

Creating and Administering Zone Clusters With Oracle® Solaris Cluster 4.4

ORACLE®

Part No: E80122
August 2018

Part No: E80122

Copyright © 2018, Oracle and/or its affiliates. All rights reserved.

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

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

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

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

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

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

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

This software or hardware and documentation may provide access to or information 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.

Access to Oracle Support

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

Référence: E80122

Copyright © 2018, Oracle et/ou ses affiliés. Tous droits réservés.

Ce logiciel et la documentation qui l'accompagne sont protégés par les lois sur la propriété intellectuelle. Ils sont concédés sous licence et soumis à des restrictions d'utilisation et de divulgation. Sauf 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, ou la documentation qui l'accompagne, est livré sous licence au Gouvernement des Etats-Unis, ou à quiconque qui aurait souscrit la licence de ce logiciel pour le compte du Gouvernement des Etats-Unis, la notice suivante s'applique :

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

Ce logiciel ou matériel a été développé pour un usage général dans le cadre d'applications de gestion des informations. Ce logiciel ou matériel n'est pas conçu ni n'est destiné à être utilisé dans des applications à risque, notamment dans des applications pouvant causer un risque de dommages corporels. Si vous utilisez ce logiciel ou ce 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 et Java sont des marques déposées d'Oracle Corporation et/ou de ses affiliés. Tout autre nom mentionné peut correspondre à des marques appartenant à d'autres propriétaires qu'Oracle.

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

Accès aux services de support Oracle

Les clients Oracle qui ont souscrit un contrat de support ont accès au support électronique via My Oracle Support. Pour plus d'informations, visitez le site <http://www.oracle.com/pls/topic/lookup?ctx=acc&id=info> ou le site <http://www.oracle.com/pls/topic/lookup?ctx=acc&id=trs> si vous êtes malentendant.

Contents

- Using This Documentation** 7

- 1 Creating Zone Clusters** 9
 - Overview of Creating and Configuring a Zone Cluster 9
 - Creating and Configuring a Zone Cluster 10
 - Creating a Zone Cluster 10
 - Importing Zones into Zone Clusters 24
 - Adding File Systems to a Zone Cluster 29
 - Adding Local File Systems to a Specific Zone-Cluster Node 45
 - Adding Storage Devices to a Zone Cluster 48

- 2 Administering the Zone Cluster** 57
 - Overview of Zone Cluster Administrative Tasks 57
 - ▼ How to Configure a Zone Cluster from the Unified Archive 59
 - ▼ How to Install a Zone Cluster from the Unified Archive 60
 - ▼ How to Add a Network Address to a Zone Cluster 61
 - Adding a Node to a Cluster or Zone Cluster 62
 - ▼ How to Add a Node to an Existing Cluster or Zone Cluster 64
 - ▼ How to Remove a File System From a Zone Cluster 65
 - ▼ How to Remove a Storage Device From a Zone Cluster 68

- Index** 71

Using This Documentation

- **Overview** – Describes how to configure and administer zone clusters on an Oracle Solaris Cluster configuration.
- **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://www.oracle.com/pls/topic/lookup?ctx=E69294_01.

Feedback

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

Creating Zone Clusters

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

- [“Overview of Creating and Configuring a Zone Cluster” on page 9](#)
- [“Creating and Configuring a Zone Cluster” on page 10](#)

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 1 Task Map: Create a Zone Cluster

Task	Instructions
Plan the zone cluster configuration.	“Zone Clusters” in <i>Installing and Configuring an Oracle Solaris Cluster 4.4 Environment</i>
If using Trusted Extensions, install and configure the Trusted Extensions software.	“How to Install and Configure Trusted Extensions” on page 10
Create the zone cluster.	“How to Create a Zone Cluster (clsetup)” on page 13
<i>(Optional)</i> Add a file system for use by multiple nodes of the zone cluster.	“Adding File Systems to a Zone Cluster” on page 29
<i>(Optional)</i> Add a file system dedicated to a single node of the zone cluster.	“Adding Local File Systems to a Specific Zone-Cluster Node” on page 45
<i>(Optional)</i> Add a storage device for use by multiple nodes or a single node of the zone cluster.	“Adding Storage Devices to a Zone Cluster” on page 48
<i>(Optional)</i> Validate the zone cluster.	“How to Validate a Basic Cluster Configuration” in <i>Administering an Oracle Solaris Cluster 4.4 Configuration</i>

Creating and Configuring a Zone Cluster

This section provides the following information and procedures to create and configure a zone cluster.

- [“Creating a Zone Cluster” on page 10](#)
- [“Importing Zones into Zone Clusters” on page 24](#)
- [“Adding File Systems to a Zone Cluster” on page 29](#)
- [“Adding Local File Systems to a Specific Zone-Cluster Node” on page 45](#)
- [“Adding Storage Devices to a Zone Cluster” on page 48](#)

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.

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

This section contains the following procedures:

- [“How to Install and Configure Trusted Extensions” on page 10](#)
- [“How to Create a Zone Cluster \(`clsetup`\)” on page 13](#)
- [“How to Configure a Zone Cluster to Use Trusted Extensions” on page 22](#)

▼ How to Install and Configure Trusted Extensions

This procedure prepares the global cluster to use the Trusted Extensions feature of Oracle Solaris with zone clusters. If you do not plan to enable Trusted Extensions, proceed to [“Creating a Zone Cluster” on page 10](#).

Perform this procedure on each node in the global cluster.

Before You Begin Perform the following tasks:

- Ensure that the Oracle Solaris OS is installed to support Oracle Solaris Cluster and Trusted Extensions software. See [“How to Install Oracle Solaris Software”](#) in *Installing and Configuring an Oracle Solaris Cluster 4.4 Environment* 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”](#) in *Installing and Configuring an Oracle Solaris Cluster 4.4 Environment*.

1. Assume the root role on a node of the global cluster.

2. Install and configure Trusted Extensions software.

Follow procedures in [Chapter 3, “Adding the Trusted Extensions Feature to Oracle Solaris”](#) in *Trusted Extensions Configuration and Administration*.

3. Disable the Trusted Extensions zoneshare and zoneunshare scripts.

The Trusted Extensions zoneshare and zoneunshare scripts support the ability to export home directories on the system. An Oracle Solaris Cluster configuration does not support this feature. Disable this feature by replacing each script with a symbolic link to the `/bin/true` utility.

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

4. Configure all logical-hostname that are to be used in the zone cluster.

See [“How to Create a Default Trusted Extensions System”](#) in *Trusted Extensions Configuration and Administration*.

5. (Optional) Enable remote login by the LDAP server to the global-cluster node.

a. In the `/etc/default/login` file, comment out the `CONSOLE` entry.

b. Enable remote login.

```
phys-schost# svcadm enable rlogin
```

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

Modify the account management entries by appending a Tab and typing `allow_remote` or `allow_unlabeled` respectively, as shown below.

```
other account requisite pam_roles.so.1 Tab allow_remote
other account required pam_unix_account.so.1 Tab allow_unlabeled
```

6. Modify the `admin_low` template.

- a. Assign the `admin_low` template to each IP address that does *not* belong to a Trusted Extensions machine that is used by the global zone.**

```
# tncfg -t admin_low
tncfg:admin_low> add host=ip-address1
tncfg:admin_low> add host=ip-address2
...
tncfg:admin_low> exit
```

- b. Remove the wildcard address `0.0.0.0/32` from the `tncfg` template.**

```
# tncfg -t admin_low remove host=0.0.0.0
```

7. Assign the `cipso` template to each IP address that *does* belong to a Trusted Extensions machine that is used by the global zone.

```
# tncfg -t cipso
tncfg:cipso> add host=ip-address1
tncfg:cipso> add host=ip-address2
...
tncfg:cipso> exit
```

8. Repeat [Step 1](#) through [Step 7](#) on each remaining node of the global cluster.

When all steps are completed on all global-cluster nodes, perform the remaining steps of this procedure on each node of the global cluster.

9. On each global-cluster node, add the IP address of the Trusted Extensions-enabled LDAP server to the `/etc/inet/hosts` file.

The LDAP server is used by the global zone and by the nodes of the zone cluster.

10. (Optional) Make the global-cluster node an LDAP client.

11. Add Trusted Extensions users.

See [“Creating Roles and Users in Trusted Extensions”](#) in *Trusted Extensions Configuration and Administration*.

Next Steps Create the zone cluster. Go to [“Creating a Zone Cluster” on page 10](#).

▼ How to Create a Zone Cluster (clsetup)

Perform this procedure to create a zone cluster using the `clsetup` utility.

