

For example:

```
phy-schost# zoneadm -z zone_name list -v
```

Ensure that you meet the requirements as described in [Requirements and Restrictions](#).

3. **Run the `clzonecluster configure` command on a configured zone cluster to set the necessary properties.**

```
phy-schost# clzonecluster configure zoneclustername
```

4. Enter the node scope property using the `add node` command.

```
clzc:zoneclustername:> add node
```

5. Set the physical host property in the node scope which hosts the zone to be imported using the `set physical-host` command.

```
clzc:zoneclustername:node> set physical-host=host
```

6. Run the `import-zone` command on the installed zone.

Accept the prompt to rename the zone.

```
clzc:zoneclustername:node> import-zone zonename=zone_name
```

If you are running the `import-zone` command in a non-interactive mode:

```
import-zone -y zonename=zone_name
```

7. Set the `hostname` in the node scope.

```
clzc:zoneclustername:node> set hostname=host
```

8. (Optional) Set other node scope properties and resources, if necessary.

For example:

```
# clzc:zoneclustername:node> add privnet
# clzc:zoneclustername:node:privnet> set physical=bge0
# clzc:zoneclustername:node:privnet> end
```

9. End the node scope and commit the zone configuration.**10. Verify that the zone is imported into the zone cluster.**

```
phy-schost# clzonecluster show -v zoneclustername
```

Example 6-1 Importing an Installed Zone into an Existing Zone Cluster

This example shows how to add an installed zone, `zone1`, into an existing zone cluster, `sczone`.

```
phy-schost# clzc configure sczone
clzc:sczone:> add node
clzc:sczone:node> set physical-host=phost2
clzc:sczone:node> import-zone zonename=zone1
This operation will rename the zone. Do you want to continue(Y/N)?
Y
clzc:sczone:node> set hostname=thost2-0a
clzc:sczone:node> end
clzc:sczone:> commit
clzc:sczone:> exit
#
```

How to Import an Installed Zone into a New Zone Cluster

1. Assume the `root` role on a node of the global cluster.

You perform all steps of the procedure from a node of the global cluster.

2. Verify that the zone that you want to import is in an installed state.

For example:

```
phy-schost# zoneadm -z zone_name list -v
```

Ensure that you meet the requirements as described in [Requirements and Restrictions](#).

3. Run the `clzonecluster configure` command on a configured zone cluster to set the necessary properties.

```
phy-schost# clzonecluster configure new-zoneclustername
```

4. Create the zone cluster using the `create` command.

```
clzc:new-zoneclustername:> create
```

5. Set the `zonepath`, `ip-type`, and `brand` properties of the zone cluster.

```
clzc:new-zoneclustername:> set zonepath=/zones/zonetc
clzc:new-zoneclustername:> set ip-type=shared
clzc:new-zoneclustername:> set brand=solaris
```

6. Enter node scope using the `add node` command.

```
clzc:new-zoneclustername:> add node
```

7. Set physical host in the node scope which hosts the zone to be imported using the `set physical-host` command.

```
clzc:new-zoneclustername:node> set physical-host=host
```

8. Run the `import-zone` command on the installed zone.

Accept the prompt to rename the zone.

```
clzc:new-zoneclustername:node> import-zone zonename=zone_name
```

If you are running the `import-zone` command in a non-interactive mode:

```
import-zone -y zonename=zone_name
```

9. Set the `hostname` in the node scope.

```
clzc:new-zoneclustername:node> set hostname=host
```

10. (Optional) Set other node scope properties and resources, if necessary.

For example:

```
# clzc:new-zoneclustername:node> add privnet
# clzc:new-zoneclustername:node:privnet> set physical=bge0
# clzc:new-zoneclustername:node:privnet> end
```

11. End the node scope and commit the zone configuration.

12. Verify that the zone is imported into the zone cluster.

```
phy-schost# clzonecluster show -v zoneclustername
```

Example 6-2 Importing an Installed Zone into a New Zone Cluster

This example shows how to import an installed zone, `zone1`, into a new zone cluster, `sczone-new`.

```
phy-schost# clzc configure sczone-new
clzc:sczone-new:> create
clzc:sczone-new:> set zonepath=/zones/zonetc
clzc:sczone-new:> set ip-type=shared
clzc:sczone-new:> set brand=solaris
clzc:sczone-new:> add node
clzc:sczone-new:node> set physical-host=phost2
clzc:sczone-new:node> import-zone zonename=zone1
This operation will rename the zone. Do you want to continue(Y/N)?
```

```
Y
clzc:sczone-new:node> set hostname=thost2-0a
clzc:sczone-new:node> end
clzc:sczone-new:> commit
clzc:sczone-new:> exit
#
```

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.HAStoragePlus` 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](#).

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 \(clsetup\)](#)
- [How to Add a Highly Available Local File System to a Zone Cluster \(CLI\)](#)
- [How to Add a ZFS Storage Pool to a Zone Cluster \(clsetup\)](#)
- [How to Add a ZFS Storage Pool to a Zone Cluster \(CLI\)](#)
- [How to Add a Oracle HSM Shared File System to a Zone Cluster \(CLI\)](#)
- [How to Add a Cluster File System to a Zone Cluster \(clsetup\)](#)
- [How to Add a ZFS Cluster File System to a Zone Cluster \(CLI\)](#)
- [How to Add a UFS Cluster File System to a Zone Cluster \(CLI\)](#)
- [How to Add an Oracle ACFS File System to a Zone Cluster \(CLI\)](#)

You can also use Oracle Solaris Cluster Manager to add a file system to a zone cluster. For the browser interface log-in instructions, see [How to Access Oracle Solaris Cluster Manager in Administering an Oracle Solaris Cluster 4.4 Configuration](#).

How to Add a Highly Available Local File System to a Zone Cluster (`clsetup`)

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 either the command line or Oracle Solaris Cluster Manager to perform this task. To use the command line to perform this task, see [How to Add a Highly Available Local File System to a Zone Cluster \(CLI\)](#). To use the Oracle Solaris Cluster Manager browser interface to perform this task, click Zone Clusters, click the zone cluster name to go to its page, click the Solaris Resources tab, then in the File Systems section click Add to start the file systems wizard. For Oracle Solaris Cluster Manager log-in instructions, see [How to Access Oracle Solaris Cluster Manager in Administering an Oracle Solaris Cluster 4.4 Configuration](#).

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 `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.**
5. **Choose the Add File System/Storage Device to a Zone Cluster menu item.**

The Zone Cluster Tasks Menu is displayed.

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 clsetup utility.**13. Verify the addition of the file system.**

```
phys-schost# clzonecluster 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 Planning and Administering Data Services for Oracle Solaris Cluster 4.4](#).

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 `clsetup` utility to perform this task. See [How to Add a Highly Available Local File System to a Zone Cluster \(clsetup\)](#). To add a ZFS pool to a zone cluster, instead perform procedures in [How to Add a ZFS Storage Pool to a Zone Cluster \(clsetup\)](#). 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 ZFS File System Highly Available in Planning and Administering Data Services for Oracle Solaris Cluster 4.4](#).

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 6-3 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/rdsk/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/rdsk/d1
type:                          ufs
options:                       [logging]
cluster-control:              [true]
...

```

Next Steps

Configure the file system to be highly available by using an `HASStoragePlus` resource. The `HASStoragePlus` 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 *Planning and Administering Data Services for Oracle Solaris Cluster 4.4*](#).

How to Add a ZFS Storage Pool to a Zone Cluster (`clsetup`)

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 11.4*](#) for procedures to create a ZFS pool.

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 either the command line or Oracle Solaris Cluster Manager to perform this task. To use the command line to perform this task, see [How to Add a ZFS Storage Pool to a Zone Cluster \(CLI\)](#). To use the Oracle Solaris Cluster Manager browser interface to perform this task, click Zone Clusters, click the zone cluster name to go to its page, click the Solaris Resources tab, then in the Datasets for ZFS Storage Pools section, click Add. For Oracle Solaris Cluster Manager log-in instructions, see [How to Access Oracle Solaris Cluster Manager in *Administering an Oracle Solaris Cluster 4.4 Configuration*](#).

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 `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 *Planning and Administering Data Services for Oracle Solaris Cluster 4.4*](#).

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 \(`clsetup`\)](#). 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 ZFS File System Highly Available in *Planning and Administering Data Services for Oracle Solaris Cluster 4.4*](#).

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 11.4](#) 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 6-4 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 `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 Planning and Administering Data Services for Oracle Solaris Cluster 4.4](#).

How to Add a Cluster File System to a Zone Cluster (`clsetup`)

Ensure that the cluster file system you want to add to the zone cluster is configured. See [Planning Cluster File Systems](#) and [Creating a Cluster File System](#).

