Open HA Cluster Installation Guide

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

Part No: 820–7821–10
June 2009, Revision A

090603@22510

Contents

Preface

The *Open HA Cluster Installation Guide* contains guidelines and procedures for installing the Open HA Cluster software on both SPARC® based systems and x86 based systems.

Note – This Sun Cluster release supports systems that use the SPARC and x86 families of processor architectures: UltraSPARC, SPARC64, AMD64, and Intel 64. In this document, x86 refers to the larger family of 32–bit and 64-bit x86 compatible products. Information in this document pertains to all platforms unless otherwise specified.

The instructions in this book assume knowledge of OpenSolaris software.

Using UNIX Commands

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

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

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

Typographic Conventions

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

TABLE P–1 Typographic Conventions

Typeface	Meaning	Example
`AaBbCc123`	The names of commands, files, and directories, and onscreen computer output	Edit your `.login` file. Use `ls -a` to list all files. `machine_name% you have mail.`
`AaBbCc123`	What you type, contrasted with onscreen computer output	`machine_name%` **`su`** `Password:`
aabbcc123	Placeholder: replace with a real name or value	The command to remove a file is `rm` *filename*.
AaBbCc123	Book titles, new terms, and terms to be emphasized	Read Chapter 6 in the *User's Guide*. A *cache* is a copy that is stored locally. Do *not* save the file. **Note:** Some emphasized items appear bold online.

Shell Prompts in Command Examples

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

TABLE P–2 Shell Prompts

Shell	Prompt
C shell	`machine_name%`
C shell for superuser	`machine_name#`
Bourne shell and Korn shell	`$`
Bourne shell and Korn shell for superuser	`#`

Related Documentation

Information about related Open HA Cluster 2009.06 and OpenSolaris 2009.06 software topics is available in the documentation that is listed in the following table.

Topic	Documentation
Open HA Cluster 2009.06 Release Notes	Open HA Cluster 2009.06 Release Notes
Sun Cluster 3.2 1/09 software, data services, and hardware	*Sun Cluster 3.2 1/09 Documentation Center*
OpenSolaris installation	Getting Started With OpenSolaris 2009.06
OpenSolaris system administration	OpenSolaris System Administrator Collection
OpenSolaris software development	OpenSolaris Software Developer Collection
OpenSolaris man pages	OpenSolaris Reference Manual Collection
COMSTAR iSCSI storage software	COMSTAR Administration

◆ ◆ ◆ CHAPTER 1

Planning the Open HA Cluster Configuration

This chapter provides planning information and guidelines specific to an Open HA Cluster 2009.06 configuration. The information in this chapter supplements or replaces guidelines in Chapter 1, "Planning the Open HA Cluster Configuration," for those features and functionality that are supported in an Open HA Cluster 2009.06 configuration. For information about Sun Cluster features that are not supported or are limited in an Open HA Cluster 2009.06 configuration, see Open HA Cluster 2009.06 Release Notes.

Hardware and Software

The following are the hardware and software requirements or defaults for an Open HA Cluster configuration:

- **Operating system** – An Open HA Cluster 2009.06 configuration runs only on OpenSolaris 2009.06 software.
- **Hardware platform** – An Open HA Cluster 2009.06 configuration runs on either SPARC based platforms or on 32-bit or 64-bit x86 based platforms.

  All nodes in a cluster must run on the same platform. For x86 based platforms, you cannot use both 32–bit machines and 64–bit machines in the same cluster.
- **Hardware topology** – An Open HA Cluster 2009.06 configuration consists of the following hardware components:
  - Exactly two physical cluster nodes that run on the same subnet
  - At least one network adapter per node
  - Shared storage is optional
- **Root file system** – ZFS is the default root file system.

Note – The creation of a `/globaldevices` partition for use as the global-devices namespace is incompatible with a ZFS root file system. You must either configure a `lofi` device to host the global-devices namespace, or create the `/globaldevices` partition on a UFS root file system.

- **System shell** – Korn shell 93 (`ksh93`) is the default system shell.
- **Administrator role** – By default, the initial user account has the Primary Administrator profile.
- **Network interface manager** – By default, Network Auto-Magic (NWAM) is the default network interface manager. However, NWAM is incompatible with Open HA Cluster 2009.06 software and you must disable it before you configure Open HA Cluster 2009.06 software.
- **DHCP** – Open HA Cluster software uses certain network configuration files in ways that are incompatible with running DHCP clients with IPMP. Therefore, cluster nodes cannot be DHCP clients. You must disable DHCP and instead configure a static IP address for the public network.

IPMP Groups

Observe the following guidelines for IPMP groups in an Open HA Cluster configuration:

- **Link-based IPMP groups** –At cluster installation time, automatically created IPMP groups are configured as link-based groups. If you want an IPMP group to be probe based, you must manually edit the `/etc/hostname.`*adapter* file on each node to add test addresses.
- **LogicalHostname and SharedAddress resources** – If you configure a `LogicalHostname` or `SharedAddress` resource with a hostname that uses a single adapter, the automatically created IPMP group for that adapter is configured for link-based monitoring. You can afterwards modify the `/etc/hostname.`*adapter* files for these IPMP groups to make them probe based.

Private Interconnect

Observe the following guidelines for the private interconnect in an Open HA Cluster configuration:

- **Optional private interconnect** – The use of a physical private interconnect is optional. You can instead use the public network for cluster traffic by configuring virtual network interfaces, or VNICs.

- **Creation of VNICs** – To use VNICs for the cluster transport, you can either configure the VNICs in advance or use the `scinstall` utility in Custom Mode to create them when you establish the cluster. For information about manually creating a VNIC, see "How to Create a Virtual Network Interface (VNIC)" on page 16.

  When you use the `scinstall` utility in Custom Mode to create a new VNIC, you specify the following information:

  - The name of the physical adapter, or NIC, to use
  - The physical adapter's MAC address or choose automatic selection (`auto`)
  - The name to give the VNIC, using the naming convention `vnic`*N*

  The VNICs are created when cluster configuration and establishment is performed.
- **Autodiscovery of adapters** - If you use the `scinstall` utility in Custom Mode to create a VNIC for use by the first cluster node you configure, you cannot use autodiscovery of adapters for the rest of the cluster nodes. When you are prompted whether to use autodiscovery, type "No".
- **Coexistence of physical and virtual adapters** – You can use a combination of physical and virtual adapters in the cluster or on a single node. However, if there is a large difference in the bandwidth for the different NICs and VNICs, performance can be impacted by the lower-speed NICs during peak loads. Ensure that the NICs and VNICs you use in the same cluster have comparable bandwidth.
- **IP Security Architecture (IPsec)** – Only use IPsec with Internet Key Exchange (IKE) for key management. Do not use the manual-key form of key management when you configure IPsec in an Open HA Cluster configuration.

iSCSI Storage

iSCSI is a protocol that enables clients, called initiators, to send SCSI commands to SCSI storage devices, called targets, on remote servers. It is a Storage Area Network (SAN) protocol that enables the consolidation of storage into data-center storage arrays, while providing hosts with the illusion of locally attached disks. The use of iSCSI does not require special-purpose cabling. Instead, communication is run over long distances by using the existing network infrastructure.

Observe the following guidelines for configuring iSCSI storage in an Open HA Cluster configuration:

- **COMSTAR** – Only COMSTAR based iSCSI target implementations are supported in an Open HA Cluster 2009.06 configuration.
- **iSCSI target location** – A disk that is exported as an iSCSI target must be a local disk that is directly attached to the cluster node that hosts the iSCSI target. You cannot use a disk as an iSCSI target if it is hosted by multiple nodes or if it is not directly attached to the cluster node.

- **Topology** – Configure the hardware connections as shown in the following diagram. This diagram shows a two-node Open HA Cluster 2009.06 configuration that uses COMSTAR and a failover ZFS storage pool to provide high availability. The arrows indicate iSCSI connections. One or more connections provide a path from each node to the same disk on Node 1. In the cluster DID namespace, this becomes a single DID device, with paths from both nodes. Similarly, one or more connections provide a path from each node to the same disk on Node 2. This creates a second DID device. The mirroring of these two DID devices by using a ZFS storage pool creates a failover ZFS file system in the Open HA Cluster configuration.

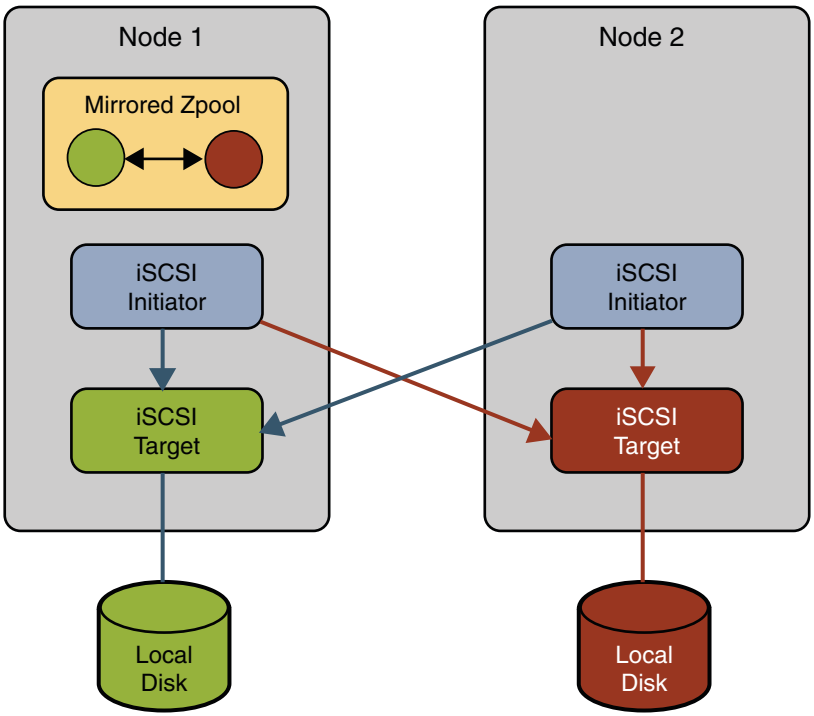

FIGURE 1–1 Cluster Topology Using Unshared COMSTAR Storage

◆ ◆ ◆ CHAPTER 2

Installing Software on Cluster Nodes

This chapter provides procedures to install Open HA Cluster 2009.06 software on cluster nodes.