To modify the zone cluster after it is installed, see [“Performing Zone Cluster Administrative Tasks” in *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.

- Before You Begin**
- Create a global cluster. See [Chapter 3, “Establishing the Global Cluster” in *Installing and Configuring an Oracle Solaris Cluster 4.4 Environment*](#).
 - Read the guidelines and requirements for creating a zone cluster. See [“Zone Clusters” in *Installing and Configuring an Oracle Solaris Cluster 4.4 Environment*](#).
 - 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” on page 10](#).
 - 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 `zonpath` 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 Geographic Edition configuration. Otherwise, this requirement is optional. For more information about this Geographic Edition requirement, see [“Disaster Recovery Framework” in *Installing and Configuring an Oracle Solaris Cluster 4.4 Environment*](#).
- If applicable, the name of the public network management object that each zone-cluster node uses to connect to the public network.

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

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

Tip - While in the `clsetup` utility, you can press the < key to return to a previous screen.

You can also use Oracle Solaris Cluster Manager to create a zone cluster. For the browser interface login instructions, see [“How to Access Oracle Solaris Cluster Manager” in *Administering an Oracle Solaris Cluster 4.4 Configuration*](#).

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

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

2. Ensure that the node of the global cluster is in cluster mode.

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

--- Node Status ---

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

3. Start the `clsetup` utility.

```
phys-schost# clsetup
```

The Main Menu is displayed.

4. **Choose the Zone Cluster menu item.**
5. **Choose the Create a Zone Cluster menu item.**
6. **Type the name of the zone cluster you want to add.**
A zone cluster name can contain ASCII letters (a-z and A-Z), numbers, a dash, or an underscore. The maximum length of the name is 20 characters.
7. **Choose the property you want to change.**
You can set the following properties:

Property	Description
<code>zonepath=zone-cluster-node-path</code>	Specifies the path to the zone cluster node. For example, <code>/zones/sczone</code> .
<code>brand=brand-type</code>	Specifies the <code>solaris</code> or <code>labeled</code> zones brand used in the zone cluster. Note - To use Trusted Extensions, you must use only the <code>labeled</code> brand.
<code>ip-type=value</code>	Specifies the type of network IP address used by the zone cluster. 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 <code>cluster(8CL)</code> man page.
<code>enable_priv_net=value</code>	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> . 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: <code>ip-type=exclusive</code> (<code>solaris</code> brand only) – Communication between zone cluster nodes uses the specified <code>privnet</code> resources. The <code>privnet</code> resources are either Virtual Network Interfaces (VNICs) for the Ethernet type of private network adapters, or InfiniBand (IB) partitions for the IB type of private network adapters. The VNICs or IB partitions are automatically created by the wizard over each private network adapter of the global cluster, and used to configure a zone cluster. The VNICs or IB partitions that the wizard generates use the following naming conventions: For the Ethernet type: <code>private-network-interface-name_zone-cluster-name_vnic0</code> .

Property	Description
	<p>For the IB type: <i>private-network-interface-name_zone-cluster-name_ibp0</i>.</p> <p>For example, the private network interfaces of the global cluster are net2 and net3, and the zone cluster name is zone1. If net2 and net3 are Ethernet type network interfaces, the two VNICs that are created for the zone cluster will have the names net2_zone1_vnic0 and net3_zone1_vnic0.</p> <p>If net2 and net3 are IB type network interfaces, the two IB partitions created for the zone cluster will have the names net2_zone1_ibp0 and net3_zone1_ibp0.</p>

8. (Optional) Choose the Zone System Resource Control property that you want to change.

You can set the following properties:

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

9. (Optional) Choose the Zone CPU Resource Control property that you want to change.

You can set the following properties:

Property	Description
<i>scope=scope-type</i>	Specifies whether the ncpus property used in a zone cluster is dedicated-cpu or capped-cpu.
<i>ncpus=value</i>	<p>Specifies the limit for the scope type.</p> <ul style="list-style-type: none"> ■ If the scope property is set to <i>dedicated-cpu</i>, the ncpus property sets a limit on the number of CPUs that should be assigned for this zone's exclusive use. The zone will create a pool and processor set

Property	Description
	<p>when it boots. See the pooladm(8) and poolcfg(8) man pages for more information on resource pools.</p> <ul style="list-style-type: none"> ■ If the scope 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), pooladm(8), and poolcfg(8) man pages for more information on resource pools.

10. (Optional) Choose the capped-memory property that you want to change.

You can set the following properties:

Property	Description
<code>physical=value</code>	Specifies the GByte limit for physical memory.
<code>swap=value</code>	Specifies the GByte limit for swap memory.
<code>locked=value</code>	Specifies the GByte limit for locked memory.

You can also use 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 login instructions, see “[How to Access Oracle Solaris Cluster Manager](#)” in *Administering an Oracle Solaris Cluster 4.4 Configuration*.

11. 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>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.

12. Specify the network addresses for the zone cluster.

The network addresses can be used to configure a logical hostname cluster resources in the zone cluster. The network address is in the zone cluster global scope.

13. 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=exclusive
set enable_priv_net=true
add capped-memory
set physical=2G
end
add node
set physical-host=phys-schost-1
set hostname=zc-host-1
add net
set address=172.1.1.1
set physical=net0
end
add net
set address=172.1.1.2
end
```

Zone cluster, zc2 has been created and configured successfully.

Continue to install the zone cluster(yes/no) ?

14. Type yes to continue.

The clsetup utility performs a standard configuration of a zone cluster and you cannot specify any options.

15. When finished, exit the clsetup utility.

16. 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
```

17. 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.

18. 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
-c <i>config-profile.xml</i>	Includes system configuration information. The -c <i>config-profile.xml</i> option provides a configuration profile for all non-global zones of the zone cluster. Using this option changes only the hostname of the zone, which is unique for each zone in the zone cluster. All profiles must have a .xml extension. The contents of the file is a line-delimited list of the commands to be specified to the interactive <code>clzonecluster</code> utility. See the Example section of the clzonecluster(8CL) man page for an example of the profile contents.
-M <i>manifest.xml</i>	Specifies a custom Automated Installer manifest that you configure to install the necessary packages on all zone-cluster nodes. Use this option if the base global-cluster nodes for the zone-cluster are not all installed with the same Oracle Solaris Cluster packages but you do not want to change which packages are on the base nodes. If the <code>clzonecluster install</code> command is run without the -M option, zone-cluster installation fails on a base node if it is missing a package that is installed on the issuing base node.

For more information, see the [clzonecluster\(8CL\)](#) man page.

- 19. 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 20](#).

Note - In the following steps, the non-global zone `znode` and `zone-cluster-name` share the same name.

- **For an exclusive-IP labeled brand zone cluster, perform the following steps.**

Configure only *one* zone-cluster node at a time.

- a. **Boot the non-global zone of one zone-cluster node.**

```
phys-schost# zoneadm -z znode boot
```

- b. **Unconfigure the Oracle Solaris instance and reboot the zone.**

```
phys-schost# zlogin znode
znode# sysconfig unconfigure
znode# reboot
```

The `zlogin` session terminates during the reboot.

- c. **Issue the `zlogin` command and progress through the interactive screens.**

```
phys-schost# zlogin -C znode
```

- d. **When finished, exit the zone console.**

For information about methods to exit from a non-global zone, see [“Exiting 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 znode halt
```

- f. **Repeat the preceding steps for each remaining zone-cluster node.**

- For a solaris 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 [“Exiting a Non-Global Zone”](#) in *Creating and Using Oracle Solaris Zones*.

- d. Repeat [Step 192b](#) through [Step 192c](#) for each remaining zone-cluster node.

20. Boot the zone cluster.

Installation of the zone cluster might take several minutes.

```
phys-schost# clzonecluster boot zone-cluster-name
```

21. 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 complete Trusted Extensions configuration, go to [“How to Configure a Zone Cluster to Use Trusted Extensions”](#) on page 22.

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

- [“Adding File Systems to a Zone Cluster”](#) on page 29
- [“Adding Local File Systems to a Specific Zone-Cluster Node”](#) on page 45
- [“Adding Storage Devices to a Zone Cluster”](#) on page 48

▼ How to Configure a Zone Cluster to Use Trusted Extensions

After you create a labeled brand zone cluster, perform the following steps to finish configuration to use Trusted Extensions.

1. Complete IP-address mappings for the zone cluster.

Perform this step on each node of the zone cluster.

a. From a node of the global cluster, display the node's ID.