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:

- ZFS cluster file system - The ZFS storage pool is imported in the global zone, and ZFS automatically mounts filesystem datasets configured in the pool.
- 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.
- Oracle HSM 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 either the command line or the Oracle Solaris Cluster Manager to perform this task. To 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 \(CLI\)](#)
- [How to Add a ZFS Cluster File System to a Zone Cluster \(CLI\)](#)
- [How to Add a Oracle HSM Shared File System to a Zone Cluster \(CLI\)](#)
- [How to Add an Oracle ACFS File System to a Zone Cluster \(CLI\)](#)

To use the Oracle Solaris Cluster Manager browser interface to perform this task, click Zone Clusters, click the zone cluster name to go to its page, click the Solaris Resources tab, then in the File Systems section click Add to start the file systems wizard. For Oracle Solaris Cluster Manager log-in instructions, see [How to Access Oracle Solaris Cluster Manager in Administering an Oracle Solaris Cluster 4.4 Configuration](#).

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. (For non-ZFS file systems only) 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 the [lofiadm\(8\)](#) man page.

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 *Planning and Administering Data Services for Oracle Solaris Cluster 4.4*](#).

How to Add a ZFS Cluster File System to a Zone Cluster (CLI)

This procedure shows how to add a ZFS 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 \(clsetup\)](#).

1. Assume the `root` role on a node of the global cluster that hosts the zone cluster.

Perform all steps of this procedure from a global cluster node.

2. On the global cluster, configure the ZFS pool and the file system that you want to use in the zone cluster.

```
phys-schost# zpool create poolname device
phys-schost# zfs create poolname/fsname
```

For information on how to set up a ZFS pool, see the `zfs(8)` and the `zpool(8)` man pages.

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=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
```

4. Verify the addition of the LOFS file system.

```
phys-schost# clzonecluster show -v zone-cluster-name
```

5. (Optional) Create an HAStoragePlus resource to manage and monitor the global ZFS pool in the global zone.

Although this step is considered to be optional, it is recommended to perform this step as a best practice.

```
phys-schost# clresource create -t HAStoragePlus \
-p GlobalZpools=poolname \
-g resource-group-name global-zpool-hastorageplus-resource-name
```

Bring the resource group online.

Proceed to Step 6 only if you skipped Step 5. Otherwise, proceed to Step 7.

6. **If the global ZFS pool is not managed by HAStoragePlus, then create the zpool device group and bring it online.**

```
phys-schost# cldevicegroup create -p poolaccess=global \
-n phys-schost-1,phys-schost-2 -t zpool global-zfs-pool
phys-schost# cldevicegroup online global-zfs-pool
```

For more information, see [How to Configure a zpool for Globally Mounted ZFS File Systems Without HAStoragePlus](#) in *Administering an Oracle Solaris Cluster 4.4 Configuration*.

7. **Ensure that the global zone file system is globally mounted by looking at the mount flags for the word, global.**

```
phys-schost# mount -p | grep global-cluster-mount-point
global-zfs-dataset - global-cluster-mount-point zfs - no
rw,devices,setuid,nonbmand,exec,
rstchown,noxattr,atime,global
```

8. **Configure an HAStoragePlus resource in the zone cluster to perform the loopback mount of the global file system in the zone cluster node.**

Note:

The global file system will be mounted on the node(s) where the HAStoragePlus is brought online. If you require the filesystem online on multiple zone cluster nodes, the resource group must be scalable.

```
ZC-schost# clresource create -g resource-group-name
-t HAStoragePlus \
-p FileSystemMountPoints=
zone-cluster-lofs-mountpoint hastorageplus-resource-name
```

9. **If an HAStoragePlus resource was created to manage and monitor the global ZFS pool in the global zone as shown in Step 5, then set the offline restart dependency.**

```
phys-schost# clresource set -Z zone-cluster-name \
-p Resource_dependencies_offline_restart=global:global-zpool-hastorageplus-
resource-name
\
hastorageplus-resource-name
```

Example 6-5 Adding a ZFS Cluster File System to a Zone Cluster (CLI)

The following example shows how to add a ZFS cluster file system in global ZFS poolglobalpool1, with mount point /globalpool1/apache to a zone cluster sczone. The file system is available to the zone cluster using the loopback mount mechanism at the mount point /zone/apache.

```
phys-schost-1# mount -p | grep 'globalpool1/apache'
...
```

```

globalpool1/apache - /globalpool1/apache zfs - no
rw,devices,setuid,nonbmand,exec,rstchown,noxattr,atime,global
...
phys-schost-1# clzonecluster configure sczone
clzc:zone-cluster-name> add fs
clzc:zone-cluster-name:fs> set dir=/zone/apache
clzc:zone-cluster-name:fs> set special=/globalpool1/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:                       /globalpool1/apache
raw:
type:                          lofs
options:                       []
cluster-control:              true
...
phys-schost-1# clresourcegroup create gpoolRG
phys-schost-1# clresource create -t HASStoragePlus -p GlobalZpools=globalpool1 \
-g gpoolRG gpoolResource
phys-schost-1# mount -p | grep '/globalpool1/apache'
globalpool1/apache - /globalpool1/apache zfs - no
rw,devices,setuid,nonbmand,exec,rstchown,noxattr,atime,global

phys-schost-1# clresource create -Z sczone -g zone-rg -t HASStoragePlus \
-p FileSystemMountPoints=/globalpool1/apache -p \
Resource_dependencies_offline_restart=global:gpoolResource zc-hasp-resource

```

Data service resources in the zone cluster that require access to the global ZFS file system must declare a resource dependency on this `HASStoragePlus` resource `zc-hasp-resource`.

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 \(clsetup\)](#).

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 the [lofiadm\(8\)](#) man page.

5. Verify the addition of the LOFS file system.

```
phys-schost# clzonecluster show -v zone-cluster-name
```

Example 6-6 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 Planning and Administering Data Services for Oracle Solaris Cluster 4.4](#).

How to Add a Oracle HSM Shared File System to a Zone Cluster (CLI)

Perform this task to add a Oracle HSM 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 \(clsetup\)](#).

At this time, Oracle HSM 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 Oracle HSM 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 Oracle HSM shared file system that you want to use in the zone cluster.**
Follow procedures for shared file systems in your Oracle HSM 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 6-7 Adding a Oracle HSM Shared File System as a Direct Mount to a Zone Cluster (CLI)

The following example shows the Oracle HSM 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)

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](#).

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 \(clsetup\)](#).

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](#).

This section contains the following procedures:

- [How to Add a Local File System to a Specific Zone-Cluster Node \(CLI\)](#)
- [How to Add a Local ZFS Storage Pool to a Specific Zone-Cluster Node \(CLI\)](#)

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 \(clsetup\)](#) or [How to Add a Highly Available Local File System to a Zone Cluster \(CLI\)](#).

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 6-8 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 \(clsetup\)](#) or [How to Add a ZFS Storage Pool to a Zone Cluster \(CLI\)](#).

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 6-9 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 to add a storage device to a zone cluster. For the browser interface log-in instructions, see [How to Access Oracle Solaris Cluster Manager in *Administering an Oracle Solaris Cluster 4.4 Configuration*](#).

This section contains the following procedures:

- [How to Add a Global Storage Device to a Zone Cluster \(clsetup\)](#)
- [How to Add a DID Device to a Zone Cluster \(CLI\)](#)
- [How to Add a Disk Set to a Zone Cluster \(CLI\)](#)
- [How to Add a Raw-Disk Device to a Specific Zone-Cluster Node \(CLI\)](#)

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 either the command line or the Oracle Solaris Cluster Manager browser interface to perform this task. To use the command line to perform this task, see the following procedures:

- [How to Add a DID Device to a Zone Cluster \(CLI\)](#)
- [How to Add a Disk Set to a Zone Cluster \(CLI\)](#)

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 \(CLI\)](#). To use the Oracle Solaris Cluster Manager browser interface to perform this task, click Zone Clusters, click the zone cluster name to go to its page, click the Solaris Resources tab, then in the Storage Devices section click Add to start the storage device wizard. For Oracle Solaris Cluster Manager log-in instructions, see [How to Access Oracle Solaris Cluster Manager in Administering an Oracle Solaris Cluster 4.4 Configuration](#).

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 Administering an Oracle Solaris Cluster 4.4 Configuration](#).

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 \(clsetup\)](#).

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
```

Note:

To reduce downtime, you can reboot the zone cluster nodes one by one. See [Shutting Down and Booting a Single Node in a Cluster in *Administering an Oracle Solaris Cluster 4.4 Configuration*](#).

Example 6-10 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 \(clsetup\)](#).