The following procedures are in this chapter:

Installing the Software

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

▼ How to Install OpenSolaris Software on the Cluster Nodes

Perform this procedure to install OpenSolaris 2009.06 software on each cluster node.

1 **Connect to a machine that you will install as a cluster node.**

2 **Become superuser.**

3 **Install the OpenSolaris 2009.06 software.**

Follow instructions in OpenSolaris Automated Installer Guide. For x86 based platforms, you can alternatively follow instructions in Installing OpenSolaris From the Live CD.

4 **If you will use COMSTAR, install iSCSI packages.**

```
phys-schost# pkg install SUNWstmf SUNWiscsi SUNWiscsit
```

Next Steps Go to "How to Prepare to Download Open HA Cluster Software" on page 14.

▼ How to Prepare to Download Open HA Cluster Software

Perform this procedure on each machine that you intend to install as a cluster node. In addition, if you intend to use a quorum server as a quorum device, perform this procedure on the machine that you intend to install as the quorum server.

Before You Begin Ensure that OpenSolaris 2009.06 software is installed. See "How to Install OpenSolaris Software on the Cluster Nodes" on page 13.

1. **Become superuser on a machine to which you want to download software.**

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

2. **Open a web browser.**

3. **Register with `pkg.sun.com` and obtain the necessary key and certificate for the Open HA Cluster repository, `pkg.sun.com/opensolaris/ha-cluster/`.**

   The following commands show the creation of the directory `/var/pkg/ssl`to contain the downloaded key file and certificate file for the Open HA Cluster repository.

   ```
   phys-schost# mkdir -m 0755 -p /var/pkg/ssl
   phys-schost# cp -i ~/Download/Open_HA_Cluster_2009.06.key.pem /var/pkg/ssl
   phys-schost# cp -i ~/Download/Open_HA_Cluster_2009.06.certificate.pem /var/pkg/ssl
   ```

   For more information, see Using Keys and Certificates for Repositories.

4. **Set the location of the Open HA Cluster 2009.06 package repository.**

   Specify the location of the key file and the certificate file that you obtained in the previous step.

   ```
   phys-schost# /usr/bin/pkg set-publisher \
   -k /var/pkg/ssl/Open_HA_Cluster_2009.06.key.pem \
   -c /var/pkg/ssl/Open_HA_Cluster_2009.06.certificate.pem \
   -O https://pkg.sun.com/opensolaris/ha-cluster/ ha-cluster
   ```

5. **Verify the `ha-cluster` publisher and repository location.**

   ```
   phys-schost# /usr/bin/pkg publisher
   PUBLISHER                         URL
   opensolaris.org                   http://pkg.opensolaris.org/release
   ha-cluster                        https://pkg.sun.com/opensolaris/ha-cluster/
   ```

Next Steps If you want to install a machine as a quorum server to use as the quorum device in your cluster, go next to "How to Install and Configure Quorum Server Software" on page 15.

Otherwise, if you will use virtual network interfaces (VNICs) in the cluster private interconnect and want to preconfigure the VNICs, go to "How to Create a Virtual Network Interface (VNIC)" on page 16. You can alternatively create VNICs during cluster configuration by running the `scinstall` utility in Custom Mode.

Otherwise, go to "How to Install Open HA Cluster 2009.06 Software" on page 17.

▼ How to Install and Configure Quorum Server Software

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

Before You Begin Perform the following tasks:

- Ensure that you are registered with `pkg.sun.com` and that the machine to install has the necessary setup to download packages. See "How to Prepare to Download Open HA Cluster Software" on page 14.
- Ensure that the machine that you choose for the quorum server has at least 1 Mbyte of disk space available for Quorum Server software installation.
- Ensure that the quorum-server machine is connected to a public network that is accessible to the cluster nodes.
- Disable the spanning tree algorithm on the Ethernet switches for the ports that are connected to the cluster public network where the quorum server will run.

1. **Become superuser on the machine to install with Quorum Server software.**

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

2. **Install the Quorum Server package.**

   ```
   quorumserver# /usr/bin/pkg install ha-cluster-quorum-server-full
   ```

3. **(Optional) Add the Quorum Server binary location to your `PATH` environment variable.**

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

4. **(Optional) Add the Quorum Server man-page location to your `MANPATH` environment variable.**

   ```
   quorumserver# MANPATH=$MANPATH:/usr/cluster/man
   ```

5. **Configure the quorum server.**

   Add the following entry to the `/etc/scqsd/scqsd.conf` file to specify configuration information about the quorum server.

Identify the quorum server by using at least one of either an instance name or a port number. You must provide the port number, but the instance name is optional.

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

/usr/cluster/lib/sc/scqsd [-d *quorumdirectory***] [-i** *instancename***] -p** *port*

`-d` *quorumdirectory*	The path to the directory where the quorum server can store quorum data. The quorum-server process creates one file per cluster in this directory to store cluster-specific quorum information. By default, the value of this option is `/var/scqsd`. This directory must be unique for each quorum server that you configure.
`-i` *instancename*	A unique name that you choose for the quorum-server instance.
`-p` *port*	The port number on which the quorum server listens for requests from the cluster.

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

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

8 Start the newly configured quorum server and its SMF service.

```
quorumserver# svcadm enable svc:/system/cluster/quorumserver:default
```

Next Steps If you will use virtual network interfaces (VNICs) in the cluster private interconnect and want to preconfigure the VNICs, go to "How to Create a Virtual Network Interface (VNIC)" on page 16. You can alternatively create VNICs during cluster configuration by running the `scinstall` utility in Custom Mode.

Otherwise, go to "How to Install Open HA Cluster 2009.06 Software" on page 17.

▼ How to Create a Virtual Network Interface (VNIC)

Perform this optional procedure to create virtual network interfaces (VNICs) to configure in the cluster interconnect.

Note – If you intend to run the `scinstall` utility in Typical mode to establish the cluster, use this procedure to preconfigure the VNICs that you will use in the private interconnect.

You can alternatively use the `scinstall` utility in Custom mode to configure VNICs for you during initial cluster configuration.

- **Create a VNIC.**

  Follow procedures in "How to Create a Virtual Network Interface" in *System Administration Guide: Network Interfaces and Network Virtualization*.

Note – To configure additional VNICs after you establish the cluster, use the `clsetup` utility.

Next Steps Install Open HA Cluster 2009.06 software on the cluster nodes. Go to "How to Install Open HA Cluster 2009.06 Software" on page 17.

▼ How to Install Open HA Cluster 2009.06 Software

Perform each step in this procedure as superuser on each cluster node.

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

Note – You cannot add or remove individual packages that are part of the `ha-cluster-minimal` framework software packages except by complete reinstallation or uninstallation. See "How to Uninstall Open HA Cluster Software" on page 53 for procedures to remove the cluster framework packages.

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

Before You Begin Ensure on each node that OpenSolaris 2009.06 software is installed and that NWAM is disabled. See "How to Install OpenSolaris Software on the Cluster Nodes" on page 13 for instructions.

1. **Connect to a machine to install as a cluster node.**

2. **Install the Open HA Cluster 2009.06 software.**

   ```
   phys-schost# /usr/bin/pkg install package
   ```

The following table lists the primary group packages for Open HA Cluster 2009.06 software and the principal features that each group package contains. You must install at least the `ha-cluster-minimal` group package.

Feature	ha-cluster-full	ha-cluster-framework-full	ha-cluster-data-services-full	ha-cluster-minimal
Framework	X	X	X	X
Agents	X		X	
Localization	X	X	X	
Framework Man Pages	X	X		
Data Service Man Pages	X		X	
Agent Builder	X	X		X
Generic Data Service	X	X	X	

3 Verify that the package installed successfully.

Output is similar to the following example, which checks the installation state of the `ha-cluster-full` group package.

```
$ /usr/bin/pkg info -r ha-cluster-full
          Name: ha-cluster-full
       Summary: Sun Cluster full installation group package
      Category: System/HA Cluster
         State: Installed
     Publisher: ha-cluster
       Version: 2009.6
...
```

4 (Optional) Add the Open HA Cluster binary location to your `PATH` environment variable.

```
phys-schost# PATH=$PATH:/usr/cluster/bin
```

5 (Optional) Add the Open HA Cluster man-page location to your `MANPATH` environment variable.

```
phys-schos# MANPATH=$MANPATH:/usr/cluster/man
```

Next Steps Establish the new cluster. Go to "How to Configure Open HA Cluster Software on All Nodes (`scinstall`)" on page 19.

◆ ◆ ◆ CHAPTER 3

Establishing the Cluster

This chapter provides procedures for how to establish a cluster.

The following procedures are in this chapter:

Establishing a New Cluster

This section provides information and procedures to establish a new cluster.

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

Perform this procedure from one node of the cluster to configure Open HA Cluster software on both nodes of the cluster.

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

Before You Begin Perform the following tasks:

- Ensure that Open HA Cluster software packages are installed on each node. See "How to Install Open HA Cluster 2009.06 Software" on page 17.
- Determine which mode of the `scinstall` utility you will use, Typical or Custom.

Note – Use Custom mode to have the `scinstall` utility create a new virtual network interface (VNIC) for the cluster private interconnect.

You can use either Typical or Custom mode if you have preconfigured VNICs.

For the Typical installation of Open HA Cluster software, `scinstall` automatically specifies the following configuration defaults.