```
phys-schost# cat /etc/cluster/nodeid  
N
```

b. Log in to a zone-cluster node on the same global-cluster node.

Ensure that the SMF service has been imported and all services are up before you log in.

c. Determine the IP addresses used by this zone-cluster node for the private interconnect.

The cluster software automatically assigns these IP addresses when the cluster software configures a zone cluster.

In the `ifconfig -a` output, locate the `clprivnet0` logical interface that belongs to the zone cluster. The value for `inet` is the IP address that was assigned to support the use of the cluster private interconnect by this zone cluster.

```
zc1# ifconfig -a  
lo0:3: flags=20010008c9<UP,LOOPBACK,RUNNING,NOARP,MULTICAST,IPv4,VIRTUAL> mtu 8232  
index 1
```

```

zone zc1
inet 127.0.0.1 netmask ff000000
net0: flags=1000843<UP,BROADCAST,RUNNING,MULTICAST,IPv4> mtu 1500 index 2
inet 10.11.166.105 netmask fffffff0 broadcast 10.11.166.255
groupname sc_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 fffffff0 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 zc1
inet 10.11.166.160 netmask fffffff0 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 ffffffff8 broadcast 172.16.0.23
ether 0:0:0:0:0:2
clprivnet0:3: flags=1009843<UP,BROADCAST,RUNNING,MULTICAST,MULTI_BCAST,PRIVATE,IPv4>
  mtu 1500 index 7
zone zc1
inet 172.16.0.22 netmask fffffffc broadcast 172.16.0.23

```

d. Add to the zone-cluster node's `/etc/inet/hosts` file the following addresses of the zone-cluster node.

- The hostname for the private interconnect, which is `clusternodeN-priv`, where *N* is the global-cluster node ID

```
172.16.0.22    clusternodeN-priv
```

- Each net resource that was specified to the `clzonecluster` command when you created the zone cluster

e. Repeat on the remaining zone-cluster nodes.

2. Authorize communication with zone-cluster components.

Create new entries for the IP addresses used by zone-cluster components and assign each entry a CIPSO template. These IP addresses which exist in the zone-cluster node's `/etc/inet/hosts` file are as follows:

- Each zone-cluster node private IP address
- All `cl_privnet` IP addresses in the zone cluster
- Each logical-hostname public IP address for the zone cluster
- Each shared-address public IP address for the zone cluster

```
phys-schost# tncfg -t cipso
tncfg:cipso> add host=ipaddress1
tncfg:cipso> add host=ipaddress2
...
tncfg:cipso> exit
```

For more information about CIPSO templates, see [“How to Configure a Different Domain of Interpretation”](#) in *Trusted Extensions Configuration and Administration*.

3. Set IP strict multihoming to weak.

Perform the following commands on each node of the zone cluster.

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

For more information about the `hostmodel` property, see [“hostmodel \(IPv4 or IPv6\)”](#) in *Oracle Solaris 11.4 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”](#) on page 29
- [“Adding Local File Systems to a Specific Zone-Cluster Node”](#) on page 45
- [“Adding Storage Devices to a Zone Cluster”](#) on page 48

See Also If you want to update the software on a zone cluster, follow procedures in [“Updating to a New Oracle Solaris Cluster Version”](#) in *Updating Your Oracle Solaris Cluster 4.4 Environment*. These procedures include special instructions for zone clusters, where needed.

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.
- 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” on page 25](#).

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 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” on page 25](#).
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/zonec
clzc:new-zoneclustername:> set ip-type=exclusive
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 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/zonec
clzc:sczone-new:> set ip-type=exclusive
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”](#) on page 45.

This section provides the following procedures to add file systems for use by the zone cluster:

- [“How to Add a Highly Available Local File System to a Zone Cluster \(clsetup\)”](#) on page 29
- [“How to Add a Highly Available Local File System to a Zone Cluster \(CLI\)”](#) on page 31
- [“How to Add a ZFS Storage Pool to a Zone Cluster \(clsetup\)”](#) on page 33
- [“How to Add a ZFS Storage Pool to a Zone Cluster \(CLI\)”](#) on page 36
- [“How to Add a StorageTek QFS Shared File System to a Zone Cluster \(CLI\)”](#) on page 42
- [“How to Add a Cluster File System to a Zone Cluster \(clsetup\)”](#) on page 37
- [“How to Add a UFS Cluster File System to a Zone Cluster \(CLI\)”](#) on page 40
- [“How to Add an Oracle ACFS File System to a Zone Cluster \(CLI\)”](#) on page 44

You can also use Oracle Solaris Cluster Manager to add a file system to a zone cluster. For the browser interface login 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 \(CLI\)”](#) on page 31.

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 login 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.**
The Zone Cluster Tasks Menu is displayed.
5. **Choose the Add File System/Storage Device to a Zone Cluster menu item.**
The Select Zone Cluster menu is displayed.
6. **Choose the zone cluster where you want to add the file system.**
The Storage Type Selection menu is displayed.
7. **Choose the File System menu item.**
The File System Selection for the Zone Cluster menu is displayed.
8. **Choose the file system you want to add to the zone cluster.**

The file systems in the list are those that are configured on the shared disks and can be accessed by the nodes where the zone cluster is configured. You can also type **e** to manually specify all properties for a file system.

The Mount Type Selection menu is displayed.

9. Choose the Loopback mount type.

The File System Properties for the Zone Cluster menu is displayed.

10. Change the properties that you are allowed to change for the file system you are adding.

Note - For UFS file systems, enable logging.

When finished, type **d** and press Return.

11. Type c to save the configuration change.

The results of your configuration change are displayed.

12. When finished, exit the `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 \(clsetup\)”](#) on page 29.

To add a ZFS pool to a zone cluster, instead perform procedures in [“How to Add a ZFS Storage Pool to a Zone Cluster \(clsetup\)”](#) on page 33. 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 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/rdisk/d1
clzc:sczone:fs> set type=ufs
clzc:sczone:fs> add options [logging]
clzc:sczone:fs> end
clzc:sczone> verify
clzc:sczone> commit
clzc:sczone> exit

phys-schost-1# clzonecluster show -v sczone
...
Resource Name:                fs
dir:                          /global/oracle/d1
special:                       /dev/md/oracle/dsk/d1
raw:                           /dev/md/oracle/rdisk/d1
type:                          ufs
options:                       [logging]
cluster-control:              [true]
...
```

Next Steps Configure the file system to be highly available by using an `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)

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 `HASStoragePlus` 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 \(CLI\)” on page 36](#).

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 login instructions, see [“How to Access Oracle Solaris Cluster Manager” in *Administering an Oracle Solaris Cluster 4.4 Configuration*](#).

Before You Begin Ensure that the ZFS pool is connected on shared disks that are connected to all nodes of the zone cluster. See [Managing ZFS File Systems in Oracle Solaris 11.4](#) for procedures to create a ZFS pool.

1. **Assume the root role on a node of the global cluster that hosts the zone cluster.**
You perform all steps of this procedure from a node of the global cluster.

2. **Start the `clsetup` utility.**

```
phys-schost# clsetup
```

The Main Menu is displayed.

Tip - To return to a previous screen, type the < key and press Return.

3. **Choose the Zone Cluster menu item.**
The Zone Cluster Tasks Menu is displayed.
4. **Choose the Add File System/Storage Device to a Zone Cluster menu item.**
The Select Zone Cluster menu is displayed.
5. **Choose the zone cluster where you want to add the ZFS storage pool.**
The Storage Type Selection menu is displayed.

6. Choose the ZFS menu item.

The ZFS Pool Selection for the Zone Cluster menu is displayed.

7. Choose the ZFS pool you want to add to the zone cluster.

The ZFS pools in the list are those that are configured on the shared disks and can be accessed by the nodes where the zone cluster is configured. You can also type **e** to manually specify properties for a ZFS pool.

The ZFS Pool Dataset Property for the Zone Cluster menu is displayed. The selected ZFS pool is assigned to the name property.

8. Type **d and press Return.**

The Review File Systems/Storage Devices for the Zone Cluster menu is displayed.

9. Type **c to save the configuration change.**

The results of your configuration change are displayed. For example:

```
>>> Result of Configuration Change to the Zone Cluster(sczone) <<<
```

```
Adding file systems or storage devices to sczone zone cluster...
```

```
The zone cluster is being created with the following configuration
```

```
/usr/cluster/bin/clzonecluster configure sczone
add dataset
set name=myzpool5
end
```

```
Configuration change to sczone zone cluster succeeded.
```

10. When finished, exit the clsetup utility.**11. Verify the addition of the file system.**

```
phys-schost# clzonecluster show -v zoneclustername
```

12. To make the ZFS storage pool highly available, configure the pool with an 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 \(clsetup\)”](#) on page 33.