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/*dsk/*
```

Specifies the full logical device path of the disk set

```
match=/dev/md/shared/N/*dsk/*
```

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 6-11 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/*dsk/*
clzc:sczone:device> end
clzc:sczone> add device
clzc:sczone:device> set match=/dev/md/shared/3/*dsk/*
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 \(clsetup\)](#) or [How to Add a DID Device to a Zone Cluster \(CLI\)](#).

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 6-12 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
```

7

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 *Administering an Oracle Solaris Cluster 4.4 Configuration*](#).

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](#)
- [How to Uninstall Oracle Solaris Cluster Quorum Server Software](#)
- [How to Uninstall Oracle Solaris Cluster Software From a Cluster Node](#)
- [How to Unconfigure a Zone Cluster](#)
- [How to Uninstall Cluster Software Packages from an Automated Installer Server](#)
- [How to Uninstall a Specific Oracle Solaris Cluster Package](#)



Caution:

The `/var/cluster.sav/` directory is created and used during current and future Oracle Solaris Cluster installation, update, and uninstallation operations. This directory and its contents are for Oracle Solaris Cluster internal use only, and must not be deleted by the user. User removal of this directory, or any of its contents, would pose a high risk of compromising the cluster during any future installation, update, or uninstallation of Oracle Solaris Cluster software.

How to Unconfigure Oracle Solaris Cluster Software to Correct Installation Problems

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.

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](#), do not perform this procedure. Instead, go to [How to Uninstall Oracle Solaris Cluster Software From a Cluster 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 Administering an Oracle Solaris Cluster 4.4 Configuration](#) 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 [About Run Level Booting in Booting and Shutting Down Oracle Solaris 11.4 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\(8CL\)](#) 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\(8\)](#) 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 Task Map: Installing the Software. 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 Managing Hardware With Oracle Solaris Cluster 4.4](#) and the removal procedure in the Oracle Solaris Cluster guide for your storage array.

How to Uninstall Oracle Solaris Cluster Software From a Cluster Node

Perform this procedure to unconfigure Oracle Solaris Cluster software from a global-cluster node before you disconnect it from a fully established cluster configuration. You can use this procedure to uninstall software from the last remaining node of a cluster.

Note:

To uninstall Oracle Solaris Cluster software from a node that has not yet joined the cluster or is still in installation mode, do not perform this procedure. Instead, go to [How to Unconfigure Oracle Solaris Cluster Software to Correct Installation Problems in *Installing and Configuring an Oracle Solaris Cluster 4.4 Environment*](#).

The `phys-schost#` prompt reflects a global-cluster prompt. Perform this procedure on a global cluster.

This procedure provides the long forms of the Oracle Solaris Cluster commands. Most commands also have short forms. Except for the long and short forms of the command names, the commands are identical.

1. Ensure that you have correctly completed all prerequisite tasks in the task map to remove a cluster node.

See [Removing a Node From a Cluster in *Administering an Oracle Solaris Cluster 4.4 Configuration*](#).

Ensure that you have removed the node from the cluster configuration by using `clnode remove` before you continue with this procedure. Other steps might include adding the node you plan to uninstall to the cluster's node-authentication list, uninstalling a zone cluster, and so on.

Note:

To unconfigure the node but leave Oracle Solaris Cluster software installed on the node, do not proceed further after you run the `clnode remove` command.

2. Assume the `root` role on the node to uninstall.

3. If your node has a dedicated partition for the global devices name space, reboot the global-cluster node into noncluster mode.

- On a SPARC based system, run the following command.

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

- On an x86 based system, run the following commands.

```
# shutdown -g0 -y -i0
...
<<< Current Boot Parameters >>>
Boot path: /pci@0,0/pci8086,2545@3/pci8086,1460@1d/pci8086,341a@7,1/
sd@0,0:a
Boot args:
```

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

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

4. In the `/etc/vfstab` file, remove all globally mounted file-system entries except the `/global/.devices` global mounts.

5. Reboot the node into noncluster mode.

- On SPARC based systems, perform the following command:

```
ok boot -x
```

- On x86 based systems, perform the following commands:

- 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 [About Run Level Booting in Booting and Shutting Down Oracle Solaris 11.4 Systems](#).

- In the boot parameters screen, use the arrow keys to select the `kernel` entry and type `e` to edit the entry.
- Add `-x` to the command to specify that the system boot into noncluster mode.
- Press Enter to accept the change and return to the boot parameters screen. The screen displays the edited command.
- 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.

6. 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 /
```

7. To unconfigure the node and remove Oracle Solaris Cluster software, run the following command:

```
phys-schost# scinstall -r [-b bename]
```

-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 *bootenvironmentname*

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(8)` man page for more information.

8. **If you intend to reinstall the Oracle Solaris Cluster software on this node after the uninstall completes, reboot the node to boot into the new boot environment.**
9. **If you do not intend to reinstall the Oracle Solaris Cluster software on this cluster, disconnect the transport cables and the transport switch, if any, from the other cluster devices.**
 - a. **If the uninstalled node is connected to a storage device that uses a parallel SCSI interface, install a SCSI terminator to the open SCSI connector of the storage device after you disconnect the transport cables.**

If the uninstalled node is connected to a storage device that uses Fibre Channel interfaces, no termination is necessary.
 - b. **Follow the documentation that shipped with your host adapter and server for disconnection procedures.**

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\(8CL\)](#) 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 Oracle Solaris Cluster Quorum Server Software

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 Administering an Oracle Solaris Cluster 4.4 Configuration](#).

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 Administering an Oracle Solaris Cluster 4.4 Configuration](#).
- On each quorum-server host computer, stop the quorum server by following the steps in [How to Stop a Quorum Server in Administering an Oracle Solaris Cluster 4.4 Configuration](#).

1. **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. **Clean up or remove the quorum server directories.**

By default, this directory is `/var/scqsd`.

How to Uninstall Cluster Software Packages from an Automated Installer Server

Perform this task to remove Oracle Solaris Cluster software installation packages from an Automated Installer (AI) installation server.

1. **On the AI server, assume a role that provides `solaris.cluster.admin` authorization.**
2. **Uninstall Oracle Solaris Cluster software packages from the AI server.**

```
installserver# pkg uninstall ha-cluster/*
```

 **Caution:**

This command uninstalls all `ha-cluster` packages that are installed on the AI server.

How to Uninstall a Specific Oracle Solaris Cluster Package

Perform this task to uninstall one or more individual Oracle Solaris Cluster software packages.

1. **On the cluster node where the package to remove is installed, assume a role that provides `solaris.cluster.admin` authorization.**
2. **If the package to remove is associated with an SMF service, ensure that the service is disabled.**

Check the status of the SMF service. If the service is online, disable it, then check the status again to verify that the SMF service is disabled. Checking for a disabled service might take more time to complete.

```
phys-schost# svcs -a | grep service-name
online 15:50:26 svc:/service-name:default
phys-schost# svcadm disable svc:/service-name:default
```

```
phys-schost# svcs -a | grep service-name
disabled 15:55:33 svc:/service-name:default
```

For a list of Oracle SolarisCluster SMF services, see [Oracle Solaris OS Feature Restrictions in Administering an Oracle Solaris Cluster 4.4 Configuration](#).

3. **Check whether the package to remove has any dependencies.**

```
phys-schost# pkg content -o type,pkg.name -t depend -a fmri=package '*'
TYPE PKG.NAME
group package-1
require package-2
```

4. **If the package to remove has a require dependency with another package, such as group package `ha-cluster-full` uninstall that package.**

Uninstallation of the package to remove would fail if its require dependency package still exists. Uninstalling a group package does not uninstall its dependency packages.

```
phys-schost# pkg uninstall require-dependency-package
```

 **Note:**

A group dependency with the package to uninstall does not prevent uninstallation of the package. When the package is uninstalled, it is put in the avoid list of the group dependency. For more information about uninstalling packages with dependencies, see [Uninstalling Packages in Updating Systems and Adding Software in Oracle Solaris 11.4](#).

5. **Uninstall the package to remove.**

```
phys-schost# pkg uninstall package
```

6. **Verify that the package is uninstalled.**

```
phys-schost# pkg list package
pkg list: no packages matching 'package' installed
```

Example 7-1 Uninstalling the Oracle Solaris Cluster Manager Packages

The following example shows the removal from one cluster node of the Oracle Solaris Cluster rManager software packages, `ha-cluster/system/manager` and `ha-cluster/system/manager-wls`. Each package has an associated SMF service, `system/cluster/manager` and `system/cluster/manager-wls`, that is first disabled. The packages have a require dependency on the `ha-cluster-full` group package.

Check for active SMF services.

```
phys-schost-1# svcs -a | grep manager
online 15:50:26 svc:/system/cluster/manager:default
online 15:50:26 svc:/system/cluster/manager-wls:default
```

Disable the manager SMF service and verify it is disabled.