Component	Default Value
Private-network address	`172.16.0.0`
Private-network netmask	`255.255.240.0`
Cluster-transport adapters	Exactly two adapters
Cluster-transport switches	`switch1` and `switch2`
Global fencing	Enabled
Global-devices file-system name	Looks for a `/globaldevices` partition, then prompts you to configure a `lofi` device
Installation security (DES)	Limited

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

Component	Description/Example	Answer	
Cluster Name	What is the name of the cluster that you want to establish?		
Cluster Nodes	List the name of the other cluster node planned for the initial cluster configuration.		
Cluster Transport Adapters and Cables	What are the names of the two cluster-transport adapters that attach the node to the private interconnect? *(To specify preconfigured VNICs, select Other from the list of adapters.)*	*First*	*Second*

Component	Description/Example	Answer	
(VLAN adapters only)	Will this be a dedicated cluster transport adapter? *(Answer No if using tagged VLAN adapters.)*	Yes | No	Yes | No
	If no, what is the VLAN ID for this adapter?		
Quorum Configuration	Do you want to disable automatic quorum device selection? *(Answer Yes if any shared storage is not qualified to be a quorum device or if you want to configure a quorum server as a quorum device.)*	Yes | No	
Check	Do you want to interrupt cluster creation for `cluster check` errors?	Yes | No	
lofi Device	Do you want to use a lofi device? *(Answer Yes.)*	Yes	

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

Component	Description/Example	Answer	
Cluster Name	What is the name of the cluster that you want to establish?		
Cluster Nodes	List the name of the other cluster node planned for the initial cluster configuration.		
Authenticating Requests to Add Nodes	Do you need to use DES authentication?	No | Yes	
Minimum Number of Private Networks	Should this cluster use at least two private networks?	Yes | No	
Point-to-Point Cables	Does this cluster use switches?	Yes | No	
Cluster Switches	Transport switch name: Defaults: `switch1` and `switch2`	*First*	*Second*
Cluster Transport Adapters and Cables	Node name (*the node from which you run scinstall*):		
	Transport adapter name, VNIC name (`vnic`*N*), or create a new VNIC:	*First*	*Second*
(VLAN adapters only)	Will this be a dedicated cluster transport adapter? *(Answer No if using tagged VLAN adapters.)*	Yes | No	Yes | No
	If no, what is the VLAN ID for this adapter?		

Component	Description/Example	Answer	
	Where does each transport adapter connect to (*a switch or another adapter*)? Switch defaults: `switch1` and `switch2`	*First*	*Second*
	If a transport switch, do you want to use the default port name?	Yes | No	Yes | No
	If no, what is the name of the port that you want to use?		
	Do you want to use autodiscovery to list the available adapters for the other nodes? *(If VNICs are configured on any cluster node, autodiscovery is available only if VNICs are preconfigured on all nodes.)* If no, supply the following information for each additional node:	Yes | No	
Specify for each additional node	Node name:		
	Transport adapter name:	*First*	*Second*
(VLAN adapters only)	Will this be a dedicated cluster transport adapter? *(Answer No if using tagged VLAN adapters.)*	Yes | No	Yes | No
	If no, what is the VLAN ID for this adapter?		
	Where does each transport adapter connect to (*a switch or another adapter*)? Defaults: `switch1` and `switch2`	*First*	*Second*
	If a transport switch, do you want to use the default port name?	Yes | No	Yes | No
	If no, what is the name of the port that you want to use?		
Network Address for the Cluster Transport	Do you want to accept the default network address (`172.16.0.0`)?	Yes | No	
	If no, which private network address do you want to use?	___.___.___.___	
	Do you want to accept the default netmask?	Yes | No	
	If no, what are the maximum numbers of nodes, private networks, and zone clusters that you expect to configure in the cluster? **Note –** Zone clusters are not available in the Open HA Cluster 2009.06 release.	_____ nodes _____ networks _____ zone clusters	
	Which netmask do you want to use? *(Choose from the values calculated by scinstall or supply your own.)*	___.___.___.___	
Global Fencing	Do you want to disable global fencing? *(Answer No unless the shared storage does not support SCSI reservations or unless you want systems that are outside the cluster to access the shared storage.)*	Yes | No	Yes | No

Component	Description/Example	Answer	
Quorum Configuration	Do you want to disable automatic quorum device selection? *(Answer Yes if any shared storage is not qualified to be a quorum device or if you want to configure a quorum server as a quorum device.)*	Yes | No	Yes | No
Global Devices File System (*specify for each node*)	Do you want to use the default name of the global-devices file system (`/globaldevices`)? *(Answer No.)*	Yes | No	
	If no, do you want to use an already-existing file system? *(Answer No.)*	Yes | No	
	What is the name of the file system that you want to use? *(Leave blank.)*		
Check	Do you want to interrupt cluster creation for `cluster check` errors?	Yes | No	

Note – For the global-devices file system, use only a `lofi` device. Do not attempt to configure a dedicated `/globaldevices` partition. Respond "No" to all prompts that ask whether to use or create a file system. After you decline to configure a file system, the `scinstall` utility prompts you to create a `lofi` device.

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

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

1 On each node to configure in a cluster, become superuser.

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

2 On each node, disable Network Auto-Magic (NWAM).

NWAM activates a single network interface and disables all others. For this reason, NWAM cannot coexist with Open HA Cluster 2009.06 software and you must disable it before you configure or run your cluster.

a. On each cluster node, determine whether NWAM is enabled or disabled.

```
phys-schost# svcs -a | grep /network/physical
```

- If NWAM is enabled, output is similar to the following:

```
online          Mar_13   svc:/network/physical:nwam
disabled        Mar_13   svc:/network/physical:default
```

- If NWAM is disabled, output is similar to the following:

```
disabled         Mar_13  svc:/network/physical:nwam
online           Mar_13  svc:/network/physical:default
```

b. If NWAM is enabled on a node, disable it.

```
phys-schost# svcadm disable svc:/network/physical:nwam
phys-schost# svcadm enable svc:/network/physical:default
```

3 On each node, configure each public-network adapter.

a. Determine which adapters are on the system.

```
phys-schost# dladm show-link
```

b. Plumb an adapter.

```
phys-schost# ifconfig adapter plumb up
```

c. Assign an IP address and netmask to the adapter.

```
phys-schost# ifconfig adapter IPaddress netmask + netmask
```

d. Verify that the adapter is up.

Ensure that the comment output contains the `UP` flag.

```
phys-schost# ifconfig -a
```

e. Create a configuration file for the adapter.

This file ensures that the configuration of the adapter persists across reboots.

```
phys-schost# vi /etc/hostname.adapter
IPaddress
```

f. Repeat Step b through Step e for each public-network adapter on both nodes.

g. On both nodes, add an entry to the `/etc/inet/hosts` file for each public-network adapter that you configured on each node.

```
phys-schost# vi /etc/inet/hosts
hostname IPaddress
```

h. If you use a naming service, add the hostname and IP address of each public-network adapter that you configured.

i. Reboot each node.

```
phys-schost# /usr/sbin/shutdown -y -g0 -i6
```

j. Verify that all adapters are configured and up.

```
phys-schost# ifconfig -a
```

4 On each node, enable the minimal RPC services that are necessary to enable the interactive `scinstall` utility.

When OpenSolaris software is installed, a restricted network profile is automatically configured. This profile is too restrictive for the cluster private network to function. To enable private-network functionality, run the following commands:

```
phys-schost# svccfg
svc:> select network/rpc/bind
svc:/network/rpc/bind> setprop config/local_only=false
svc:/network/rpc/bind> quit

phys-schost# svcadm refresh network/rpc/bind:default
phys-schost# svcprop network/rpc/bind:default | grep local_only
```

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

For more information about re-enabling network services, see "Planning Network Security" in *Solaris 10 5/08 Installation Guide: Planning for Installation and Upgrade*.

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

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

6 Type the option number for Create a New Cluster or Add a Cluster Node and press the Return key.

```
 *** Main Menu ***

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

      * 1) Create a new cluster or add a cluster node
      * 2) Print release information for this cluster node

      * ?) Help with menu options
      * q) Quit

    Option:  1
```

The New Cluster and Cluster Node Menu is displayed.

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

The Typical or Custom Mode menu is displayed.

8 Type the option number for either Typical or Custom and press the Return key.

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

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

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

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

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

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

11 From one node, verify that all nodes have joined the cluster.

```
phys-schost# /usr/cluster/bin/clnode status
```

Output resembles the following.

```
=== Cluster Nodes ===

--- Node Status ---

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

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

12 (Optional) Enable the automatic node reboot feature.

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

a. Enable automatic reboot.

```
phys-schost# /usr/cluster/bin/clnode set -p reboot_on_path_failure=enabled
```

`-p`	Specifies the property to set
`reboot_on_path_failure=enable`	Enables automatic node reboot if failure of all monitored disk paths occurs.

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

```
phys-schost# /usr/cluster/bin/clnode show
=== Cluster Nodes ===

Node Name:                                      node
```

```
...
  reboot_on_path_failure:                          enabled
...
```

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

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

```
exclude:lofs
```

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

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

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

See "The Loopback File System" in *System Administration Guide: Devices and File Systems* for more information about loopback file systems.

Example 3–1 Configuring Open HA Cluster Software on All Nodes

The following example shows the `scinstall` progress messages that are logged as `scinstall` completes configuration tasks on the two-node cluster, `schost`. The cluster is installed from `phys-schost-1` by using the `scinstall` utility in Typical Mode. The other cluster node is `phys-schost-2`. The adapter name is `e1000g0`. No `/globaldevices` partition exists, so the global-devices namespace is created on a `lofi` device. Automatic quorum-device selection is not used.

```
*** Create a New Cluster ***
Tue Apr 14 10:36:19 PDT 2009

    Attempting to contact "phys-schost-1" ...

    Searching for a remote configuration method ...

scrcmd -N phys-schost-1 test isfullyinstalled
The Sun Cluster framework software is installed.
scrcmd to "phys-schost-1" - return status 1.
```

```
rsh phys-schost-1 -n "/bin/sh -c '/bin/true; /bin/echo SC_COMMAND_STATUS=\$?'"
phys-schost-1: Connection refused
rsh to "phys-schost-1" failed.

ssh root@phys-schost-1 -o "BatchMode yes" -o "StrictHostKeyChecking yes"
-n "/bin/sh -c '/bin/true; /bin/echo SC_COMMAND_STATUS=\$?'"
No RSA host key is known for phys-schost-1 and you have requested strict checking.
Host key verification failed.
ssh to "phys-schost-1" failed.

    The Sun Cluster framework is able to complete the configuration
    process without remote shell access.


    Checking the status of service network/physical:nwam ...