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 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 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 Cluster File System to a Zone Cluster (clsetup)

The `clsetup` utility discovers and displays the available file systems that are configured on the cluster nodes where the selected zone cluster is configured. When you use the `clsetup` utility to add a file system, the file system is added in cluster scope.

You can add the following types of cluster file systems to a zone cluster:

- UFS cluster file system - You specify the file system type in the `/etc/vfstab` file, using the `global` mount option. This file system can be located on the shared disk or on a Solaris Volume Manager device.
- StorageTek QFS shared file system - You specify the file system type in the `/etc/vfstab` file, using the `shared` mount option.
- ACFS - Discovered automatically, based on the `ORACLE_HOME` path you provide.

Note - Alternatively, you can use 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 \(CLI\)” on page 40](#)
- [“How to Add a StorageTek QFS Shared File System to a Zone Cluster \(CLI\)” on page 42](#)
- [“How to Add an Oracle ACFS File System to a Zone Cluster \(CLI\)” on page 44](#)

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 login instructions, see [“How to Access Oracle Solaris Cluster Manager” in *Administering an Oracle Solaris Cluster 4.4 Configuration*](#).

Before You Begin Ensure that the cluster file system you want to add to the zone cluster is configured. See [“Planning Cluster File Systems” in *Installing and Configuring an Oracle Solaris Cluster 4.4 Environment*](#) and Chapter 5, [“Creating a Cluster File System” in *Installing and Configuring an Oracle Solaris Cluster 4.4 Environment*](#).

1. **Assume the root role on a node of the global cluster that hosts the zone cluster.**
You perform all steps of this procedure from a node of the global cluster.
2. **On each node of the global cluster that hosts a zone-cluster node, add an entry to the `/etc/vfstab` file for the file system that you want to mount on the zone cluster.**

```
phys-schost# vi /etc/vfstab
```

- **For a UFS entry, include the `global` mount option, similar to the following example:**

```
/dev/md/datadg/dsk/d0 /dev/md/datadg/rdisk/d0 /global/fs ufs 2 no global, logging
```

- **For a shared QFS entry, include the `shared` mount option, similar to the following example:**

```
Data-cz1 - /db_qfs/Data1 samfs - no shared,notrace
```

3. **On the global cluster, start the `clsetup` utility.**

```
phys-schost# clsetup
```

The Main Menu is displayed.

Tip - To return to a previous screen, type the < key and press Return.

4. Choose the Zone Cluster menu item.

The Zone Cluster Tasks Menu is displayed.

5. Choose the Add File System/Storage Device to a Zone Cluster menu item.

The Select Zone Cluster menu is displayed.

6. Choose the zone cluster where you want to add the file system.

The Storage Type Selection menu is displayed.

7. Choose the File System menu item.

The File System Selection for the Zone Cluster menu is displayed.

8. Choose a file system from the list.

You can also type **e** to manually specify all properties for a file system. If you are using an ACFS file system, you can select **Discover ACFS** and then specify the **ORACLE_HOME** directory.

The Mount Type Selection menu is displayed.

9. Choose the Loopback file system mount type for the zone cluster.

If you chose ACFS in [Step 7](#), the `clsetup` utility skips this step because ACFS supports only the direct mount type.

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 UFS Cluster File System to a Zone Cluster (CLI)

Perform this procedure to add a UFS cluster file system for use by a zone cluster.

Note - Alternatively, you can use the `clsetup` utility to perform this task. See [“How to Add a Cluster File System to a Zone Cluster \(`clsetup`\)” on page 37](#).

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.

5. Verify the addition of the LOFS file system.

```
phys-schost# clzonecluster show -v zone-cluster-name
```

Example 5 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 StorageTek QFS Shared File System to a Zone Cluster (CLI)

Perform this task to add a StorageTek QFS shared file system for use by a zone cluster.

Note - Alternatively, you can use the `clsetup` utility to perform this task. See [“How to Add a Cluster File System to a Zone Cluster \(clsetup\)” on page 37](#).

At this time, StorageTek QFS shared file systems are only supported for use in clusters that are configured with Oracle RAC. On clusters that are not configured with Oracle RAC, you can use a single-machine StorageTek QFS file system that is configured as a highly available local file system.

- 1. Assume the root role on a node of the global cluster that hosts the zone cluster.**
You perform all steps of this procedure from a node of the global cluster.
- 2. On the global cluster, configure the StorageTek QFS shared file system that you want to use in the zone cluster.**

Follow procedures for shared file systems in your StorageTek QFS documentation.

3. **On each node of the global cluster that hosts a zone-cluster node, add an entry to the `/etc/vfstab` file for the file system that you want to mount on the zone cluster.**
4. **Add the file system to the zone cluster configuration.**

```
phys-schost# clzonecluster configure zone-cluster-name
clzc:zone-cluster-name> add fs
clzc:zone-cluster-name: fs> set dir=mount-point
clzc:zone-cluster-name: fs> set special=QFS-file-system-name
clzc:zone-cluster-name: fs> set type=samfs
clzc:zone-cluster-name: fs> end
clzc:zone-cluster-name> verify
clzc:zone-cluster-name> commit
clzc:zone-cluster-name> exit
```

5. **Verify the addition of the file system.**

```
phys-schost# clzonecluster show -v zone-cluster-name
```

Example 6 Adding a StorageTek QFS Shared File System as a Direct Mount to a Zone Cluster (CLI)

The following example shows the StorageTek QFS shared file system `Data-cz1` added to the zone cluster `sczone`. From the global cluster, the mount point of the file system is `/zones/sczone/root/db_qfs/Data1`, where `/zones/sczone/root/` is the zone's root path. From the zone-cluster node, the mount point of the file system is `/db_qfs/Data1`.

```
phys-schost-1# vi /etc/vfstab
#device    device    mount    FS      fsck    mount    mount
#to mount  to fsck   point    type    pass    at boot  options
#
Data-cz1  -    /zones/sczone/root/db_qfs/Data1 samfs - no shared,notrace
```

```
phys-schost-1# clzonecluster configure sczone
clzc:sczone> add fs
clzc:sczone: fs> set dir=/db_qfs/Data1
clzc:sczone: fs> set special=Data-cz1
clzc:sczone: fs> set type=samfs
clzc:sczone: fs> end
clzc:sczone> verify
clzc:sczone> commit
clzc:sczone> exit
```

```
phys-schost-1# clzonecluster show -v sczone
...
Resource Name:                                fs
```

```
dir: /db_qfs/Data1
special: Data-cz1
raw:
type: samfs
options: []
...
```

▼ How to Add an Oracle ACFS File System to a Zone Cluster (CLI)

Perform this procedure to add an Oracle ACFS file system for use by a zone cluster.

Note - Alternatively, you can use the `clsetup` utility to perform this task. See [“How to Add a Cluster File System to a Zone Cluster \(clsetup\)”](#) on page 37.

Before You Begin Ensure that the Oracle ACFS file system is created and ready for use by a zone cluster. See [“How to Create an Oracle ACFS File System”](#) in *Installing and Configuring an Oracle Solaris Cluster 4.4 Environment*.

1. **Assume the root role or become an administrator that provides `solaris.cluster.admin` and `solaris.cluster.modify` authorizations.**

2. **Add the Oracle ACFS file system to the zone cluster.**

Perform this step from the global zone of one node.

```
# clzonecluster configure zonecluster
clzc:zonecluster> add fs
clzc:zonecluster:fs> set dir=mountpoint
clzc:zonecluster:fs> set special=/dev/asm/volume-dev-path
clzc:zonecluster:fs> set type=acfs
clzc:zonecluster:fs> end
clzc:zonecluster> exit
```

3. **Verify that the file system is added to the zone cluster.**

```
# clzonecluster show zonecluster
...
Resource Name: fs
dir: mountpoint
special /dev/asm/volume-dev-path
raw:
type: acfs
options: []
```

```
cluster-control:      true
...
```

Adding Local File Systems to a Specific Zone-Cluster Node

This section describes how to add file systems that are dedicated to a single zone-cluster node. To instead configure file systems for use by the entire zone cluster, go to [“Adding File Systems to a Zone Cluster” on page 29](#).

This section contains the following procedures:

- [“How to Add a Local File System to a Specific Zone-Cluster Node \(CLI\)” on page 45](#)
- [“How to Add a Local ZFS Storage Pool to a Specific Zone-Cluster Node \(CLI\)” on page 47](#)

▼ How to Add a Local File System to a Specific Zone-Cluster Node (CLI)

Perform this procedure to add a local file system to a single, specific zone-cluster node of a specific zone cluster. The file system is not managed by Oracle Solaris Cluster software but is instead passed to the underlying Oracle Solaris zone.