```
phys-schost-1# svcadm disable svc:/system/cluster/manager:default
phys-schost-1# svcs -a | grep manager
disabled 15:53:33 svc:/system/cluster/manager:default
online 15:50:26 svc:/system/cluster/manager-wls:default
```

Disable the manager-wls SMF service and verify it is disabled.

```
phys-schost-1# svcadm disable svc:/system/cluster/manager-wls:default
phys-schost-1# svcs -a | grep manager
disabled 15:53:33 svc:/system/cluster/manager:default
disabled 15:56:43 svc:/system/cluster/manager-wls:default
```

Check for package dependencies.

```
phys-schost-1# pkg content -o type, pkg.name -t depend -a fmri=ha-cluster/system/
manager
' * '
TYPE PKG.NAME
require ha-cluster-full
...
phys-schost-1# pkg content -o type, pkg.name -t depend -a fmri=ha-cluster/system/
managerwls
' * '
TYPE PKG.NAME
require ha-cluster-full
...
```

Uninstall the require dependency package.

```
phys-schost-1# pkg uninstall ha-cluster-full
```

Uninstall the Oracle Solaris Cluster Manager packages.

```
phys-schost-1# pkg uninstall ha-cluster/system/manager ha-cluster/system/manager-
wls
```

Verify that the packages are no longer installed.