/usr/cluster/lib/scadmin/lib/cmd_test isnwamenabled

scrcmd -N phys-schost-1 test isnwamenabled
    Plumbing network address 172.16.0.0 on adapter e1000g0 >> NOT DUPLICATE ... done
    Plumbing network address 172.16.0.0 on adapter e1000g0 >> NOT DUPLICATE ... done
    Testing for "/globaldevices" on "phys-schost-2" ...

/globaldevices is not a directory or file system mount point.
Cannot use "/globaldevices" on "phys-schost-2".


    Testing for "/globaldevices" on "phys-schost-1" ...

scrcmd -N phys-schost-1 chk_globaldev fs /globaldevices
/globaldevices is not a directory or file system mount point.


/globaldevices is not a directory or file system mount point.
Cannot use "/globaldevices" on "phys-schost-1".


scrcmd -N phys-schost-1 chk_globaldev lofi /.globaldevices 100m

---------------------------------
- Cluster Creation -
---------------------------------

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

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

    Cluster check report is displayed
...

scrcmd -N phys-schost-1 test isinstalling
"" is not running.

scrcmd -N phys-schost-1 test isconfigured
Sun Cluster is not configured.

    Configuring "phys-schost-1" ...

scrcmd -N phys-schost-1 install -logfile /var/cluster/logs/install/scinstall.log.2895
-k -C schost -F -G lofi -T node=phys-schost-2,node=phys-schost-1,authtype=sys
-w netaddr=172.16.0.0,netmask=255.255.240.0,maxnodes=64,maxprivatenets=10,
numvirtualclusters=12 -A trtype=dlpi,name=e1000g0 -B type=direct
ips_package_processing: ips_postinstall...
ips_package_processing: ips_postinstall done

Initializing cluster name to "schost" ... done
Initializing authentication options ... done
Initializing configuration for adapter "e1000g0" ... done
Initializing private network address options ... done

Plumbing network address 172.16.0.0 on adapter e1000g0 >> NOT DUPLICATE ... done

Setting the node ID for "phys-schost-1" ... done (id=1)

Verifying that NTP is configured ... done
Initializing NTP configuration ... done

Updating nsswitch.conf ... done

Adding cluster node entries to /etc/inet/hosts ... done


Configuring IP multipathing groups ...done


Verifying that power management is NOT configured ... done
Unconfiguring power management ... done
/etc/power.conf has been renamed to /etc/power.conf.041409104821
Power management is incompatible with the HA goals of the cluster.
Please do not attempt to re-configure power management.
```

```
Ensure network routing is disabled ... done
Network routing has been disabled on this node by creating /etc/notrouter.
Having a cluster node act as a router is not supported by Sun Cluster.
Please do not re-enable network routing.

Please reboot this machine.

Log file - /var/cluster/logs/install/scinstall.log.2895

scrcmd -N phys-schost-1 test hasbooted
This node has not yet been booted as a cluster node.
    Rebooting "phys-schost-1" ...
```

Troubleshooting **Unsuccessful configuration** – If one or more nodes cannot join the cluster, or if the wrong configuration information was specified, first attempt to rerun this procedure. If that does not correct the problem, perform the procedure "How to Uninstall Open HA Cluster Software" on page 53 on each misconfigured node to remove it from the cluster configuration. Then rerun this procedure.

Next Steps If you did not yet configure a quorum device in your cluster, go to "How to Configure Quorum Devices" on page 30.

Otherwise, go to "How to Verify the Quorum Configuration and Installation Mode" on page 34.

▼ How to Configure Quorum Devices

Note – If you chose automatic quorum configuration when you established the cluster, do not perform this procedure. Instead, proceed to "How to Verify the Quorum Configuration and Installation Mode" on page 34.

Perform this procedure one time only, after the new cluster is fully formed. Use this procedure to assign quorum votes and then to remove the cluster from installation mode.

Before You Begin If you intend to configure a quorum server as a quorum device, do the following:

- Install the Quorum Server software on the quorum server host machine and start the quorum server. For information about installing and starting the quorum server, see "How to Install and Configure Quorum Server Software" on page 15.

- Ensure that network switches that are directly connected to cluster nodes meet one of the following criteria:
  - The switch supports Rapid Spanning Tree Protocol (RSTP).
  - Fast port mode is enabled on the switch.

  One of these features is required to ensure immediate communication between cluster nodes and the quorum server. If this communication is significantly delayed by the switch, the cluster interprets this prevention of communication as loss of the quorum device.
- Have available the following information:
  - A name to assign to the configured quorum device
  - The IP address of the quorum server host machine
  - The port number of the quorum server

1 If you intend to use a quorum server and the public network uses variable-length subnetting, also called Classless Inter-Domain Subnetting (CIDS), on each node of the cluster modify netmask file entries for the public network.

If you use classful subnets, as defined in RFC 791, you do not need to perform this step.

a. Add to the `/etc/inet/netmasks` file an entry for each public subnet that the cluster uses.

The following is an example entry that contains a public-network IP address and netmask:

```
10.11.30.0    255.255.255.0
```

b. Append `netmask` + `broadcast` + to the hostname entry in each `/etc/hostname.`*adapter* file.

nodename **`netmask + broadcast +`**

2 On one node, become superuser.

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

3 To use a shared disk as a quorum device, verify device connectivity to the cluster nodes and choose the device to configure.

a. From one node of the cluster, display a list of all the devices that the system checks.

You do not need to be logged in as superuser to run this command.

```
phys-schost-1# /usr/cluster/bin/cldevice list -v
```

Output resembles the following:

```
DID Device          Full Device Path
----------          ----------------
```

```
d1                  phys-schost-1:/dev/rdsk/c0t0d0
d2                  phys-schost-1:/dev/rdsk/c0t6d0
d3                  phys-schost-2:/dev/rdsk/c1t1d0
d3                  phys-schost-1:/dev/rdsk/c1t1d0
...
```

b. Ensure that the output shows all connections between cluster nodes and storage devices.

c. Determine the global device-ID name of each shared disk that you are configuring as a quorum device.

Note – Any shared disk that you choose must be qualified for use as a quorum device.

Use the `scdidadm` output from Step a to identify the device–ID name of each shared disk that you are configuring as a quorum device. For example, the output in Step a shows that global device `d3` is shared by `phys-schost-1` and `phys-schost-2`.

4 To use a shared disk that does not support the SCSI protocol, ensure that fencing is disabled for that shared disk.

a. Display the fencing setting for the individual disk.

```
phys-schost# /usr/cluster/bin/cldevice show device

=== DID Device Instances ===
DID Device Name:                                    /dev/did/rdsk/dN
...
  default_fencing:                                   nofencing
...
```

- **If fencing for the disk is set to `nofencing` or `nofencing-noscrub`, fencing is disabled for that disk. Go to Step 5.**
- **If fencing for the disk is set to `pathcount` or `scsi`, disable fencing for the disk. Skip to Step c.**
- **If fencing for the disk is set to `global`, determine whether fencing is also disabled globally. Proceed to Step b.**

  Alternatively, you can simply disable fencing for the individual disk, which overrides for that disk whatever value the `global_fencing` property is set to. Skip to Step c to disable fencing for the individual disk.

b. Determine whether fencing is disabled globally.

```
phys-schost# /usr/cluster/bin/cluster show -t global
```

```
=== Cluster ===
Cluster name:                                   cluster
...
   global_fencing:                                nofencing
...
```

- **If global fencing is set to `nofencing` or `nofencing-noscrub`, fencing is disabled for the shared disk whose `default_fencing` property is set to `global`. Go to Step 5.**
- **If global fencing is set to `pathcount` or `prefer3`, disable fencing for the shared disk. Proceed to Step c.**

Note – If an individual disk has its `default_fencing` property set to `global`, the fencing for that individual disk is disabled only while the cluster-wide `global_fencing` property is set to `nofencing` or `nofencing-noscrub`. If the `global_fencing` property is changed to a value that enables fencing, then fencing becomes enabled for all disks whose `default_fencing` property is set to `global`.

---

c. Disable fencing for the shared disk.

```
phys-schost# /usr/cluster/bin/cldevice set \
-p default_fencing=nofencing-noscrub device
```

d. Verify that fencing for the shared disk is now disabled.

```
phys-schost# /usr/cluster/bin/cldevice show device
```

5 Start the `clsetup` utility.

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

The Initial Cluster Setup screen is displayed.

Note – If the Main Menu is displayed instead, initial cluster setup was already successfully performed. Skip to Step 10.

---

6 At the prompt `Do you want to add any quorum disks?`, type `Yes`.

7 Specify what type of device you want to configure as a quorum device.

Note – NAS devices are not a supported option for quorum devices in an Open HA Cluster 2009.06 configuration. Reference to NAS devices in the following table are for information only.

Quorum Device Type	Description
`shared_disk`	Sun NAS device or shared disk
`quorum_server`	Quorum server
`netapp_nas`	Network Appliance NAS device

8 Specify the name of the device to configure as a quorum device.

For a quorum server, also specify the following information:

- The IP address of the quorum server host
- The port number that is used by the quorum server to communicate with the cluster nodes

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

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

10 Quit the `clsetup` utility.

Next Steps Verify the quorum configuration and that installation mode is disabled. Go to "How to Verify the Quorum Configuration and Installation Mode" on page 34.

Troubleshooting **Interrupted clsetup processing** - If the quorum setup process is interrupted or fails to be completed successfully, rerun `clsetup`.

▼ How to Verify the Quorum Configuration and Installation Mode

Perform this procedure to verify that quorum configuration was completed successfully and that cluster installation mode is disabled.

You do not need to be superuser to run these commands.

1 From any node, verify the device and node quorum configurations.

```
phys-schost% /usr/cluster/bin/clquorum list
```

Output lists each quorum device, if used, membership type, and each node.

2 From any node, verify that cluster installation mode is disabled.