Note - To add a highly available local file system to a zone cluster, perform procedures in [“How to Add a Highly Available Local File System to a Zone Cluster \(clsetup\)” on page 29](#) or [“How to Add a Highly Available Local File System to a Zone Cluster \(CLI\)” on page 31](#).

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 7 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 \(c1setup\)”](#) on page 29 or [“How to Add a ZFS Storage Pool to a Zone Cluster \(CLI\)”](#) on page 36.

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 8 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 login 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 \(clsetup\)”](#) on page 49
- [“How to Add a DID Device to a Zone Cluster \(CLI\)”](#) on page 51
- [“How to Add a Disk Set to a Zone Cluster \(CLI\)”](#) on page 52
- [“How to Add a Raw-Disk Device to a Specific Zone-Cluster Node \(CLI\)”](#) on page 54

▼ 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 \(CLI\)”](#) on page 51
- [“How to Add a Disk Set to a Zone Cluster \(CLI\)”](#) on page 52

To add a raw-disk device to a specific zone-cluster node, go instead to [“How to Add a Raw-Disk Device to a Specific Zone-Cluster Node \(CLI\)”](#) on page 54.

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 login 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 \(`clsetup`\)”](#) on page 49.

1. Assume the `root` role on a node of the global cluster that hosts the zone cluster.

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

2. Identify the DID device to add to the zone cluster.

The device you add must be connected to all nodes of the zone cluster.

```
phys-schost# cldevice list -v
```

3. Add the DID device for use by the zone cluster.

Note - An asterisk (*) is used as a wildcard character in the path name.

```
phys-schost# clzonecluster configure zone-cluster-name
clzc:zone-cluster-name> add device
clzc:zone-cluster-name:device> set match=/dev/did/*dsk/dNs*
clzc:zone-cluster-name:device> end
clzc:zone-cluster-name> verify
clzc:zone-cluster-name> commit
clzc:zone-cluster-name> exit
```

```
match=/dev/did/*dsk/dNs*
```

Specifies the full device path of the DID device

4. Reboot the zone cluster.

The change becomes effective after the zone cluster reboots.

```
phys-schost# clzonecluster reboot zone-cluster-name
```

Example 9 Adding a DID Device to a Zone Cluster (CLI)

The following example adds the DID device d10 to the sczone zone cluster.

```
phys-schost-1# clzonecluster configure sczone
clzc:sczone> add device
clzc:sczone:device> set match=/dev/did/*dsk/d10s*
clzc:sczone:device> end
clzc:sczone> verify
clzc:sczone> commit
clzc:sczone> exit
```

```
phys-schost-1# clzonecluster reboot sczone
```

▼ How to Add a Disk Set to a Zone Cluster (CLI)

Perform this procedure to add an entire Solaris Volume Manager disk set to a zone cluster.

Note - Alternatively, you can use the `clsetup` utility to perform this task. See [“How to Add a Global Storage Device to a Zone Cluster \(clsetup\)”](#) on page 49.

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 10 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 \(clsetup\)”](#) on page 49 or [“How to Add a DID Device to a Zone Cluster \(CLI\)”](#) on page 51.

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 11 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
```


◆◆◆ CHAPTER 2

Administering the Zone Cluster

This chapter provides administrative procedures that affect an entire global cluster or a zone cluster:

For information about adding or removing a node from the cluster, see [Chapter 8](#), “Administering Cluster Nodes” in *Administering an Oracle Solaris Cluster 4.4 Configuration*.

Overview of Zone Cluster Administrative Tasks

You can perform administrative tasks on a zone cluster, such as moving the zone path, preparing a zone cluster to run applications, and cloning a zone cluster. All of these commands must be performed from a node of the global cluster.

You can create a new zone cluster or add a file system or storage device to an existing zone cluster by using the `clsetup` utility to launch the zone cluster configuration wizard. The zones in a zone cluster are configured when you run `clzonecluster install -c` to configure the profiles. See “[Creating and Configuring a Zone Cluster](#)” in *Installing and Configuring an Oracle Solaris Cluster 4.4 Environment* for instructions about using the `clsetup` utility or the `-c config_profile` option.

You can also use the Oracle Solaris Cluster Manager browser interface to create a zone cluster, or add a file system or storage device to it. You can also use the Oracle Solaris Cluster Manager browser interface to edit the zone cluster's `Resource Security` property. Click `Zone Clusters`, click the name of the zone cluster to go to its page, then click the `Solaris Resources` tab to administer zone-cluster components, or click `Properties` to administer zone-cluster properties. 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*.

Note - The Oracle Solaris Cluster commands that you run only from a node in the global cluster are not valid for use with zone clusters. See the appropriate Oracle Solaris Cluster man page for information about the valid use of a command in zone clusters.

TABLE 2 Zone Cluster Tasks

Task	Instructions
Move the zone path to a new zone path	<code>clzonecluster move -f zonepath zone-cluster-name</code>
Prepare the zone cluster to run applications	<code>clzonecluster ready -n nodename zone-cluster-name</code>
Restore the nodes from the United Archives	“How to Restore a Node from the Unified Archive” in <i>Administering an Oracle Solaris Cluster 4.4 Configuration</i>
Configure or install a zone cluster from the United Archive	<p>“How to Configure a Zone Cluster from the Unified Archive” on page 59</p> <p>“How to Install a Zone Cluster from the Unified Archive” on page 60</p> <p>Use a command:</p> <pre>clzonecluster clone -Z target-zone-cluster-name [-m copymethod] source-zone-cluster-name</pre> <p>Halt the source zone cluster before you use the <code>clone</code> subcommand. The target zone cluster must already be configured.</p>
Add a network address to a zone cluster	“How to Add a Network Address to a Zone Cluster” on page 61
Add a node to a zone cluster	“How to Add a Node to an Existing Cluster or Zone Cluster” in <i>Administering an Oracle Solaris Cluster 4.4 Configuration</i>
Remove a node from a zone cluster	“How to Remove a Node From a Zone Cluster” in <i>Administering an Oracle Solaris Cluster 4.4 Configuration</i>
Remove a zone cluster	“How to Unconfigure a Zone Cluster” in <i>Installing and Configuring an Oracle Solaris Cluster 4.4 Environment</i>
Remove a file system from a zone cluster	“How to Remove a File System From a Zone Cluster” on page 65
Remove a storage device from a zone cluster	“How to Remove a Storage Device From a Zone Cluster” on page 68
Restore zone cluster nodes from the Unified Archive	“How to Restore a Node from the Unified Archive” in <i>Administering an Oracle Solaris Cluster 4.4 Configuration</i>
Troubleshoot a node uninstallation	“Troubleshooting a Node Uninstallation” in <i>Administering an Oracle Solaris Cluster 4.4 Configuration</i>
Create, set up, and manage the Oracle Solaris Cluster SNMP Event MIB	“Creating, Setting Up, and Managing the Oracle Solaris Cluster SNMP Event MIB” in <i>Administering an Oracle Solaris Cluster 4.4 Configuration</i>

▼ How to Configure a Zone Cluster from the Unified Archive

Use the `clzonecluster` command to launch an interactive utility to configure a `solaris` or labeled branded zone cluster from the Unified Archive. The `clzonecluster configure` utility lets you specify a *recovery* archive or a *clone* archive.

If you prefer to use the command line rather than the interactive utility to configure a zone cluster from an archive, use the `clzonecluster configure -f command-file` command. See the [clzonecluster\(8CL\)](#) man page for more information.

Note - If the zone cluster that you want to install was already configured using other supported methods, you do not have to configure the zone cluster from a Unified Archive.

1. Create a recovery or clone 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. For more information about using the `archiveadm` command, see the [archiveadm\(8\)](#) man page.

2. Assume the root role on a node of the global cluster that will host the zone cluster.

3. Configure the zone cluster from the recovered or cloned archive in the Unified Archive.