```
phys-schost# pkg list manager
pkg list: no packages matching 'manager' installed
phys-schost# pkg list manager-wls
pkg list: no packages matching 'manager-wls' installed
```

Index

Symbols

`/etc/hosts.allow` setting, [3-11](#), [3-16](#), [3-30](#), [3-36](#)

`/etc/inet/hosts` file

- configuring, [2-6](#)
- planning, [1-6](#)

`/etc/inet/ntp.conf` file

- configuring, [3-73](#)

`/etc/inet/ntp.conf.include` file

- creating, [3-74](#), [3-75](#)

`/etc/inet/ntp.conf.sc` file

- configuring, [3-74](#), [3-75](#)

`/etc/lvm/md.tab` file, [4-7](#)

`/etc/vfstab` file

- adding mount points, [5-2](#)
- verifying the configuration, [5-3](#)

`/global` directory, [1-31](#)

`/opt/SUNWcluster/bin/pconsole` utility

- installing the software, [2-8](#)

`/usr/cluster/bin/claccess` command

- adding nodes to the authorized-node list, [7-2](#)
- removing nodes from the authorized-node list, [3-57](#)

`/usr/cluster/bin/cldevice` command

- updating the `global-devices` namespace, [4-3](#)
- verifying command processing, [4-3](#)

`/usr/cluster/bin/clnode` command

- viewing private host names, [3-73](#)

`/usr/cluster/bin/clquorumserver` command

- starting the quorum server, [2-12](#)

`/usr/cluster/bin/clresource` command

- disabling resources, [7-7](#)
- listing resources, [7-6](#)
- taking resource groups offline

 - CR 6349292, [7-6](#)

`/usr/cluster/bin/clresourcegroup` command

- deleting resource groups, [7-7](#)
- unmanaging resource groups, [7-7](#)

`/usr/cluster/bin/clsetup` utility

- adding
 - cluster interconnects, [3-48](#)
 - file systems to a zone cluster, [6-34](#)
 - global storage devices to a zone cluster, [6-46](#)
 - highly available local file systems to a zone cluster, [6-28](#)
 - ZFS storage pools to a zone cluster, [6-32](#)
- changing private host names, [3-73](#)
- creating a zone cluster, [6-6](#)
- navigation tip, [6-6](#)
- post installation setup, [3-69](#)

`/usr/cluster/bin/cluster check` command

- validating the cluster, [3-78](#)

`/usr/cluster/bin/cluster` command

- adding
 - nodes, [3-60](#)
 - zone clusters to the private IP-address range, [1-14](#), [3-49](#)
- creating new global clusters, [3-13](#)

`/usr/cluster/bin/clzonecluster` command

- booting the zone cluster, [6-16](#)
- halting the zone cluster, [7-7](#)
- installing the zone cluster, [6-12](#)

`/usr/cluster/bin/scinstall` command

- adding nodes, [3-53](#)
- creating the global cluster, [3-4](#)
- creating the global cluster by using Automated Installer, [3-19](#)
- unconfiguring Oracle Solaris Cluster software, [7-1](#)

A

adapters

- local MAC address, [1-8](#)
- Oracle VM Server for SPARC guest domains, [1-16](#)
- tagged VLAN
 - cluster interconnect guidelines, [1-16](#)
 - public-network guidelines, [1-7](#)

adapters (*continued*)

VLAN

private-network guidelines, [1-14](#)

adding, [2-2](#)

cluster-file-system mount points to new nodes, [2-6](#)

drives to a disk set, [4-5](#)

file systems to a zone cluster, [6-28](#)

local file systems to a zone cluster, [6-43](#)

mediator hosts, [4-10](#)

storage devices to zone clusters, [6-46](#)

address property, [6-11](#)

admin_low template

Trusted Extensions, [6-4](#)

administrative console

installing pconsole software, [2-8](#)

IP addresses, [1-7](#)

affinities

scalable device-group resource groups, [5-17](#)

AI, [3-19](#)

allowed-address property, [6-11](#)

applications

installing on a cluster file system, [1-27](#)

archiveadm

create an archive, [3-32](#), [3-37](#)

authentication

authorized-node list, [3-57](#)

authorized-node list

adding nodes, [7-2](#)

removing nodes, [3-57](#)

Automated Installer

installing Oracle Solaris and Oracle Solaris Cluster software, [3-19](#)

manifest, [3-27](#), [3-34](#), [3-39](#)

troubleshooting installation, [3-32](#)

automatic power-saving shutdown

restriction, [1-3](#)

automatic reboot on disk-path failures

disk-path failures, [3-10](#)

B

booting

into noncluster mode, [7-2](#)

brand property, [6-8](#)

brand zones

zone clusters, [1-23](#)

BUI

tasks you can perform

add a file system to a zone cluster, [6-28](#)

add storage a zone cluster, [6-46](#)

create a zone cluster, [6-6](#)

view capped memory for a zone cluster, [6-10](#)

C

capped-cpu scope type

zone clusters, [6-10](#)

centralized installer, [2-16](#)

changing

private hostnames, [3-72](#)

private IP-address range, [3-49](#)

cipso template

zone cluster IP addresses, [6-23](#)

claccess command

adding nodes to the authorized-node list, [7-2](#)

removing nodes from the authorized-node list, [3-57](#)

cldevice command

updating the global-devices namespace, [4-3](#)

verifying command processing, [4-3](#)

clinstall program

installing software packages, [2-16](#)

clnode command

viewing private host names, [3-73](#)

clquorumserver command

starting the quorum server, [2-12](#)

clresource

group command, [5-14](#)

clresource command

disabling resources, [7-7](#)

listing resources, [7-6](#)

taking resource groups offline, [7-6](#)

clresourcegroup command

bringing online resource groups, [5-15](#)

creating

Oracle ACFS resource group, [5-20](#)

scalable resource group, [5-17](#)

SUNW.wait_zc_boot resource group, [5-20](#)

deleting resource groups, [7-7](#)

unmanaging resource groups, [7-7](#)

clsetup utility

adding

cluster interconnects, [3-48](#)

file systems to a zone cluster, [6-34](#)

global storage devices to a zone cluster, [6-46](#)

highly available local file systems to a zone cluster, [6-28](#)

ZFS storage pools to a zone cluster, [6-32](#)

changing private hostnames, [3-73](#)

creating a zone cluster, [6-6](#)

navigation tip, [6-6](#)

post installation setup, [3-69](#)

cluster check command

validating the cluster, [3-78](#)

- cluster check command (*continued*)
 - vfstab file check, [5-3](#)
 - cluster command
 - adding
 - nodes, [3-60](#)
 - zone clusters to the private IP-address range, [1-14](#), [3-49](#)
 - creating new global clusters, [3-13](#)
 - cluster file systems
 - adding to a new node, [2-6](#)
 - adding to a zone cluster, [6-39](#)
 - adding to a zone cluster with clsetup, [6-34](#)
 - installing applications, [1-27](#)
 - LOFS restriction, [1-26](#)
 - mount options, [5-2](#)
 - planning, [1-25](#), [1-31](#)
 - restrictions
 - communication end points, [1-27](#)
 - fattach command, [1-27](#)
 - forcedirectio, [1-32](#)
 - LOFS, [1-3](#)
 - quotas, [1-26](#)
 - zone clusters, [1-26](#)
 - zone file system, [1-26](#)
 - verifying the configuration, [5-3](#)
 - cluster file systems, UFS
 - configuring, [5-1](#)
 - cluster file systems, ZFS
 - configuring, [5-4](#)
 - cluster interconnects
 - changing the private IP-address range, [3-49](#)
 - configuring on a single-host cluster, [3-48](#)
 - planning, [1-15](#)
 - cluster name, [1-12](#)
 - cluster nodes, [1-12](#)
 - clzonecluster command
 - booting the zone cluster, [6-16](#)
 - halting the zone cluster, [7-7](#)
 - installing the zone cluster, [6-12](#)
 - communication end points
 - restriction on cluster file systems, [1-27](#)
 - Compatibility Guide, [2-3](#)
 - config/enable_tcpwrappersdisabling, [3-8](#), [3-13](#), [3-55](#), [3-61](#)
 - config/enable_tcpwrappersenabling, [3-57](#), [3-62](#)
 - configuration guide, [2-3](#)
 - configure-allowed-address property, [6-11](#)
 - configuring
 - additional global-cluster nodes
 - by using an XML file, [3-60](#)
 - by using scinstall, [3-53](#)
 - cluster interconnects on a single-host cluster, [3-48](#)
 - disk sets, [4-3](#)
 - configuring (*continued*)
 - IPMP groups, [2-7](#)
 - md.tab file, [4-7](#)
 - Network Time Protocol (NTP), [3-73](#)
 - new global clusters
 - by using an XML file, [3-13](#)
 - by using Automated Installer, [3-19](#)
 - by using scinstall, [3-4](#)
 - Oracle ACFS file systems, [5-8](#)
 - Packet Filter feature of Oracle Solaris, [2-25](#)
 - quorum devices, [3-67](#)
 - quorum server software, [2-10](#)
 - Solaris Volume Manager, [4-1](#)
 - state database replicas, [4-2](#)
 - troubleshooting, [3-12](#), [3-19](#), [3-59](#), [3-63](#)
 - Trusted Extensions for use with zone clusters, [6-3](#), [6-22](#)
 - Trusted Extensions software, [6-3](#)
 - UFS cluster file systems, [5-1](#)
 - user work environment, [2-24](#)
 - ZFS cluster file systems, [5-4](#)
 - zone clusters, [6-2](#)
 - console-access devices
 - IP addresses, [1-7](#)
 - planning, [1-7](#)
 - cpu-shares property, [6-10](#)
 - create
 - archive, [3-32](#)
 - creating
 - multiple-owner volume-manager framework resource group, [5-14](#)
 - Oracle RAC framework resource group, [5-14](#)
 - solaris10 brand zone clusters, [6-17](#)
 - zone clusters, [6-6](#)
- ## D
-
- data services
 - support in solaris10 brand zone clusters, [6-13](#)
 - data-services installation
 - by using the clinstall command, [2-16](#)
 - by using the pkg command, [2-20](#)
 - dedicated-cpu scope type