```
phys-schost% /usr/cluster/bin/cluster show -t global | grep installmode
  installmode:                                    disabled
```

Cluster installation and creation is complete.

Next Steps If you want to configure a failover ZFS file system that uses COMSTAR iSCSI storage, go to one of the following procedures:

- "How to Configure iSCSI Storage Using COMSTAR and Single Paths" on page 35
- "How to Configure iSCSI Storage Using COMSTAR and Multiple Paths" on page 40

Otherwise, if you want to use IP Security Architecture (IPsec) to provide secure TCP/IP communication on the cluster interconnect, go to "How to Configure IP Security Architecture (IPsec) on the Cluster Private Interconnect" on page 42.

Otherwise, configure the data services that you want to run on your cluster. Go to "Configuring Data Services" on page 44.

▼ How to Configure iSCSI Storage Using COMSTAR and Single Paths

Perform this procedure to configure OpenSolaris Common Multiprotocol SCSI TARget (COMSTAR) on locally attached storage, to share access among multiple cluster nodes. This procedure uses single paths between iSCSI initiators and iSCSI targets and also configures a mirrored ZFS storage pool to provide high availability.

Note – If you use multiple paths between iSCSI initiators and iSCSI targets, instead go to "How to Configure iSCSI Storage Using COMSTAR and Multiple Paths" on page 40.

Before You Begin Ensure that the storage configuration meets Open HA Cluster 2009.06 requirements. See "iSCSI Storage" on page 11.

1 On each node, perform the required procedures from Configuring an iSCSI Storage Array With COMSTAR (Task Map) that are listed in the following table, observing the Special Instructions.

Task	Documentation	Special Instructions
1. Perform basic setup.	Getting Started with COMSTAR	To create the SCSI logical unit, perform the procedure How to Create a Disk Partition SCSI Logical Unit. If you specify a whole disk instead of a slice to the `sbdadm create-lu` command, run the `cldevice clear` command afterwards to clear the DID namespace.

Task	Documentation	Special Instructions
2. Configure iSCSI target ports.	How to Configure iSCSI Target Ports	Create a target for each private-network adapter on each node.
3. Configure the iSCSI target.	How to Configure an iSCSI Target for Discovery	Use either static discovery or SendTargets. Do not use dynamic discovery.
4. Make a logical unit available.	How to Make Logical Units Available for iSCSI and iSER	

5. Configure an initiator system to access target storage.	How to Configure an iSCSI Initiator	■ Specify the node's `clprivnet` IP address as the target system. To determine the IP address of the `clprivnet` interface, run the following command. Output is similar to the following:
		`phys-schost# ifconfig clprivnet0`
		`clprivnet0:`
		`flags=1009843<UP,BROADCAST,RUNNING,MULTICAST,\`
		`MULTI_BCAST,PRIVATE,IPv4>`
		`mtu 1500 index 5`
		`inet 172.16.4.1 netmask fffffe00 broadcast \`
		`172.16.5.255`
		`ether 0:0:0:0:0:1`
		■ When completed, on each node update and populate the global-devices namespace.
		`phys-schost# scdidadm -r`
		`phys-schost# cldevice populate`

2 Disable fencing for each newly created device.

```
phys-schost# /usr/cluster/bin/cldevice set -p default_fencing=nofencing-noscrub device
```

Alternatively, disable fencing globally for all devices in the cluster. Do this if there are no shared devices in the cluster that are being used as a quorum device.

```
phys-schost# /usr/cluster/bin/cluster set -p global_fencing=nofencing-noscrub
```

3 List the DID mappings for the devices in the cluster.

Output is similar to the following, which shows a path from each node to each device:

```
phys-schost# /usr/cluster/bin/cldevice list -v
DID Device          Full Device Path
----------          ----------------
...
d3                  phys-schost-1:/dev/rdsk/c14t1d0s4
d3                  phys-schost-2:/dev/rdsk/c14t1d0s4
d4                  phys-schost-1:/dev/rdsk/c15t8d0s4
d4                  phys-schost-2:/dev/rdsk/c15t8d0s4
...
```

4 **From one node, create a mirrored ZFS storage pool from the DID devices that you created on each node.**

For the device path name, combine `/dev/did/dsk/`, the DID device name, and slice `s2`.

```
phys-schost# zpool create pool mirror /dev/did/dsk/dNs2 /dev/did/dsk/dYs2
```

5 **Configure the mirrored ZFS storage pool as an HAStoragePlus resource.**

```
phys-schost# /usr/cluster/bin/clresourcegroup resourcegroup
phys-schost# /usr/cluster/bin/clresourcetype register HAStoragePlus
phys-schost# /usr/cluster/bin/clresource create -g resourcegroup -t HAStoragePlus \
-p Zpools=pool resource
phys-schost# /usr/cluster/bin/clresourcegroup manage resourcegroup
phys-schost# /usr/cluster/bin/clresourcegroup online resourcegroup
```

Example 3–2 Configuring iSCSI Storage Using COMSTAR and Single Paths

This example shows the steps involved to configure COMSTAR based iSCSI storage and a mirrored ZFS storage pool, `zpool-1`. The locally attached disk for the node `phys-schost-1` is `/dev/rdsk/c1t0d0s4` and for `phys-schost-2` is `/dev/rdsk/c1t8d0s4`. The IP address of the `clprivnet0` interface is `172.16.4.1`.

Static discovery of the iSCSI target is configured. Procedures performed on `phys-schost-1` to configure an iSCSI initiator and target are also performed on `phys-schost-2`. After the `devfsadm` command attaches the disks as iSCSI targets, `/dev/rdsk/c1t0d0s4` becomes `/dev/rdsk/c14t0d0s4` on the initiator side and `/dev/rdsk/c1t8d0s4` becomes `/dev/rdsk/c15t8d0s4`.

The cluster does not use any shared disks, so fencing is turned off globally for all disks in the cluster. The resource group `rg-1` is configured with HAStoragePlus resource `hasp-rs` the mirrored ZFS storage pool `zpool-1`.

Enable and verify the STMF service

```
phys-schost-1# svcadm enable stmf
phys-schost-1# svcs stmf
online    15:59:53 svc:/system/stmf:default
```

Repeat on phys-schost-2

Create and verify disk-partition SCSI logical units on each node

```
phys-schost-1# sbdadm create-lu /dev/rdsk/c1t0d0s4
Created the following LU:

              GUID                    DATA SIZE           SOURCE
--------------------------------  -------------------  ----------------
600144f05b4c460000004a1d9dd00001      73407800320      /dev/rdsk/c1t0d0s4
root@phys-schost-1:#
-------------------------
```

```
phys-schost-2# sbdadm create-lu /dev/rdsk/c1t8d0s4
Created the following LU:

              GUID                    DATA SIZE           SOURCE
--------------------------------  -------------------  -----------------
600144f07d15cd0000004a202e340001      73407800320      /dev/rdsk/c1t8d0s4
root@phys-schost-2:#
-------------------------
```

Enable the iSCSI target SMF service

```
phys-schost-1# svcadm enable -r svc:/network/iscsi/target:default
phys-schost-1# svcs -a | grep iscsi
online  14:21:25 svc:/network/iscsi/target:default
```

Repeat on phys-schost-2

Configure each iSCSI target for static discovery

```
phys-schost-1# itadm create-target
Target: iqn.1986-03.com.sun:02:97c1caa8-5732-ec53-b7a2-a722a946fead
successfully created
phys-schost-1# itadm list-target
TARGET NAME                                                  STATE   SESSIONS
iqn.1986-03.com.sun:02:97c1caa8-5732-ec53-b7a2-a722a946fead  online  0
```

Repeat on phys-schost-2 for the other iSCSI target

Make the logical units available

```
phys-schost-1# sbdadm list-lu
phys-schost-1# stmfadm add-view 600144f05b4c460000004a1d9dd00001
```

Repeat on phys-schost-2 for the other logical unit's GUID

Configure iSCSI initiators to access target storage

```
phys-schost-1# iscsiadm modify discovery --static enable
phys-schost-1# iscsiadm list discovery
Discovery:
Static: enabled
Send Targets: disabled
iSNS: disabled
phys-schost-1# ifconfig clprivnet0
    clprivnet0:
...
    inet 172.16.4.1 netmask fffffe00 broadcast 172.16.5.255
...
phys-schost-1# iscsiadm add static-config \
iqn.1986-03.com.sun:02:97c1caa8-5732-ec53-b7a2-a722a946fead,172.16.4.1
phys-schost-1# iscsiadm list static-config
Static Configuration Target:
iqn.1986-03.com.sun:02:97c1caa8-5732-ec53-b7a2-a722a946fead,172.16.4.1:3260
phys-schost-1# devfsadm -i iscsi
phys-schost-1# format -e
```

```
phys-schost-1# iscsiadm list target
Target: iqn.1986-03.com.sun:02:97c1caa8-5732-ec53-b7a2-a722a946fead
        Alias: -
        TPGT: 1
        ISID: 4000002a0000
        Connections: 1
```

Repeat on phys-schost-2 for this target

Repeat on both nodes for the other target

Update and populate the global-devices namespace on each node

```
phys-schost-1# scdidadm -r
phys-schost-1# cldevice populate
```

Repeat on phys-schost-2

Disable fencing for all disks in the cluster

```
phys-schost-1# /usr/cluster/bin/cluster set -p global_fencing=nofencing-noscrub
```

Create a mirrored ZFS storage pool

```
phys-schost-1/usr/cluster/bin/cldevice list -v
DID Device          Full Device Path
----------          ----------------
...
d3                  phys-schost-1:/dev/rdsk/c14t0d0s4
d3                  phys-schost-2:/dev/rdsk/c14t0d0s4
d4                  phys-schost-1:/dev/rdsk/c15t8d0s4
d4                  phys-schost-2:/dev/rdsk/c15t8d0s4
...
phys-schost-1# zpool create zpool-1 mirror /dev/did/dsk/d3s2 /dev/did/dsk/d4s2
```

Configure the mirrored ZFS storage pool as an HAStoragePlus resource

```
phys-schost# /usr/cluster/bin/clresourcegroup rg-1
phys-schost# /usr/cluster/bin/clresourcetype register HAStoragePlus
phys-schost# /usr/cluster/bin/clresource create -g rg-1 -t HAStoragePlus \
-p Zpools=zpool-1 hasp-rs
phys-schost# /usr/cluster/bin/clresourcegroup manage rg-1
phys-schost# /usr/cluster/bin/clresourcegroup online rg-1
```

Next Steps If you want to use IP Security Architecture (IPsec) to provide secure TCP/IP communication on the cluster interconnect, go to "How to Configure IP Security Architecture (IPsec) on the Cluster Private Interconnect" on page 42.

Otherwise, configure the data services that you want to run on your cluster. Go to "Configuring Data Services" on page 44.

▼ How to Configure iSCSI Storage Using COMSTAR and Multiple Paths

Perform this procedure to configure OpenSolaris Common Multiprotocol SCSI TARget (COMSTAR) on locally attached storage, to share access among multiple cluster nodes. This procedure uses multiple paths between iSCSI initiators and iSCSI targets and also configures a mirrored ZFS storage pool to provide high availability. This procedure optionally includes configuring the I/O multipathing feature (MPxIO).

Note – If you use single paths between iSCSI initiators and iSCSI targets, go instead to "How to Configure iSCSI Storage Using COMSTAR and Single Paths" on page 35.

Before You Begin Ensure that the storage configuration meets Open HA Cluster 2009.06 requirements. See "iSCSI Storage" on page 11.

1 (Optional) If you intend to use I/O multipathing (MPxIO), on each node ensure that the I/O multipathing feature is enabled for iSCSI.

The feature is enabled when the `mpxio-disable` property is set to `no`.