```
phys-schost-1# clzonecluster configure zone-cluster-name
```

The `clzonecluster configure zone-cluster-name` command launches the interactive utility, where you can specify `create -a archive [other-options-such-as-"x"]`. The archive can be either a clone archive or a recovery archive.

Note - The zone cluster members must be added to the configuration before a zone cluster can be created.

The `configure` subcommand uses the `zonecfg` command to configure a zone on each specified machine. The `configure` subcommand lets you specify properties that apply to each node of the zone cluster. These properties have the same meaning as established by the `zonecfg` command for individual zones. The `configure` subcommand supports the configuration of properties that are unknown to the `zonecfg` command. The `configure` subcommand launches an interactive

shell if you do not specify the `-foption`. The `-f` option takes a command file as its argument. The `configure` subcommand uses this file to create or modify zone clusters non-interactively.

▼ How to Install a Zone Cluster from the Unified Archive

You can install a zone cluster from the Unified Archive. The `clzonecluster install` utility lets you specify the absolute path of the archive to use for the installation. The absolute path of the archive should be accessible on all the physical nodes of the cluster where the zone cluster will be installed. The Unified Archive installation can use either a *recovery* archive or a *clone* archive.

If you prefer to use the command line rather than the interactive utility to install a zone cluster from an archive, use the `clzonecluster create -a archive -z archived-zone` command. See the [clzonecluster\(8CL\)](#) man page for more information.

1. Create a recovery or clone 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. For more information about using the `archiveadm` command, see the [archiveadm\(8\)](#) man page.

2. Assume the `root` role on a node of the global cluster that will host the zone cluster.

3. Install the zone cluster from the recovered or cloned archive from the Unified Archive.

```
phys-schost-1# clzonecluster install -a absolute_path_to_archive zone-cluster-name
```

The absolute path of the archive should be accessible on all the physical nodes of the cluster where the zone cluster will be installed. If you have an HTTPS Unified Archive location, specify the SSL certificate, Certificate Authority (CA) certificate, and key files using `-x cert|ca-cert|key=file`.

The Unified Archives do not contain zone cluster node resources. Node resources are specified when the cluster is configured. When you configure a zone cluster from a global zone by using the Unified Archives, you must set the `zonpath`.

If the Unified Archive contains multiple zones, use the `zone-cluster-name` to specify the zone name of the source of the installation. See the [clzonecluster\(8CL\)](#) man page for more information.

Note - If the source you used to create the Unified Archive does not contain the Oracle Solaris Cluster packages, you must run `pkg install ha-cluster-packages` (substituting the specific package name, such as `ha-cluster-minimal` or `ha-cluster-framework-full`). You will need to boot the zone, run the `zlogin` command, and then run the `pkg install` command. This action installs the same packages on the target zone cluster as the global cluster.

4. Boot the new zone cluster.

```
phys-schost-1# clzonecluster boot zone-cluster-name
```

▼ How to Add a Network Address to a Zone Cluster

This procedure adds a network address for use by an existing zone cluster. A network address is used to configure logical host or shared IP address resources in the zone cluster. You can run the `clsetup` utility multiple times to add as many network addresses as you need.

Note - You can also add a network address to a zone cluster by using the Oracle Solaris Cluster Manager browser interface. Click Zone Clusters, click the name of the zone cluster to go to its page, then click the Solaris Resources tab to administer zone-cluster components. 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.**
2. **On the global cluster, configure the cluster file system that you want to use in the zone cluster.**

Start the `clsetup` utility.

```
phys-schost# clsetup
```

The Main Menu is displayed.

3. **Choose the Zone Cluster menu item.**
4. **Choose the Add Network Address to a Zone Cluster menu item.**
5. **Choose the zone cluster where you want to add the network address.**
6. **Choose the property to specify the network address you want to add.**

`address=value`

Specifies the network address used to configure logical host or shared IP address resources in the zone cluster. For example, `192.168.100.101`.

The following types of network addresses are supported:

- A valid IPv4 address, optionally followed by / and a prefix length.
- A valid IPv6 address, which must be followed by / and a prefix length.
- A hostname which resolves to an IPv4 address. Hostnames that resolve to IPv6 addresses are not supported.

See the [zonecfg\(8\)](#) man page for more information about network addresses.

7. To add an additional network address, type `a`.

8. 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 network address to the zone cluster...
```

```
The zone cluster is being created with the following configuration
```

```
/usr/cluster/bin/clzonecluster configure sczone
add net
set address=phys-schost-1
end
```

```
All network address added successfully to sczone.
```

9. When finished, exit the `clsetup` utility.

Adding a Node to a Cluster or Zone Cluster

This section describes how to add a node to a global cluster or a zone cluster. You can create a new zone-cluster node on a node of the global cluster that hosts the zone cluster, as long as that global-cluster node does not already host a node of that particular zone cluster.

Note - The node that you add must run the same version of the Oracle Solaris Cluster software as the cluster it is joining.

Specifying an IP address and NIC for each zone cluster node is optional.

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

1. 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.
2. The cluster software will activate any logical host IP address on any NIC.

If the original zone cluster node did not have a IP address or NIC specified, then you do not need to specify that information for the new zone cluster node.

In this chapter, `phys - schost#` reflects a global-cluster prompt. The `clzonecluster` interactive shell prompt is `clzc: schost>`.

The following table lists the tasks to perform to add a node to an existing cluster. Perform the tasks in the order shown.

TABLE 3 Task Map: Adding a Node to an Existing Global or Zone Cluster

Task	Instructions
Install the host adapter on the node and verify that the existing cluster interconnects can support the new node	Managing Hardware With Oracle Solaris Cluster 4.4
Add shared storage	<p>Add shared storage manually by following the instructions in the Managing Hardware With Oracle Solaris Cluster 4.4.</p> <p>You can also use Oracle Solaris Cluster Manager to add a shared storage device to a zone cluster. Navigate in Oracle Solaris Cluster Manager to the page for the zone cluster and click the Solaris Resources tab. 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.</p>
Add the node to the authorized node list	<code>claccess allow -h node-being-added</code>
Install and configure the software on the new cluster node	Chapter 2, “Installing Software on Global-Cluster Nodes” in Installing and Configuring an Oracle Solaris Cluster 4.4 Environment
Add the new node to an existing cluster	“How to Add a Node to an Existing Cluster or Zone Cluster” in Administering an Oracle Solaris Cluster 4.4 Configuration
If the cluster is configured in an Oracle Solaris Cluster disaster recovery framework partnership, configure the new node as an active participant in the configuration	“How to Add a New Node to a Cluster in a Partnership” in Administering the Disaster Recovery Framework for Oracle Solaris Cluster 4.4

▼ How to Add a Node to an Existing Cluster or Zone Cluster

Before adding an Oracle Solaris host or a virtual machine to an existing global cluster or a zone cluster, ensure that the node has all of the necessary hardware correctly installed and configured, including an operational physical connection to the private cluster interconnect.

For hardware installation information, refer to the [Managing Hardware With Oracle Solaris Cluster 4.4](#) or the hardware documentation that shipped with your server.

This procedure enables a machine to install itself into a cluster by adding its node name to the list of authorized nodes for that cluster.

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. **On a current global-cluster member, assume the `root` role on the current cluster member. Perform these steps from a node of a global cluster.**
 2. **Ensure that you have correctly completed all prerequisite hardware installation and configuration tasks that are listed in the task map for “[Task Map: Adding a Node to an Existing Global or Zone Cluster](#)” in [Administering an Oracle Solaris Cluster 4.4 Configuration](#).**
 3. **Install and configure the software on the new cluster node.**
4. **Use the `scinstall` utility on the new node to configure that node in the cluster.**
 5. **To manually add a node to a zone cluster, you must specify the Oracle Solaris host and the virtual node name.**

You must also specify a network resource to be used for public network communication on each node. In the following example, the zone name is `sczone`, and `sc_ipmp0` is the IPMP group name.

```
clzc:sczone>add node
clzc:sczone:node>set physical-host=phys-cluster-3
clzc:sczone:node>set hostname=hostname3
```



```

clzc:sczone:node>add net
clzc:sczone:node:net>set address=hostname3
clzc:sczone:node:net>set physical=sc_ipmp0
clzc:sczone:node:net>end
clzc:sczone:node>end
clzc:sczone>exit

```

For detailed instructions on configuring the node, see [“Creating and Configuring a Zone Cluster” in *Installing and Configuring an Oracle Solaris Cluster 4.4 Environment*](#).

6. **After you configure the node, reboot the node into cluster mode and install the zone cluster on the node.**

```
# clzonecluster install zone-cluster-name
```

7. **To prevent any new machines from being added to the cluster, from the `clsetup` utility type the number for the option to instruct the cluster to ignore requests to add new machines.**

Press the Return key.

Follow the `clsetup` prompts. This option tells the cluster to ignore all requests over the public network from any new machine that is trying to add itself to the cluster.

8. **Quit the `clsetup` utility.**

See Also [`clsetup\(8CL\)` man page](#).

For a complete list of tasks for adding a cluster node, see [“Task Map: Adding a Node to an Existing Global or Zone Cluster” in *Administering an Oracle Solaris Cluster 4.4 Configuration*](#), “Task Map: Adding a Cluster Node”.