 - zone clusters, [6-10](#)
 - default router, [6-11](#)
 - defrouter property, [6-11](#)
 - dependencies
 - ScalDeviceGroup resource type, [5-18](#)
 - SUNW.ScalDeviceGroup resource type, [5-18](#)
 - SUNW.vucmm_svm resource type, [5-15](#)
 - vucmm_svm resource type, [5-15](#)
 - device groups
 - planning, [1-25](#)

device groups (*continued*)
 replicated disks, [1-26](#)
 resource group affinities, [5-17](#)
 setting the replication property, [4-4](#)

DID devices
 adding to a zone cluster, [6-48](#)

disabling
 installation mode, [3-70](#)
 resources, [7-6](#), [7-7](#)

disk drives, [4-1](#)

disk sets
 adding drives, [4-5](#)
 adding to a zone cluster, [6-49](#)
 configuring, [4-3](#)
 repartitioning drives, [4-6](#)
 reserved names, [1-33](#)

disk strings
 dual-string mediator requirements, [4-10](#)

disk-path monitoring
 default, [3-10](#)

disks, [4-5](#)

disksets, [4-1](#)

domain console network interfaces
 IP addresses, [1-7](#)

drives
 adding to disk sets, [4-5](#)
 mirroring differing device sizes, [1-34](#)
 repartitioning, [4-6](#)

dual-string mediators
 adding hosts, [4-10](#)
 overview, [4-9](#)
 planning, [1-33](#)
 repairing data, [4-11](#)
 status, [4-11](#)

E

enable_priv_net property, [6-9](#)

enabling
 automatic reboot on disk-path failures, [3-10](#)
 LOFS requirements, [1-3](#)
 NTP, [3-73](#)
 Oracle Solaris I/O multipathing, [2-7](#)
 remote login, [6-4](#)

enabling automatic reboot, [3-10](#)

examples
 activating volumes in the md.tab file, [4-9](#)
 adding
 drives to a disk set, [4-6](#)
 local file systems to a zone-cluster node,
[6-44](#)
 local ZFS pools to a zone-cluster node,
[6-45](#)
 mediator hosts, [4-11](#)

examples (*continued*)
 adding (*continued*)
 raw-disk devices to a zone-cluster node,
[6-52](#)

adding devices to a zone cluster
 DID devices, [6-49](#)
 disk sets, [6-50](#)

adding file systems to a zone cluster
 cluster file systems, [6-40](#)
 highly available local file systems, [6-31](#)
 Oracle HSM shared file systems, [6-42](#)
 ZFS cluster file systems, [6-38](#)
 ZFS storage pools, [6-34](#)

configuring Oracle Solaris Cluster software
 on additional nodes by using
 scinstall, [3-58](#)
 on all nodes by using an XML file, [3-18](#)
 on all nodes by using scinstall, [3-11](#)

creating a cluster file system, [5-3](#)

creating a disk set, [4-5](#)

creating state database replicas, [4-2](#)

md.tab file, [4-8](#)

uninstalling
 individual Solaris Cluster packages, [7-10](#)

updating SCSI quorum devices after adding
 a node, [3-66](#)

exclusive IP type
 restriction for zone clusters, [1-23](#)

explorer command, [3-79](#)

F

failover file systems, [1-26](#)

Fair Share Scheduler shares, [6-10](#)

fattach command restriction, [1-27](#)

fencing protocols, [1-17](#)
 NAS devices, [1-10](#)
 SCSI quorum devices, [1-19](#)
 software quorum, [1-19](#)

file system logging
 planning, [1-33](#)

file systems
 adding to a zone cluster, [6-28](#)
 highly available local file systems, [6-28](#)
 ZFS storage pools, [6-32](#)

forcedirectio command
 restriction, [1-32](#)

G

global devices
 file system limitation to UFS, [1-4](#)
 namespace
 updating, [4-3](#)

global devices (*continued*)
 planning, [1-25](#)
 restrictions on lofi device, [1-4](#)
 space requirement for lofi device, [1-4](#)

global fencing, [1-17](#)

global file systems, [5-1](#), [5-3](#)

global storage devices
 adding to a zone cluster, [6-46](#)

Global_zone resource-typeproperty, [1-23](#)

global-cluster nodes
 adding nodes
 by using an XML file, [3-60](#)
 by using `scinstall`, [3-53](#)
 updating quorum devices after, [3-64](#)

creating the global cluster
 by using an XML file, [3-13](#)
 by using Automated Installer, [3-19](#)
 by using `scinstall`, [3-4](#)

IP addresses, [1-7](#)

naming, [1-12](#)

node ID number assignment, [1-12](#)

Oracle VM Server for SPARC domains, [2-3](#)

planning, [1-12](#)

planning for zone clusters, [1-21](#)

verifying
 status, [3-10](#), [3-16](#), [3-30](#), [3-36](#), [3-41](#), [3-57](#)

H

HA for NFS
 coexistence with highly available local file systems, [3-11](#), [3-16](#), [3-29](#), [3-58](#)
 restriction with LOFS, [1-3](#), [1-26](#)

hardware RAID
 internal disk mirroring, [2-13](#)

HASStoragePlus, [1-23](#)
 mounting file systems in zone clusters, [1-23](#)

high-priority processes
 restriction, [1-11](#)

highly available local file systems, [1-23](#)
 adding to a zone cluster, [6-28](#)
 alternative to cluster file systems, [1-26](#)
 coexistence with HA for NFS, [3-11](#), [3-16](#), [3-29](#), [3-58](#)
 LOFS restriction with HA for NFS, [1-3](#)
 quotas support, [1-26](#)

highly available ZFS storage pools
 adding to a zone cluster, [6-32](#)

hostname property, [6-11](#)

hosts file
 configuring, [2-6](#)
 planning, [1-6](#)

hot-spare disks
 planning, [1-33](#)

I

importing
 zones into zone clusters, [6-24](#)

initialization files, [2-24](#)

installation mode
 disabling, [3-70](#)

installing, [2-2](#)
 data services
 by using the `clinstall` command, [2-16](#)
 by using the `pkg` command, [2-20](#)
 Oracle HSM software, [2-24](#)
 Oracle Solaris Cluster software, [2-20](#)
 Oracle Solaris Cluster software
 (`clinstall`), [2-16](#)
 Oracle Solaris software
 alone, [2-4](#)
 with Oracle Solaris Cluster software,
[3-19](#)
 Oracle VM Server for SPARC software, [2-14](#)
`pconsole` software, [2-8](#)
 quorum server software, [2-10](#)
 Solaris Volume Manager software, [4-1](#)
 Trusted Extensions software, [6-3](#)

installing and configuring
 Oracle Solaris and Oracle Solaris Cluster
 from unified archives, [3-32](#)

instantiating
 resource types, [5-18](#)

integrated mirroring, [2-13](#)

internal hardware disk mirroring, [2-13](#)

IP addresses
 adding to a naming service, [2-4](#)
 changing the private IP-address range, [3-49](#)
 guidelines for zone clusters, [1-22](#)
 planning the private network, [1-13](#)
 planning the public network, [1-6](#)

IP network multipathing (IPMP), [1-8](#)

IP strict multihoming
`hostmodel` property, [6-23](#)
 Trusted Extensions, [6-23](#)
 zone clusters, [6-23](#)

IP type
 zone clusters, [1-23](#)

IP-forwarding
 restrictions for IP forwarding, [1-11](#)

`ip-type` property, [6-8](#)

`ipadm` listing private IP addresses, [3-11](#), [3-16](#),
[3-30](#), [3-36](#), [3-57](#), [3-62](#)

IPMP
 configuring groups, [2-7](#)
 planning the public network, [1-8](#)

IPv6 addresses
 private network restriction, [1-14](#), [1-16](#)
 public-network use, [1-7](#)

J
java

- minimum version, [1-2](#)

L
labeled brand zone clusters

- configuring with Trusted Extensions, [6-22](#)
- options, [6-12](#)
- planning, [1-21](#)

LDAP configuration for Trusted Extensions, [6-5](#)**licenses**

- planning, [1-6](#)

lightweight processes

- zone clusters, [6-10](#)

link-based IPMP groups, [1-8](#)**live migration**

- restriction for Oracle VM Server for SPARC software, [1-20](#)

local**file systems**

- adding to a zone cluster, [6-30](#)

- MAC address, [1-8](#)

local file systems

- adding to a specific zone-cluster node, [6-43](#)
- adding to a zone cluster, [6-43](#)

local ZFS storage pools

- adding to a specific zone-cluster node, [6-45](#)

locked memory, [6-10](#)**locked property, [6-10](#)****lofi device**

- space requirement, [1-4](#)
- use restrictions, [1-4](#)

LOFS

- restriction, [1-3](#), [1-26](#)

log files

- Oracle Solaris Cluster installation, [3-10](#)

logging for cluster file systems

- planning, [1-33](#)

logical addresses

- planning, [1-8](#)

logical network interfaces

- restriction, [1-16](#)

loopback file system (LOFS)

- restriction, [1-3](#), [1-26](#)

LWPs

- zone clusters, [6-10](#)

M
manifest

- Automated Installer, [3-27](#), [3-34](#), [3-39](#)

max-lwps property, [6-10](#)

- max-msg-ids property, [6-10](#)

max-sem-ids

- property, [6-10](#)

- max-shm-ids property, [6-10](#)

- max-shm-memory property, [6-10](#)

md.tab file

- configuring, [4-7](#)

mediators, [4-9](#)**message queue IDs**

- zone clusters, [6-10](#)

mirroring

- differing device sizes, [1-34](#)

- internal disks, [2-13](#)

- multihost disks, [1-34](#)

- planning, [1-34](#)

- root pool planning, [1-34](#)

mount options for cluster file systems

- Oracle HSM software, [6-35](#)

- requirements, [5-2](#)

- UFS, [1-30](#), [6-30](#), [6-35](#)

mount points

- adding to new nodes, [2-6](#)

- cluster file systems, [1-31](#)

- modifying the `/etc/vfstab` file, [5-2](#)

- nested, [1-31](#)

MPxIO

- Oracle Solaris I/O multipathing, [2-7](#)

multihost disks

- mirroring, [1-34](#)

- planning, [1-33](#)

multiple-owner volume-manager framework

- resource group

- configuring, [5-14](#)

multiported disks, [1-33](#)**multiuser services**

- verifying, [3-10](#), [3-15](#), [3-57](#)

N
naming conventions

- cluster, [1-12](#)

- global-cluster nodes, [1-12](#)