```
phys-schost# cat /kernel/drv/iscsi.conf
...
mpxio-disable="no";
```

For more information about I/O multipathing, see *Solaris Fibre Channel Storage Configuration and Multipathing Support Guide*.

2 Determine the IP address of each adapter that is used for the private interconnect.

You will specify these addresses later when you create iSCSI target ports. Output is similar to the following:

```
phys-schost# /usr/cluster/bin/clinterconnect status
=== Cluster Transport Paths ===

Endpoint1               Endpoint2               Status
---------               ---------               ------
phys-schost-1:adapter1    phys-schost-2:adapter1    Path online
phys-schost-1:adapter2    phys-schost-2:adapter2    Path online

phys-schost# ifconfig adapter1
nge1: flags=201008843<UP,BROADCAST,RUNNING,MULTICAST,PRIVATE,IPv4,CoS> mtu
1500 index 3
        inet 172.16.1.1 netmask ffffff80 broadcast 172.16.1.127
        ether 0:14:4f:8d:9b:3
phys-schost# ifconfig adapter2
e1000g1: flags=201008843<UP,BROADCAST,RUNNING,MULTICAST,PRIVATE,IPv4,CoS>
```

```
mtu 1500 index 4
        inet 172.16.0.129 netmask ffffff80 broadcast 172.16.0.255
        ether 0:15:17:35:9b:a1
```

3 On each node, perform the procedures that are listed in Configuring an iSCSI Storage Array With COMSTAR (Task Map).

Observe the following additional instructions when you configure a COMSTAR iSCSI target in an Open HA Cluster 2009.06 configuration:

Task	Documentation	Special Instructions
1. Perform basic setup.	Getting Started with COMSTAR	To create the SCSI logical unit, perform the procedure How to Create a Disk Partition SCSI Logical Unit. If you specify a whole disk instead of a slice to the `sbdadm create-lu` command, run the `cldevice clear` command afterwards to clear the DID namespace.
2. Configure iSCSI target ports.	How to Configure iSCSI Target Ports	Create a target for each private-network adapter on each node.
3. Configure the iSCSI target.	How to Configure an iSCSI Target for Discovery	Use either static discovery or SendTargets. Do not use dynamic discovery.
4. Make a logical unit available.	How to Make Logical Units Available for iSCSI and iSER	

5. Configure an initiator system to access target storage.	How to Configure an iSCSI Initiator	■ Specify the node's `clprivnet` IP address as the target system. To determine the IP address of the `clprivnet` interface, run the following command. Output is similar to the following:
		`phys-schost# ifconfig clprivnet0`
		`clprivnet0:`
		`flags=1009843<UP,BROADCAST,RUNNING,MULTICAST,\`
		`MULTI_BCAST,PRIVATE,IPv4>`
		`mtu 1500 index 5`
		`inet 172.16.4.1 netmask fffffe00 broadcast \`
		`172.16.5.255`
		`ether 0:0:0:0:0:1`
		■ When completed, on each node update and populate the global-devices namespace.
		`phys-schost# scdidadm -r`
		`phys-schost# cldevice populate`

4 Disable fencing for each of the newly created devices.

```
phys-schost# /usr/cluster/bin/cldevice set -p default_fencing=nofencing-noscrub device
```

5 **From one node, create a mirrored ZFS storage pool from the DID devices that you created on each node.**

```
phys-schost# zpool create pool mirror /dev/did/dsk/dNsX /dev/did/dsk/dYsX
```

6 **From one node, configure the mirrored ZFS storage pool as an HAStoragePlus resource.**

```
phys-schost# /usr/cluster/bin/clresourcegroup resourcegroup
phys-schost# /usr/cluster/bin/clresourcetype register HAStoragePlus
phys-schost# /usr/cluster/bin/clresource create -g resourcegroup -t HAStoragePlus  \
-p Zpools=pool resource
phys-schost# /usr/cluster/bin/clresourcegroup manage resourcegroup
phys-schost# /usr/cluster/bin/clresourcegroup online resourcegroup
```

Next Steps If you want to use IP Security Architecture (IPsec) to provide secure TCP/IP communication on the cluster interconnect, go to "How to Configure IP Security Architecture (IPsec) on the Cluster Private Interconnect" on page 42.

Otherwise, configure the data services that you want to run on your cluster. Go to "Configuring Data Services" on page 44.

▼ How to Configure IP Security Architecture (IPsec) on the Cluster Private Interconnect

You can configure IP Security Architecture (IPsec) for the private-interconnect interface to provide secure TCP/IP communication on the cluster interconnect.

For information about IPsec, see Part IV, "IP Security," in *System Administration Guide: IP Services* and the `ipsecconf(1M)` man page. For information about the `clprivnet` interface, see the `clprivnet(7)` man page.

Perform this procedure on each cluster node that you want to configure to use IPsec.

1 **Become superuser.**

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

2 **On each node, determine the IP address of the `clprivnet` interface.**

```
phys-schost# ifconfig clprivnet0
```

3 If you use virtual NICs (VNICs) to route private interconnect communication over the public network, also determine the IP address of the physical interfaces that the VNICs use.

a. Display the status of all transport paths in the cluster and the physical interfaces that are used.

Output is similar to the following:

```
phys-schost# /usr/cluster/bin/clinterconnect status
-- Cluster Transport Paths --

                    Endpoint               Endpoint               Status
                    --------               --------               ------
  Transport path:   phys-schost-1:adapter1 phys-schost-2:adapter1 Path online
  Transport path:   phys-schost-1:adapter2 phys-schost-2:adapter2 Path online
```

b. Identify the IP address of each interface that is used on each node.

```
phys-schost-1# ifconfig adapter
phys-schost-2# ifconfig adapter
```

4 On each node, configure the `/etc/inet/ipsecinit.conf` policy file and add Security Associations (SAs) between each pair of private-interconnect IP addresses that you want to use IPsec.

Follow the instructions in "How to Secure Traffic Between Two Systems With IPsec" in *System Administration Guide: IP Services*. In addition, observe the following guidelines:

- Ensure that the values of the configuration parameters for these addresses are consistent on all the partner nodes.
- Configure each policy as a separate line in the configuration file.
- To implement IPsec without rebooting, follow the instructions in the procedure's example, "Securing Traffic With IPsec Without Rebooting."

For more information about the `sa unique` policy, see the `ipsecconf(1M)` man page.

a. In each file, add one entry for each `clprivnet` IP address in the cluster to use IPsec.

Include the `clprivnet` private-interconnect IP address of the local node.

b. If you use VNICs, also add one entry for the IP address of each physical interface that is used by the VNICs.

c. (Optional) To enable striping of data over all links, include the `sa unique` policy in the entry.

This feature helps the driver to optimally utilize the bandwidth of the cluster private network, which provides a high granularity of distribution and better throughput. The private-interconnect interface uses the Security Parameter Index (SPI) of the packet to stripe the traffic.

5 **On each node, edit the `/etc/inet/ike/config` file to set the `p2_idletime_secs` parameter.**

Add this entry to the policy rules that are configured for cluster transports. This setting provides the time for security associations to be regenerated when a cluster node reboots, and limits how quickly a rebooted node can rejoin the cluster. A value of 30 seconds should be adequate.

```
phys-schost# vi /etc/inet/ike/config
...
{
    label "clust-priv-interconnect1-clust-priv-interconnect2"
...
p2_idletime_secs 30
}
...
```

Next Steps Configure the data services that you want to run on your cluster. Go to "Configuring Data Services" on page 44.

Configuring Data Services

This section provides information to configure data services that are supported with Open HA Cluster 2009.06 software.

The following table lists the location of information to install and configure each supported data service. Use these procedures to configure data services for the Open HA Cluster 2009.06 release, except for the following changes:

- Install application software as described by the application's installation instructions for OpenSolaris environments.
- Install the data-service agent by following instructions in "How to Prepare to Download Open HA Cluster Software" on page 14 and "How to Install Open HA Cluster 2009.06 Software" on page 17.

Data Service	Documentation
Data Service for Apache	*Sun Cluster Data Service for Apache Guide for Solaris OS*
Data Service for Apache Tomcat	*Sun Cluster Data Service for Apache Tomcat Guide for Solaris OS*
Data Service for DHCP	*Sun Cluster Data Service for DHCP Guide for Solaris OS*
Data Service for DNS	*Sun Cluster Data Service for DNS Guide for Solaris OS*
Data Service for Glassfish	*Sun Cluster Data Service for Sun Java System Application Server Guide for Solaris OS*

Data Service	Documentation
Data Service for Kerberos	*Sun Cluster Data Service for Kerberos Guide for Solaris OS*
Data Service for MySQL	*Sun Cluster Data Service for MySQL Guide for Solaris OS*
Data Service for NFS	*Sun Cluster Data Service for NFS Guide for Solaris OS*
Data Service for Samba	*Sun Cluster Data Service for Samba Guide for Solaris OS*
Data Service for Solaris Containers	"How to Configure the HA-Containers Zone Boot Component for ipkg Brand Zones" on page 45 *Sun Cluster Data Service for Solaris Containers Guide for Solaris OS*

How to Configure the HA-Containers Zone Boot Component for ipkg Brand Zones

Perform this procedure to configure the zone boot component (`sczbt`) of the Solaris Containers data service to use `ipkg` brand non-global zones. Use this procedure instead of the instructions for `sczbt` that are in *Sun Cluster Data Service for Solaris Containers Guide for Solaris OS*. All other procedures in the Solaris Containers data-service manual are valid for an Open HA Cluster 2009.06 configuration.

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

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

2. **Create a resource group.**

   ```
   phys-schost-1# /usr/cluster/bin/clresourcegroup create resourcegroup
   ```

3. **Create a mirrored ZFS storage pool to be used for the HA zone root path.**

   ```
   phys-schost-1# zpool create -m mountpoint pool mirror /dev/rdsk/cNtXdY \
   /dev/rdsk/cNtXdZ
   phys-schost# zpool export pool
   ```

4. **Register the HAStoragePlus resource type.**

   ```
   phys-schost-1# /usr/cluster/bin/clresourcetype register SUNW.HAStoragePlus
   ```

5. **Create an HAStoragePlus resource.**

   Specify the ZFS storage pool and the resource group that you created.

   ```
   phys-schost-1# /usr/cluster/bin/clresource create -t SUNW.HAStoragePlus \
   -g resourcegroup -p Zpools=pool hasp-resource
   ```

6 Bring the resource group online.