To add a node to an existing resource group, see the [Planning and Administering Data Services for Oracle Solaris Cluster 4.4](#).

▼ How to Remove a File System From a Zone Cluster

A file system can be exported to a zone cluster using either a direct mount or a loopback mount.

Zone clusters support direct mounts for the following:

- UFS local file system
- StorageTek QFS standalone file system

- StorageTek QFS shared file system, when used to support Oracle RAC
- Oracle Solaris ZFS (exported as a data set)
- NFS from supported NAS devices

Zone clusters can manage loopback mounts for the following:

- UFS local file system
- StorageTek QFS standalone file system
- StorageTek QFS shared file system, only when used to support Oracle RAC
- UFS cluster file system

You configure an `HASStoragePlus` or `ScalMountPoint` resource to manage the mounting of the file system. For instructions on adding a file system to a zone cluster, see [“Adding File Systems to a Zone Cluster” in *Installing and Configuring an Oracle Solaris Cluster 4.4 Environment*](#).

An `HASStoragePlus` resource does not monitor a ZFS file system if the file system has its `mountpoint` property set to `none` or `legacy`, or its `canmount` property set to `off`. For all other ZFS file systems, the `HASStoragePlus` resource fault monitor checks if the file system is mounted. If the file system is mounted, the `HASStoragePlus` resource then probes the file system's accessibility by reading and writing to it, depending on the value of the `I00ption` property called `ReadOnly/ReadWrite`.

If the ZFS file system is not mounted or the probe of the file system fails, the resource fault monitor fails and the resource is set to `Faulted`. The RGM will attempt to restart it, determined by the `retry_count` and `retry_interval` properties of the resource. This action results in remounting the file system if the specific `mountpoint` and `canmount` property settings described above are not in play. If the fault monitor continues to fail and exceeds the `retry_count` within the `retry_interval`, the RGM fails the resource over to another node.

The `phys-schost#` prompt reflects a global-cluster prompt. 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.

Note - You can also use the Oracle Solaris Cluster Manager browser interface to remove a file system from a zone cluster. Click `Zone Clusters`, click the name of the zone cluster to go to its page, then click the `Solaris Resources` tab to administer zone-cluster components. 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.**

Some steps in this procedure are performed from a node of the global cluster. Other steps are performed from a node of the zone cluster.

2. Delete the resources related to the file system being removed.

- a. Identify and remove the Oracle Solaris Cluster resource types, such as `HASStoragePlus` and `SUNW.ScalMountPoint`, that are configured for the zone cluster's file system that you are removing.

```
phys-schost# clresource delete -F -Z zone-cluster-name fs_zone_resources
```

- b. If applicable, identify and remove the Oracle Solaris Cluster resources of type `SUNW.qfs` that are configured in the global cluster for the file system that you are removing.

```
phys-schost# clresource delete -F fs_global_resources
```

Use the `-F` option carefully because it forces the deletion of all the resources you specify, even if you did not disable them first. All the resources you specified are removed from the resource-dependency settings of other resources, which can cause a loss of service in the cluster. Dependent resources that are not deleted can be left in an invalid state or in an error state. For more information, see the [clresource\(8CL\)](#) man page.

Tip - If the resource group for the removed resource later becomes empty, you can safely delete the resource group.

3. Determine the path to the file-system mount point directory.

For example:

```
phys-schost# clzonecluster configure zone-cluster-name
```

4. Remove the file system from the zone-cluster configuration.

```
phys-schost# clzonecluster configure zone-cluster-name
```

```
clzc:zone-cluster-name> remove fs dir=filesystemdirectory
```

```
clzc:zone-cluster-name> commit
```

The file system mount point is specified by `dir=`.

5. Verify the removal of the file system.

```
phys-schost# clzonecluster show -v zone-cluster-name
```

Example 12 Removing a Highly Available Local File System in a Zone Cluster

This example shows how to remove a file system with a mount-point directory (/local/ufs-1) that is configured in a zone cluster called sczone. The resource is hasp-rs and is of the type HASStoragePlus.

```
phys-schost# clzonecluster show -v sczone
...
Resource Name:                fs
dir:                          /local/ufs-1
special:                       /dev/md/ds1/dsk/d0
raw:                           /dev/md/ds1/rdisk/d0
type:                          ufs
options:                       [logging]
...
phys-schost# clresource delete -F -Z sczone hasp-rs
phys-schost# clzonecluster configure sczone
clzc:sczone> remove fs dir=/local/ufs-1
clzc:sczone> commit
phys-schost# clzonecluster show -v sczone
```

Example 13 Removing a Highly Available ZFS File System in a Zone Cluster

This example shows to remove a ZFS file systems in a ZFS pool called HAZpool, which is configured in the sczone zone cluster in resource hasp-rs of type SUNW.HASStoragePlus.

```
phys-schost# clzonecluster show -v sczone
...
Resource Name:                dataset
name:                         HAZpool
...
phys-schost# clresource delete -F -Z sczone hasp-rs
phys-schost# clzonecluster configure sczone
clzc:sczone> remove dataset name=HAZpool
clzc:sczone> commit
phys-schost# clzonecluster show -v sczone
```

▼ How to Remove a Storage Device From a Zone Cluster

You can remove storage devices, such as Solaris Volume Manager disk sets and DID devices, from a zone cluster. Perform this procedure to remove a storage device from a zone cluster.

Note - You can also use the Oracle Solaris Cluster Manager browser interface to remove a storage device from a zone cluster. Click Zone Clusters, click the name of the zone cluster to go to its page, then click the Solaris Resources tab to administer zone-cluster components. 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.**
Some steps in this procedure are performed from a node of the global cluster. Other steps can be performed from a node of the zone cluster.

2. **Delete the resources related to the devices being removed.**

Identify and remove the Oracle Solaris Cluster resource types, such as SUNW.HAStoragePlus and SUNW.ScalDeviceGroup, that are configured for the zone cluster's devices that you are removing.

```
phys-schost# clresource delete -F -Z zone-cluster dev_zone_resources
```

3. **Determine the match entry for the devices to be removed.**

```
phys-schost# clzonecluster show -v zone-cluster
...
Resource Name:      device
match:              <device_match>
...
```

4. **Remove the devices from the zone-cluster configuration.**

```
phys-schost# clzonecluster configure zone-cluster
clzc:zone-cluster-name> remove device match=devices-match
clzc:zone-cluster-name> commit
clzc:zone-cluster-name> end
```

5. **Reboot the zone cluster.**

```
phys-schost# clzonecluster reboot zone-cluster
```

6. **Verify the removal of the devices.**

```
phys-schost# clzonecluster show -v zone-cluster
```

Example 14 Removing a Solaris Volume Manager Disk Set From a Zone Cluster

This example shows how to remove a Solaris Volume Manager disk set called `apachedg` configured in a zone cluster called `sczone`. The set number of the `apachedg` disk set is 3. The devices are used by the `zc_rs` resource that is configured in the cluster.

```
phys-schost# clzonecluster show -v sczone
...
Resource Name:      device
match:              /dev/md/apachedg/*dsk/*
Resource Name:      device
match:              /dev/md/shared/3/*dsk/*
...
phys-schost# clresource delete -F -Z sczone zc_rs

phys-schost# ls -l /dev/md/apachedg
lrwxrwxrwx 1 root root 8 Jul 22 23:11 /dev/md/apachedg -> shared/3
phys-schost# clzonecluster configure sczone
clzc:sczone> remove device match=/dev/md/apachedg/*dsk/*
clzc:sczone> remove device match=/dev/md/shared/3/*dsk/*
clzc:sczone> commit
clzc:sczone> end
phys-schost# clzonecluster reboot sczone
phys-schost# clzonecluster show -v sczone
```

Example 15 Removing a DID Device From a Zone Cluster

This example shows how to remove DID devices `d10` and `d11`, which are configured in a zone cluster called `sczone`. The devices are used by the `zc_rs` resource that is configured in the cluster.

```
phys-schost# clzonecluster show -v sczone
...
Resource Name:      device
match:              /dev/did/*dsk/d10*
Resource Name:      device
match:              /dev/did/*dsk/d11*
...
phys-schost# clresource delete -F -Z sczone zc_rs
phys-schost# clzonecluster configure sczone
clzc:sczone> remove device match=/dev/did/*dsk/d10*
clzc:sczone> remove device match=/dev/did/*dsk/d11*
clzc:sczone> commit
clzc:sczone> end
phys-schost#
phys-schost# clzonecluster show -v sczone
```