- private hostnames, [1-15](#)

- raw-disk devices, [5-2](#)

- tagged VLAN adapters, [1-16](#)

- zone clusters, [1-22](#)

naming service

- adding IP-address mappings, [2-4](#)

NAS devices

- configuring as quorum devices, [3-67](#)

- fencing, [1-10](#)

- troubleshooting in a zone cluster installation, [6-6](#)

- NAT and Packet Filter feature of Oracle Solaris software, [1-3](#)

- n_cpus property, [6-10](#)
 - netmasks
 - changing the private netmask, [3-49](#)
 - planning the private network, [1-13](#)
 - Network File System (NFS)
 - guidelines for cluster nodes, [1-9](#)
 - HA for NFS restriction with LOFS, [1-9](#)
 - network interface card (NIC) support for local
 - MAC addresses, [1-8](#)
 - Network Time Protocol (NTP)
 - configuring, [3-73](#)
 - enabling, [3-73](#)
 - guidelines, [1-11](#)
 - restarting, [3-73](#)
 - starting, [3-73](#)
 - network-attached storage, [3-67](#)
 - NFS, [1-9](#)
 - NIC support for local MAC addresses, [1-8](#)
 - NIS servers
 - restriction for cluster nodes, [1-11](#)
 - node IDs, [1-12](#)
 - number assignment, [1-12](#)
 - node lists
 - device groups, [1-33](#)
 - nodes, [1-12](#)
 - noncluster mode
 - booting into, [7-2](#)
 - NTP
 - configuring, [3-73](#)
 - enabling, [3-73](#)
 - guidelines, [1-11](#)
 - restarting, [3-73](#)
 - starting, [3-73](#)
 - ntp.conf file
 - configuring, [3-73](#)
 - ntp.conf.include file
 - creating, [3-74](#), [3-75](#)
 - ntp.conf.sc file
 - configuring, [3-74](#), [3-75](#)
 - num_xip_zoneclusters property, [6-8](#)
- ## O
-
- Oracle ACFS file systems
 - adding to a zone cluster, [6-42](#)
 - configuration diagrams, [5-9](#)
 - creating, [5-8](#), [5-16](#)
 - framework resource groups, [5-13](#)
 - Oracle Explorer software, [3-79](#)
 - Oracle HSM
 - adding shared file systems to a zone cluster, [6-41](#)
 - Oracle HSM software
 - installing the software, [2-24](#)
 - limited use, [1-26](#)
 - Oracle HSM software (*continued*)
 - mount options, [6-35](#)
 - restriction, [1-32](#)
 - Oracle RAC framework resource group
 - configuring for Oracle ACFS, [5-13](#)
 - Oracle Solaris
 - minimum package, [1-2](#)
 - Oracle Solaris Cluster Manager, [6-6](#)
 - adding
 - zfs storage to a zone cluster, [6-32](#)
 - tasks you can perform
 - adding cluster file system to a zone cluster, [6-34](#)
 - adding highly available local file systems to a zone cluster, [6-28](#)
 - adding shared storage to a zone cluster, [6-46](#)
 - Oracle Solaris Cluster software
 - group packages, [2-15](#), [2-20](#), [3-25](#)
 - publisher, [2-6](#), [2-8](#), [2-10](#), [2-23](#), [4-1](#)
 - uninstalling software, [7-1](#)
 - Oracle Solaris I/O multipathing software
 - enabling, [2-7](#)
 - Oracle Solaris software
 - installing
 - alone, [2-4](#)
 - Trusted Extensions feature, [6-3](#)
 - with Oracle Solaris Cluster software, [3-19](#)
 - planning, [1-2](#)
 - /varfile system, [1-4](#)
 - partitions, [1-3](#)
 - root (/) file system, [1-3](#)
 - volume managers, [1-4](#)
 - publisher, [2-6](#), [4-1](#)
 - restrictions
 - automatic power-saving shutdown, [1-3](#)
 - Packet Filter feature, [1-3](#)
 - SMF, [3-10](#), [3-15](#), [3-57](#)
 - Trusted Extensions feature, [1-24](#)
 - Oracle Solaris unified archive
 - installing and configuring Oracle Solaris and Oracle Solaris Cluster software, [3-32](#)
 - Oracle Solaris Unified Archive
 - installing Oracle Solaris and Oracle Solaris Cluster software, [3-19](#)
 - how to, [3-26](#)
 - replicate,restore, [3-19](#)
 - Oracle Solaris zones
 - fencing for NAS devices, [1-10](#)
 - quorum servers, [1-9](#)
 - Oracle VM Server for SPARC software
 - domains as cluster nodes, [2-3](#)

Oracle VM Server for SPARC software (*continued*)

- guest domains
 - Parallel Console Access (`pconsole`)
 - restriction, [2-8](#)
 - planning the address range, [1-14](#)
 - virtual adapter names, [1-16](#)
- guidelines, [1-19](#)
- installing, [2-14](#)
- live migration restriction, [1-20](#)
- mpgroup boot drive support, [1-21](#)

P

package installation

- Oracle Solaris Cluster software, [2-20](#)
- `pconsole` software, [2-8](#)

package installation (`clinstall`)

- Oracle Solaris Cluster software, [2-16](#)

Packet Filter

- configuring, [2-25](#)
- restriction on scalable data services, [2-25](#)
- restrictions, [1-3](#)

Parallel Console Access (`pconsole`) software

- installing, [2-8](#)
- Oracle VM Server for SPARC restriction, [2-8](#)

partitions

- `/var` file system, [1-4](#)
- repartitioning drives, [4-6](#)
- root (`/`) file system, [1-3](#)
- swap, [1-4](#)
- volume manager, [1-4](#)

patches, [1-6](#)

PATH

- cluster nodes
 - `/usr/cluster/bin/` directory, [2-24](#)

`pconsole` utility

- installing the software, [2-8](#)
- using, [2-16](#), [2-21](#), [3-28](#), [3-35](#), [3-40](#)

physical memory

- zone clusters, [6-10](#)

physical property, [6-10](#), [6-11](#)

`pkg` program

- installing software packages, [2-20](#)

PNM object

- automatic group creation during installation, [1-8](#)

private host names

- verifying, [3-73](#)

private hostnames

- changing, [3-72](#)
- planning, [1-15](#)
- zone clusters, [1-23](#)

private IP addresses

- listing with `ipadm`, [3-11](#), [3-16](#), [3-30](#), [3-36](#), [3-57](#), [3-62](#)

private network

- adding zone clusters to the IP-address range, [1-14](#), [3-49](#)
- changing the IP-address range, [3-49](#)
- IPv6 address restriction, [1-16](#)
- Oracle VM Server for SPARC guest domains, [1-14](#)
- planning, [1-13](#)
- use in zone clusters, [1-22](#)

probe-based IPMP groups, [1-8](#)

public network

- adding IP addresses to a naming service, [2-4](#)
- IPv6 support, [1-7](#)
- planning, [1-7](#)

publisher

- Oracle Solaris Cluster software, [2-6](#), [2-8](#), [2-10](#), [2-23](#), [4-1](#)
- Oracle Solaris software, [2-6](#), [4-1](#)

Q

QFS, [2-24](#)

quorum devices, [1-18](#)

- and mirroring, [1-34](#)
- initial configuration, [3-67](#)
- NAS devices, [3-67](#)
- planning, [1-18](#)
- quorum servers, [3-66](#)
- restriction for replicated devices, [1-19](#)
- SATA disks, [1-19](#)
- SCSI protocol setting, [1-19](#)
- software quorum protocol, [1-19](#)
- troubleshooting
 - `clsetup` failure, [3-70](#)
 - `scinstall` failure, [3-70](#)
 - vote count, [3-70](#)
- updating after adding a node, [3-64](#)
- verifying, [3-71](#)
- ZFS restriction, [1-19](#)

quorum servers, [2-10](#)

- configuring as quorum devices, [3-67](#)
- guidelines, [1-9](#)
- installation directory, [2-12](#)
- installing quorum server software, [2-10](#)
- non-global zones, [1-9](#)
- packages installed, [2-12](#)
- requirements as a quorum device, [3-66](#)
- starting, [2-12](#)
- troubleshooting, [2-12](#)

quotas

- restriction on cluster file systems, [1-26](#)

R

rac_framework resource type
 instantiating, [5-14](#)
 registering, [5-14](#)

RAID
 restriction, [1-32](#)

raidctl command, [2-13](#)

rarpd service
 restriction for cluster nodes, [1-11](#)

raw-disk devices
 adding to a zone cluster, [6-51](#)
 naming conventions, [5-2](#)

rebooting
 into noncluster mode, [7-2](#)

recovering
 unsuccessful cluster node creation, [3-12](#),
 [3-19](#), [3-59](#), [3-63](#)

registering
 resource types, [5-17](#)
 ScalDeviceGroup resource type, [5-17](#)
 SUNW.rac_framework resource type, [5-14](#)
 SUNW.ScalDeviceGroup resource type, [5-17](#)
 SUNW.vucmm_framework resource type, [5-15](#)
 SUNW.vucmm_svm resource type, [5-15](#)

remote login
 Trusted Extensions, [6-4](#)

removing, [7-1](#)
 Oracle Solaris Cluster software, [7-1](#)
 partially configured node, [7-1](#)
 zone clusters, [7-6](#)

repairing
 mediator data, [4-11](#)

replicated devices
 disk requirement, [1-26](#)
 restriction as quorum devices, [1-19](#)
 setting the replication property, [4-4](#)

replicating
 a zone cluster
 from unified archives, [3-37](#)
 a zone cluster node
 from unified archives, [3-37](#)

resource groups
 making unmanaged, [7-7](#)
 multiple-owner volume-manager framework,
 [5-14](#)
 scalable device group
 affinities, [5-17](#)

resource types
 ScalDeviceGroup
 dependencies, [5-18](#)
 SUNW.rac_framework, [5-14](#)
 SUNW.ScalDeviceGroup, [5-17](#), [5-18](#)
 dependencies, [5-18](#)

resource types (*continued*)
 SUNW.vucmm_framework, [5-15](#)
 SUNW.vucmm_svm, [5-15](#)
 dependencies, [5-15](#)

resources
 disabling, [7-6](#), [7-7](#)
 listing, [7-6](#)

restarting
 NTP, [3-73](#)

restrictions
 live migration, [1-20](#)
 lofi device, [1-4](#)

root environment
 configuring, [2-24](#)

root pool mirroring
 planning, [1-34](#)

routers
 restriction for cluster nodes, [1-10](#)

RPC service
 disabling TCP wrappers, [3-8](#), [3-13](#), [3-55](#),
 [3-61](#)
 enabling TCP wrappers, [3-57](#), [3-62](#)
 restoring external access, [2-16](#), [2-21](#)
 restricted program numbers, [1-11](#)

S

SATA disks
 configuring as quorum devices, [1-19](#)

scalable data services