```
phys-schost-1# clresourcegroup online -eM resourcegroup
```

7 Create a ZFS file-system dataset on the ZFS storage pool that you created.

You will use this file system as the zone root path for the `ipkg` brand zone that you create later in this procedure.

```
phys-schost-1# zfs create pool/filesystem
```

8 Ensure that the universally unique ID (UUID) of each node's boot-environment (BE) root dataset is the same value.

a. Determine the UUID of the node where you initially created the zone.

Output is similar to the following.

```
phys-schost-1# beadm list -H
...
b101b-SC;8fe53702-16c3-eb21-ed85-d19af92c6bbd;NR;/;756...
```

In this example output, the UUID is `8fe53702-16c3-eb21-ed85-d19af92c6bbd` and the BE is `b101b-SC`.

b. Set the same UUID on the second node.

```
phys-schost-2# zfs set org.opensolaris.libbe:uuid=uuid rpool/ROOT/BE
```

9 On both nodes, configure the `ipkg` brand non-global zone.

Set the zone root path to the file system that you created on the ZFS storage pool.

```
phys-schost# zonecfg -z zonename  \
'create ; set zonepath=/pool/filesystem/zonename ; set autoboot=false'

phys-schost# zoneadm list -cv
  ID NAME          STATUS        PATH                          BRAND    IP
   0 global        running       /                             native   shared
   - zonename      configured    /pool/filesystem/zonename     ipkg     shared
```

10 From the node that masters the HAStoragePlus resource, install the `ipkg` brand non-global zone.

Output is similar to the following:

a. Determine which node masters the HAStoragePlus resource.

```
phys-schost# /usr/cluster/bin/clresource status
=== Cluster Resources ===

Resource Name             Node Name       Status        Message
--------------            ----------      -------       -------
 hasp-resource            phys-schost-1   Online        Online
```

```
                                phys-schost-2   Offline       Offline
```

Perform the remaining tasks in this step from the node that masters the HAStoragePlus resource.

b. Install the zone on the node that masters the HAStoragePlus resource for the ZFS storage pool.

```
phys-schost-1# zoneadm -z zonename install
```

c. Verify that the zone is installed.

```
phys-schost-1# zoneadm list -cv
  ID NAME           STATUS       PATH                           BRAND    IP
   0 global         running      /                              native   shared
   - zonename       installed    /pool/filesystem/zonename      ipkg     shared
```

d. Boot the zone that you created and verify that the zone is running.

```
phys-schost-1# zoneadm -z zonename boot
phys-schost-1# zoneadm list -cv
  ID NAME           STATUS       PATH                           BRAND    IP
   0 global         running      /                              native   shared
   - zonename       running      /pool/filesystem/zonename      ipkg     shared
```

e. Open a new terminal window and log in to the zone.

f. Halt the zone.

The zone's status should return to `installed`.

```
phys-schost-1# zoneadm -z zonename halt
```

11 Switch the resource group to the other node and forcibly attach the zone.

a. Switch over the resource group.

Output is similar to the following, where `phys-schost-1` is the node that currently masters the resource group and `phys-schost-2` is the node to which you switch the resource group.

```
phys-schost-1# /usr/cluster/bin/clresourcegroup switch -n phys-schost-2 resourcegroup
```

Perform the remaining tasks in this step from the node to which you switch the resource group.

b. Forcibly attach the zone to the node to which you switched the resource group.

```
phys-schost-2# zoneadm -z zonename attach -F
```

c. Verify that the zone is installed on the node.

Output is similar to the following:

```
phys-schost-2# zoneadm list -cv
  ID NAME           STATUS       PATH                           BRAND    IP
   0 global         running      /                              native   shared
   - zonename       installed    /pool/filesystem/zonename      ipkg     shared
```

d. Boot the zone.

```
phys-schost-2# zoneadm -z zonename boot
```

e. Open a new terminal window and log in to the zone.

Perform this step to verify that the zone is functional.

```
phys-schost-2# zlogin -C zonename
```

f. Halt the zone.

```
phys-schost-2# zoneadm -z zonename halt
```

12 From one node, configure the zone-boot (`sczbt`) resource.

a. Register the `SUNW.gds` resource type.

```
phys-schost-1# /usr/cluster/bin/clresourcetype register SUNW.gds
```

b. Create a directory on the ZFS file system that you created.

You will specify this directory to store the parameter values that you set for the zone-boot resource.

```
phys-schost-1# mkdir /pool/filesystem/parameterdir
```

c. Install and configure the HA-Containers agent.

```
phys-schost# pkg install SUNWsczone
phys-schost# cd /opt/SUNWsczone/sczbt/util
phys-schost# cp -p sczbt_config sczbt_config.zoneboot-resource
phys-schost# vi sczbt_config.zoneboot-resource
      Add or modify the following entries in the file.
RS="zoneboot-resource"
RG="resourcegroup"
PARAMETERDIR="/pool/filesystem/parameterdir"
SC_NETWORK="false"
SC_LH=""
FAILOVER="true"
HAS_RS="hasp-resource"

Zonename="zonename"
Zonebrand="ipkg"
Zonebootopt=""
```

```
Milestone="multi-user-server"
LXrunlevel="3"
SLrunlevel="3"
Mounts=""
```

Save and exit the file.

d. Configure the zone-boot resource.

The resource is configured with the parameters that you set in the zone-boot configuration file.

```
phys-schost-1# ./sczbt_register -f ./sczbt_config.zoneboot-resource
```

e. Verify that the zone-boot resource is enabled.

```
phys-schost-1# /usr/cluster/bin/clresource enable zoneboot-resource
```

13 Verify that the resource group can switch to another node and the ZFS storage pool successfully starts there after the switchover.

a. Switch the resource group to another node.

```
phys-schost-2# /usr/cluster/bin/clresourcegroup switch -n phys-schost-1 resourcegroup
```

b. Verify that the resource group is now online on the new node.

Output is similar to the following:

```
phys-schost-1# /usr/cluster/bin/clresourcegroup status
=== Cluster Resource Groups ===

    Group Name                  Node Name          Suspended        Status
    ----------                  ---------          ---------        ------
    resourcegroup               phys-schost-1      No               Online
                                phys-schost-2      No               Offline
```

c. Verify that the zone is running on the new node.

```
phys-schost-1# zoneadm list -cv
ID  NAME     STATUS       PATH                           BRAND      IP
 0  global   running      /                              native     shared
 1  zonename running      /pool/filesystem/zonename      ipkg       shared
```

Example 3–3 Configuring the HA-Containers Zone Boot Component for ipkg Brand Zones

This example creates the HAStoragePlus resource `hasp-rs`, which uses a mirrored ZFS storage pool `hapool` in the resource group `zone-rg`. The storage pool is mounted on the `/hapool/ipkg` file system. The `hasp-rs` resource runs on the `ipkg` brand non-global zone `ipkgzone1`, which is configured on both `phys-schost-1` and `phys-schost-2`. The zone-boot resource `ipkgzone1-rs` is based on the `SUNW.gds` resource type.

Create a resource group.

```
phys-schost-1# /usr/cluster/bin/clresourcegroup create zone-rg
```

Create a mirrored ZFS storage pool to be used for the HA zone root path.

```
phys-schost-1# zpool create -m /ha-zones hapool mirror /dev/rdsk/c4t6d0 \
/dev/rdsk/c5t6d0
phys-schost# zpool export hapool
```

Create an HAStoragePlus resource that uses the resource group and mirrored ZFS storage pool that you created.

```
phys-schost-1# /usr/cluster/bin/clresourcetype register SUNW.HAStoragePlus
phys-schost-1# /usr/cluster/bin/clresource create -t SUNW.HAStoragePlus \
-g zone-rg -p Zpools=hapool hasp-rs
```

Bring the resource group online.

```
phys-schost-1# clresourcegroup online -eM zone-rg
```

Create a ZFS file-system dataset on the ZFS storage pool that you created.

```
phys-schost-1# zfs create hapool/ipkg
```

Ensure that the universally unique ID (UUID) of each node's boot-environment (BE) root dataset is the same value on both nodes.

```
phys-schost-1# beadm list -H
...
zfsbe;8fe53702-16c3-eb21-ed85-d19af92c6bbd;NR;/;7565844992;static;1229439064
...
phys-schost-2# zfs set org.opensolaris.libbe:uuid=8fe53702-16c3-eb21-ed85-d19af92c6bbd rpool/ROOT/zfsbe
```

Configure the ipkg brand non-global zone.