Index

A

adding

- file systems to a zone cluster, 29
- local file systems to a zone cluster, 45
- network address to a zone cluster, 61
- nodes, 62
 - to a global cluster, 64
 - to a zone cluster, 64
- storage devices to zone clusters, 48

admin_low template

- Trusted Extensions, 12

administering

- zone clusters, 57

B

brand property, 15

BUI

tasks you can perform

- add a file system to a zone cluster, 29
- add storage a zone cluster, 49
- create a zone cluster, 14
- view capped memory for a zone cluster, 17

C

capped-cpu scope type

- zone clusters, 17

cipso template

- global cluster IP addresses, 12
- zone cluster IP addresses, 23

clsetup

- adding a network address to a zone cluster, 61

- administering a zone cluster, 57

clsetup utility

adding

- file systems to a zone cluster, 37
- global storage devices to a zone cluster, 49
- highly available local file systems to a zone cluster, 29
- ZFS storage pools to a zone cluster, 33

- creating a zone cluster, 13

navigation tip, 14

cluster file systems

- adding to a zone cluster, 40
- adding to a zone cluster with clsetup, 37

clzonecluster command

- booting the zone cluster, 21
- installing the zone cluster, 19

configuring

- Trusted Extensions for use with zone clusters, 10, 22
- Trusted Extensions software, 11
- zone clusters, 10

cpu-shares property, 16

creating

- zone clusters, 13

D

dedicated-cpu scope type

- zone clusters, 16

default router, 17

defrouter property, 17

DID devices

- adding to a zone cluster, 51

direct mount
 exporting a file system to a zone cluster, 65
disk sets
 adding to a zone cluster, 52

E

enable_priv_net property, 15
enabling
 remote login, 11
examples
 adding
 local file systems to a zone-cluster node, 46
 local ZFS pools to a zone-cluster node, 48
 raw-disk devices to a zone-cluster node, 55
 adding devices to a zone cluster
 DID devices, 52
 disk sets, 53
 adding file systems to a zone cluster
 cluster file systems, 41
 highly available local file systems, 33
 StorageTek QFS shared file systems, 43
 ZFS storage pools, 36

F

Fair Share Scheduler shares, 16
file system
 removing it in a zone cluster, 57
file systems
 adding to a zone cluster, 29
 highly available local file systems, 29
 ZFS storage pools, 33

G

global storage devices
 adding to a zone cluster, 49

H

highly available local file systems

 adding to a zone cluster, 29
 highly available ZFS storage pools
 adding to a zone cluster, 33
 hostmodel property, 24
 hostname property, 17

I

importing
 zones into zone clusters, 24
installing
 Trusted Extensions software, 11
IP strict multihoming, 24
ip-type property, 15

L

labeled brand zone clusters
 configuring with Trusted Extensions, 22
 options, 19
LDAP configuration for Trusted Extensions, 12
lightweight processes
 zone clusters, 16
local
 file systems
 adding to a zone cluster, 31
 local file systems
 adding to a specific zone-cluster node, 45
 adding to a zone cluster, 45, 45
 See also file systems, adding to a zone cluster
 local ZFS storage pools
 adding to a specific zone-cluster node, 47
 locked memory, 17
 locked property, 17
 loopback mount
 exporting a file system to a zone cluster, 65
LWPs
 zone clusters, 16

M

max-lwps property, 16

- max-msg-ids property, 16
 - max-sem-ids
 - property, 16
 - max-shm-ids property, 16
 - max-shm-memory property, 16
 - message queue IDs
 - zone clusters, 16
 - mount options for cluster file systems
 - StorageTek QFS software, 38
 - UFS, 31, 38
- N**
- NAS devices
 - troubleshooting in a zone cluster installation, 14
 - ncpus property, 16
 - network address
 - adding to a zone cluster, 61
 - nodes
 - adding, 62
 - num_xip_zoneclusters property, 15
- O**
- Oracle ACFS file systems
 - adding to a zone cluster, 44
 - Oracle Solaris Cluster Manager *See* BUI
 - adding
 - zfs storage to a zone cluster, 33
 - tasks you can perform
 - add a file system to a zone cluster, 57
 - add a network address to a zone cluster, 61
 - add a storage device to a zone cluster, 57
 - add shared storage to a zone cluster, 63
 - adding cluster file system to a zone cluster, 37
 - adding highly available local file systems to a zone cluster, 29
 - adding shared storage to a zone cluster, 49
 - create a zone cluster, 57
 - edit the zone cluster's Resource Security property, 57
 - remove a file system from a zone cluster, 66
 - remove a storage device from a zone cluster, 69
 - Oracle Solaris software
 - installing
 - Trusted Extensions feature, 10
- P**
- physical memory
 - zone clusters, 17
 - physical property, 17, 17
- R**
- raw-disk devices
 - adding to a zone cluster, 54
 - remote login
 - Trusted Extensions, 11
- S**
- scope property, 16
 - semaphore IDs
 - zone clusters, 16
 - shared memory
 - zone clusters, 16
 - shared memory IDs
 - zone clusters, 16
 - software updates
 - zone clusters, 24
 - solaris brand zone clusters
 - options, 19
 - Solaris Volume Manager
 - disk sets
 - adding to a zone cluster, 52
 - storage devices
 - adding global storage devices to a zone cluster, 49
 - adding raw-disk devices to a zone cluster, 54
 - adding to zone clusters, 48
 - StorageTek QFS
 - adding shared file systems to a zone cluster, 42
 - StorageTek QFS software
 - mount options, 38

swap
 memory for zone clusters, 17
 swap property, 17
sysid configuration for zone clusters, 20

T

tncfg template, 12
troubleshooting
 NAS devices in a zone cluster, 14
Trusted Extensions
 admin_low template, 12
 cipso template
 global cluster, 12
 zone cluster, 23
 configuring, 10
 configuring a zone cluster, 22
 installation, 11
 installing, 10
 IP strict multihoming, 24
 LDAP, 12
 password files, 19
 remote login, 11
 tncfg template, 12
 txzonemgr BUI, 19
 zoneshare script, 11

U

/usr/cluster/bin/clsetup utility
 adding
 file systems to a zone cluster, 37
 global storage devices to a zone cluster, 49
 highly available local file systems to a zone cluster, 29
 ZFS storage pools to a zone cluster, 33
 creating a zone cluster, 13
 navigation tip, 14
/usr/cluster/bin/clzonecluster command
 booting the zone cluster, 21
 installing the zone cluster, 19
Unified Archive

 configuring a zone cluster, 59
 installing a zone cluster, 60

Z

ZFS
 removing a file system, 65
ZFS storage pools
 adding a local pool to a zone-cluster node, 47
 adding to a zone cluster, 33, 36
zone cluster
 adding network addresses, 61
 administration, 57
 cloning, 57
 configuring from the Unified Archive, 59
 installing from the Unified Archive, 60
 moving a zonempath, 57
 preparing it for applications, 57
 removing a file system, 57
 supported direct mounts, 65
zone clusters
 adding an Oracle ACFS file system, 44
 adding devices
 DID devices, 51
 disk sets, 52
 global storage devices, 49
 raw-disk devices, 54
 adding file systems, 29
 cluster file systems, 37, 40
 highly available local file systems, 29
 highly available ZFS storage pools, 33
 local file systems, 31
 local file systems for a specific zone-cluster node, 45
 local ZFS storage pools, 47
 StorageTek QFS, 42
 ZFS storage pools, 33, 36
 adding local file systems, 45
 adding storage devices, 48
 brand property, 15
 capped-cpu, 17
 configuring, 10, 10
 configuring to use Trusted Extensions, 22

- creating, 13
- dedicated-cpu, 16
- default router, 17
- defrouter property, 17
- enable_priv_net property, 15
- Fair Share Scheduler shares, 16
- hostmodel property, 24
- hostname property, 17
- importing zones, 24
- IP strict multihoming, 24
- ip-type property, 15
- labeled brand options, 19
- lightweight processes, 16
- locked memory, 17
- locked property, 17
- message queue IDs, 16
- ncpus property, 16
- num_xip_zoneclusters property, 15
- physical memory, 17
- physical property, 17, 17
- scope property, 16
- semaphore IDs, 16
- shared memory, 16
- shared memory IDs, 16
- software updates, 24
- solaris brand options, 19
- swap memory, 17
- swap property, 17
- sysid configuration, 20
- zonepath property, 15
- zone path
 - moving, 57
- zone-cluster nodes
 - specifying IP address and NIC, 62
- zonepath property, 15
- zones
 - importing into zone clusters, 24
- zoneshare Trusted Extension script, 11