 Packet Filter restriction, [2-25](#)

scalable device groups
 resource group affinities, [5-17](#)

ScalDevice
 Group resource type
 registering, [5-17](#)

ScalDeviceGroup
 instantiating, [5-18](#)
 registering, [5-17](#)

ScalDeviceGroup resource type
 dependencies, [5-18](#)
 instantiating, [5-18](#)

scinstall command
 adding nodes, [3-53](#)
 creating the global cluster, [3-4](#)
 creating the global cluster by using
 Automated Installer, [3-19](#)
 unconfiguring Oracle Solaris Cluster
 software, [7-1](#)

scope property, [6-10](#)

SCSI devices
 correcting reservations after adding a third
 node, [3-64](#)
 fencing protocol setting in quorum devices,
 [1-19](#)

- semaphore IDs
 - zone clusters, [6-10](#)
 - shared devices
 - installing quorum devices, [3-67](#)
 - shared IP type
 - zone clusters, [1-23](#)
 - shared memory
 - zone clusters, [6-10](#)
 - shared memory IDs
 - zone clusters, [6-10](#)
 - single-host clusters
 - configuring cluster interconnects, [3-48](#)
 - naming, [1-12](#)
 - single-node clusters, [1-12](#)
 - SMF
 - verifying online services, [3-10](#), [3-15](#), [3-57](#)
 - software quorum protocol, [1-19](#)
 - software RAID
 - restriction, [1-32](#)
 - software updates
 - planning, [1-6](#)
 - zone clusters, [6-23](#)
 - solaris brand zone clusters
 - options, [6-12](#)
 - planning, [1-21](#)
 - Solaris Volume Manager
 - configuring, [4-1](#)
 - disk set reserved names, [1-33](#)
 - disk sets
 - adding drives, [4-5](#)
 - adding to a zone cluster, [6-49](#)
 - configuring, [4-3](#)
 - repartitioning drives, [4-6](#)
 - dual-string mediators
 - adding hosts, [4-10](#)
 - overview, [4-9](#)
 - repairing bad data, [4-11](#)
 - status, [4-11](#)
 - installing the software, [4-1](#)
 - md.tab file, [4-7](#)
 - mediators, [4-10](#)
 - planning, [1-33](#)
 - raw-disk device names, [5-2](#)
 - state database replicas, [4-2](#)
 - unique naming, [1-33](#)
 - volumes
 - activating, [4-8](#)
 - solaris10 brand zone clusters
 - creating, [6-17](#)
 - options, [6-13](#)
 - planning, [1-21](#)
 - supported versions, [6-13](#)
 - spanning tree algorithm
 - disabling for quorum servers, [1-9](#)
 - starting
 - clsetup utility, [3-48](#)
 - NTP, [3-73](#), [3-74](#)
 - pconsole utility, [2-16](#), [2-21](#), [3-28](#), [3-35](#), [3-40](#)
 - quorum servers, [2-12](#)
 - scinstall utility, [3-9](#)
 - state database replicas
 - configuring, [4-2](#)
 - status
 - dual-string mediators, [4-11](#)
 - verifying, [3-71](#)
 - storage devices
 - adding global storage devices to a zone cluster, [6-46](#)
 - adding raw-disk devices to a zone cluster, [6-51](#)
 - adding to zone clusters, [6-46](#)
 - Sun Explorer software, [3-79](#)
 - SUNW.rac_framework resource group
 - configuring for Oracle ACFS, [5-13](#)
 - SUNW.rac_framework resource type
 - instantiating, [5-14](#)
 - registering, [5-14](#)
 - SUNW.ScalDeviceGroup resource type
 - dependencies, [5-18](#)
 - instantiating, [5-18](#)
 - registering, [5-17](#)
 - SUNW.vucmm_framework resource type
 - instantiating, [5-15](#)
 - registering, [5-15](#)
 - SUNW.vucmm_svm resource type
 - dependencies, [5-15](#)
 - instantiating, [5-15](#)
 - registering, [5-15](#)
 - supported configurations, [2-3](#)
 - swap
 - memory for zone clusters, [6-10](#)
 - planning, [1-4](#)
 - swap property, [6-10](#)
 - switches
 - planning, [1-17](#)
 - sysid configuration for zone clusters, [6-15](#)
- ## T
-
- tagged VLAN adapters
 - cluster interconnect guidelines, [1-16](#)
 - public-network guidelines, [1-7](#)
 - TCP wrappers for RPC
 - disabling, [3-8](#), [3-13](#), [3-55](#), [3-61](#)
 - enabling, [3-57](#), [3-62](#)
 - modifying /etc/hosts.allow, [3-11](#), [3-16](#), [3-30](#), [3-36](#)

three-way mirroring, [1-34](#)
 tncfg template, [6-4](#)
 transport adapters, [1-16](#)
 transport switches
 planning, [1-17](#)
 troubleshooting
 Automated Installer installation, [3-32](#)
 configuring
 additional nodes, [3-59](#)
 new global clusters, [3-12](#), [3-19](#), [3-59](#),
 [3-63](#)
 explorer baseline record of the
 configuration, [3-79](#)
 NAS devices in a zone cluster, [6-6](#)
 quorum devices
 clsetup failure, [3-70](#)
 scinstall failure, [3-70](#)
 vote count, [3-70](#)
 quorum server installation, [2-12](#)
 removing partially configured node, [7-3](#)
 Trusted Extensions
 admin_low template, [6-4](#)
 cipso template
 zone cluster, [6-23](#)
 configuring, [6-3](#)
 configuring a zone cluster, [6-22](#)
 guidelines, [1-24](#)
 installation, [6-3](#)
 installing, [6-3](#)
 labeled brand, [1-21](#)
 LDAP, [6-5](#)
 password files, [6-12](#)
 remote login, [6-4](#)
 tncfg template, [6-4](#)
 txzonemgr BUI, [6-12](#)
 zoneshare script, [6-3](#)
 zoneshare Trusted Extension script, [6-3](#)

U

UFS logging
 planning, [1-33](#)
 unconfiguring, [7-1](#)
 Oracle Solaris Cluster software, [7-1](#)
 zone clusters, [7-6](#)
 unified archives
 archiveadm, [3-32](#)
 replicating a zone cluster, [3-37](#)
 uninstalling, [7-1](#)
 Oracle Solaris Cluster software, [7-1](#), [7-4](#)
 unique naming
 Solaris Volume Manager, [1-33](#)
 user-initialization files
 modifying, [2-24](#)

V

validating the cluster, [3-78](#)
 verifying
 automatic reboot on disk-path failures, [3-10](#)
 cldevice command processing, [4-3](#)
 cluster configuration, [3-78](#)
 cluster node status, [3-10](#), [3-16](#), [3-30](#), [3-36](#),
 [3-41](#), [3-57](#)
 Oracle Solaris Cluster software installation,
 [3-71](#)
 private host names, [3-73](#)
 quorum configurations, [3-71](#)
 SMF, [3-10](#), [3-15](#), [3-57](#)
 vfstab configuration, [5-3](#)
 vfstab file
 adding mount points, [5-2](#)
 verifying the configuration, [5-3](#)
 Virtual disk multipathing, [1-21](#)
 VLAN adapters
 cluster interconnect guidelines, [1-16](#)
 private-network guidelines, [1-14](#)
 public-network guidelines, [1-7](#)
 volume managers, [1-32](#)
 partitions for, [1-4](#)
 planning
 general, [1-32](#)
 Solaris Volume Manager, [1-33](#)
 volumes
 Solaris Volume Manager
 activating, [4-8](#)
 vucmm_framework resource type
 instantiating, [5-15](#)
 registering, [5-15](#)
 vucmm_svm resource type
 dependencies, [5-15](#)
 instantiating, [5-15](#)
 registering, [5-15](#)

Z

ZFS root pools
 internal disk mirroring, [2-13](#)
 planning mirroring, [1-34](#)
 ZFS storage pools
 adding a local pool to a zone-cluster node,
 [6-45](#)
 adding to a zone cluster, [6-32](#), [6-33](#)
 restrictions, [1-19](#)
 zone clusters
 adding an Oracle ACFS file system, [6-42](#)
 adding devices
 DID devices, [6-48](#)
 disk sets, [6-49](#)
 global storage devices, [6-46](#)

- zone clusters (*continued*)
 - adding devices (*continued*)
 - raw-disk devices, [6-51](#)
 - adding file systems, [6-28](#)
 - cluster file systems, [6-34](#), [6-39](#)
 - highly available local file systems, [6-28](#)
 - highly available ZFS storage pools, [6-32](#)
 - local file systems, [6-30](#)
 - local file systems for a specific zone-cluster node, [6-43](#)
 - local ZFS storage pools, [6-45](#)
 - Oracle HSM, [6-41](#)
 - ZFS storage pools, [6-32](#), [6-33](#)
 - adding local file systems, [6-43](#)
 - adding storage devices, [6-46](#)
 - adding to the private IP-address range, [1-14](#)
 - address property, [6-11](#)
 - allowed-address property, [6-11](#)
 - brand property, [6-8](#)
 - brands, [1-21](#)
 - capped-cpu, [6-10](#)
 - configure-allowed-address property, [6-11](#)
 - configuring, [6-2](#)
 - configuring to use Trusted Extensions, [6-22](#)
 - creating, [6-6](#)
 - dedicated-cpu, [6-10](#)
 - default IP type, [1-23](#)
 - default router, [6-11](#)
 - defrouter property, [6-11](#)
 - enable_priv_net property, [6-9](#)
 - Fair Share Scheduler shares, [6-10](#)
 - halting, [7-7](#)
 - hostmodel property, [6-23](#)
 - hostname property, [6-11](#)
 - importing zones, [6-24](#)
 - IP addresses, [1-7](#)
- zone clusters (*continued*)
 - ip-type property, [6-8](#)
 - labeled brand options, [6-12](#)
 - lightweight processes, [6-10](#)
 - locked memory, [6-10](#)
 - locked property, [6-10](#)
 - message queue IDs, [6-10](#)
 - mounting file systems with HAStoragePlus, [1-23](#)
 - naming, [1-22](#)
 - ncpus property, [6-10](#)
 - num_xip_zoneclusters property, [6-8](#)
 - Oracle ACFS file systems, [5-16](#)
 - physical memory, [6-10](#)
 - physical property, [6-10](#), [6-11](#)
 - planning, [1-21](#)
 - scope property, [6-10](#)
 - semaphore IDs, [6-10](#)
 - setting the Global_zone resource-type property, [1-23](#)
 - shared memory, [6-10](#)
 - shared memory IDs, [6-10](#)
 - software updates, [6-23](#)
 - solaris brand options, [6-12](#)
 - solaris10 brand
 - supported versions, [6-13](#)
 - solaris10 brand options, [6-13](#)
 - swap memory, [6-10](#)
 - swap property, [6-10](#)
 - sysid configuration, [6-15](#)
 - unconfiguring, [7-6](#)
 - zonepath property, [6-7](#)
- zonepath property, [6-7](#)
- zones
 - importing into zone clusters, [6-24](#)