```
phys-schost-1# zonecfg -z ipkgzone1 'create ; \
set zonepath=/hapool/ipkg/ipkgzone1 ; set autoboot=false'
phys-schost-1# zoneadm list -cv
  ID NAME             STATUS       PATH                           BRAND    IP
   0 global           running      /                              native   shared
   - ipkgzone1        configured   /hapool/ipkg/ipkgzone1         ipkg     shared
```

Repeat on phys-schost-2.

Identify the node that masters the HAStoragePlus resource, and from that node install ipkgzone1.

```
phys-schost-1# /usr/cluster/bin/clresource status
=== Cluster Resources ===

Resource Name             Node Name       Status        Message
--------------            ----------      -------        -------
 hasp-rs                  phys-schost-1   Online        Online
                          phys-schost-2   Offline       Offline
phys-schost-1# zoneadm -z ipkgzone1 install
phys-schost-1# zoneadm list -cv
  ID NAME             STATUS     PATH                           BRAND    IP
   0 global           running    /                              native   shared
```

```
   - ipkgzone1        installed  /hapool/ipkg/ipkgzone1         ipkg     shared
phys-schost-1# zoneadm -z ipkgzone1 boot
phys-schost-1# zoneadm list -cv
  ID NAME             STATUS     PATH                           BRAND    IP
   0 global           running    /                              native   shared
   - ipkgzone1        running    /hapool/ipkg/ipkgzone1         ipkg     shared
```

Open a new terminal window and log in to ipkgzone1.

```
phys-schost-1# zoneadm -z ipkgzone1 halt
```

Switch zone-rg to phys-schost-2 and forcibly attach the zone.

```
phys-schost-1# /usr/cluster/bin/clresourcegroup switch -n phys-schost-2 zone-rg
phys-schost-2# zoneadm -z ipkgzone1 attach -F
phys-schost-2# zoneadm list -cv
  ID NAME             STATUS      PATH                           BRAND    IP
   0 global           running     /                              native   shared
   - ipkgzone1        installed   /hapool/ipkg/ipkgzone1         ipkg     shared
phys-schost-2# zoneadm -z ipkgzone1 boot
```

Open a new terminal window and log in to ipkgzone1.

```
phys-schost-2# zlogin -C ipkgzone1
phys-schost-2# zoneadm -z ipkgzone1 halt
```

From one node, configure the zone-boot (sczbt) resource.

```
phys-schost-1# /usr/cluster/bin/clresourcetype register SUNW.gds
phys-schost-1# mkdir /hapool/ipkg/params
```

Install and configure the HA-Containers agent.

```
phys-schost# pkg install SUNWsczone
phys-schost# cd /opt/SUNWsczone/sczbt/util
phys-schost# cp -p sczbt_config sczbt_config.ipkgzone1-rs
phys-schost# vi sczbt_config.ipkgzone1-rs
```

Add or modify the following entries in the sczbt_config.ipkgzone1-rs file.

```
RS="ipkgzone1-rs"
RG="zone-rg"
PARAMETERDIR="/hapool/ipkg/params"
SC_NETWORK="false"
SC_LH=""
FAILOVER="true"
HAS_RS="hasp-rs"

Zonename="ipkgzone1"
Zonebrand="ipkg"
Zonebootopt=""
Milestone="multi-user-server"
LXrunlevel="3"
SLrunlevel="3"
```

```
Mounts=""
```

Save and exit the file.

Configure the ipkgzone1-rs resource.

```
phys-schost-1# ./sczbt_register -f ./sczbt_config.ipkgzone1-rs
phys-schost-1# /usr/cluster/bin/clresource enable ipkgzone1-rs
```

Verify that zone-rg can switch to another node and that ipkgzone1 successfully starts there after the switchover.

```
phys-schost-2# /usr/cluster/bin/clresourcegroup switch -n phys-schost-1 zone-rg
phys-schost-1# /usr/cluster/bin/clresourcegroup status
=== Cluster Resource Groups ===

    Group Name                  Node Name          Suspended        Status
    ----------                  ---------          ---------        ------
    zone-rg                     phys-schost-1      No               Online
                                phys-schost-2      No               Offline
phys-schost-1# zoneadm list -cv
ID  NAME      STATUS       PATH                           BRAND      IP
 0  global    running      /                              native     shared
 1  ipkgzone1 running      /hapool/ipkg/ipkgzone1         ipkg       shared
```

◆ ◆ ◆ CHAPTER 4

Uninstalling Software From the Cluster

This chapter provides procedures for uninstalling an Open HA Cluster configuration. The following procedures are in this chapter:

Uninstalling the Software

This section provides procedures to uninstall Open HA Cluster 2009.06 software from a cluster.

▼ How to Uninstall Open HA Cluster Software

Perform this procedure to uninstall Open HA Cluster 2009.06 software from the node. If the node is a configured member of a cluster, this procedure also removes the node from the cluster configuration.

Note – Do not use the `pkg uninstall` command to remove Open HA Cluster 2009.06 software from cluster nodes. Use only the `scinstall -r` command to ensure that all cluster packages, including any that were installed by explicit command, and any cluster configuration information is completely removed from the node.

---

1 **Add to the cluster's node-authentication list each node that you intend to unconfigure.**

If you are removing software from a node that is not a configured member of a cluster, skip to Step 2.

a. **On an active cluster member *other than* the node that you are unconfiguring, become superuser.**

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

b. **Specify the name of the node to add to the authentication list.**

- To add a single node, use the following command:

```
phys-schost# /usr/cluster/bin/claccess allow -h nodename
```

`-h` *nodename* Specifies the name of the node to add to the authentication list.

- To add all nodes, use the following command:

```
phys-schost# /usr/cluster/bin/claccess allow-all
```

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

You can alternatively use the `clsetup` utility to perform this task. See "How to Add a Node to the Authorized Node List" in *Sun Cluster System Administration Guide for Solaris OS* for procedures.

2 **On the node that you intend to unconfigure, become superuser.**

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

3 **Shut down the node.**

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

4 **Reboot the node into noncluster mode.**

- **On SPARC based systems, perform the following command:**

```
ok boot -x
```

- **On x86 based systems, perform the following commands:**

  a. **In the GRUB menu, use the arrow keys to select the appropriate Solaris entry and type e to edit its commands.**

     The GRUB menu appears similar to the following:

```
GNU GRUB version 0.97 (639K lower / 1047488K upper memory)
+-------------------------------------------------------------------------+
| OpenSolaris 2009.06                                                     |
| OHAC-2009-06                                                            |
|                                                                         |
+-------------------------------------------------------------------------+
Use the ^ and v keys to select which entry is highlighted.
Press enter to boot the selected OS, 'e' to edit the
commands before booting, or 'c' for a command-line.
```

     For more information about GRUB based booting, see "Booting an x86 Based System by Using GRUB (Task Map)" in *System Administration Guide: Basic Administration*.

  b. **In the boot parameters screen, use the arrow keys to select the `kernel` entry and type e to edit the entry.**

     The GRUB boot parameters screen appears similar to the following:

```
GNU GRUB version 0.97 (639K lower / 1047488K upper memory)
+-------------------------------------------------------------------------+
|...                                                                      |
| kernel /platform/i86pc/kernel/$ISADIR/unix -B $ZFS-BOOTFS,console=gr>|
|...                                                                      |
+-------------------------------------------------------------------------+
Use the ^ and v keys to select which entry is highlighted.
Press 'b' to boot, 'e' to edit the selected command in the
boot sequence, 'c' for a command-line, 'o' to open a new line
after ('O' for before) the selected line, 'd' to remove the
selected line, or escape to go back to the main menu.
```

  c. **Add `-x` to the command to specify that the system boot into noncluster mode.**

```
[ Minimal BASH-like line editing is supported. For the first word, TAB
lists possible command completions. Anywhere else TAB lists the possible
completions of a device/filename. ESC at any time exits. ]

grub edit> kernel /platform/i86pc/kernel/$ISADIR/unix \
-B $ZFS-BOOTFS,console=graphics -x
```

  d. **Press Enter to accept the change and return to the boot parameters screen.**

     The screen displays the edited command.

  e. **Type `b` to boot the node into noncluster mode.**

Note – This change to the kernel boot parameter command does not persist over the system boot. The next time you reboot the node, it will boot into cluster mode. To boot into noncluster mode instead, perform these steps to again add the `-x` option to the kernel boot parameter command.

5 **Change to a directory, such as the root (`/`) directory, that does not contain any files that are delivered by the Open HA Cluster packages.**

```
phys-schost# cd /
```

6 **Remove the node from the cluster configuration.**

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

The node is removed from the cluster configuration and Open HA Cluster software is removed from the node. For more information, see the `scinstall(1M)` man page.

7 **Repeat Step 2 through Step 6 for any additional node to unconfigure.**

Next Steps To physically remove the node from the cluster, see "How to Remove an Interconnect Component" in *Sun Cluster 3.1 - 3.2 Hardware Administration Manual for Solaris OS* and the removal procedure in the appropriate Sun Cluster Hardware Administration Collection manual for your storage array.

▼ How to Remove Quorum Server Software

Before You Begin Before you uninstall Quorum Server software, ensure that you have completed the following tasks:

- On each cluster that uses the quorum server, remove the quorum server as a quorum device. Perform the steps in "How to Remove a Quorum Device" in *Sun Cluster System Administration Guide for Solaris OS*.

  In normal operation, this step also removes the quorum server information on the quorum-server host. If communications between the cluster and the quorum -server host computer are down during this step, you must clean up the invalid quorum server information on the quorum-server host computer. Perform the steps in "Cleaning Up Stale Quorum Server Cluster Information" in *Sun Cluster System Administration Guide for Solaris OS*.
- On each quorum-server host computer, stop the quorum server by following the steps in "How to Stop a Quorum Server" in *Sun Cluster System Administration Guide for Solaris OS*.

1. **Become superuser on the quorum server host computer to uninstall.**

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

2. **Uninstall the quorum-server software.**

   ```
   quorumserver# /usr/bin/pkg uninstall ha-cluster-quorum-server-full
   ```

3. **(Optional) Clean up or remove the quorum server directories.**

   By default, this directory is `/var/scqsd`.

Index

Z