Oracle® Solaris Cluster Upgrade Guide

Part No: E24617
March 2012, E24617–02

ORACLE

120305@25097

Contents

Preface

The *Oracle Solaris Cluster Upgrade Guide* contains guidelines and procedures for upgrading the Oracle Solaris Cluster software on both SPARC based systems and x86 based systems.

Note – This Oracle Solaris Cluster release supports systems that use the SPARC and x86 families of processor architectures. In this document, "x86" refers to the larger family of x86 compatible products. Information in this document pertains to all platforms unless otherwise specified.

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

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

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

Bash is the default shell for Oracle Solaris 11. Machine names shown with the Bash shell prompt are displayed for clarity.

Using UNIX Commands

This document contains information about commands that are specific to installing and configuring Oracle Solaris Cluster data services. The document does *not* contain comprehensive information about basic UNIX commands and procedures, such as shutting down the system, booting the system, and configuring devices. Information about basic UNIX commands and procedures is available from the following sources:

- Online documentation for the Oracle Solaris Operating System
- Oracle Solaris Operating System man pages

- Other software documentation that you received with your system

Typographic Conventions

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

TABLE P–1 Typographic Conventions

Typeface	Description	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 shells that are included in the Oracle Solaris OS. Note that the default system prompt that is displayed in command examples varies, depending on the Oracle Solaris release.

TABLE P–2 Shell Prompts

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

Related Documentation

Information about related Oracle Solaris Cluster topics is available in the documentation that is listed in the following table. All Oracle Solaris Cluster documentation is available at http://www.oracle.com/technetwork/indexes/documentation/index.html.

Topic	Documentation
Hardware installation and administration	*Oracle Solaris Cluster 4.0 Hardware Administration Manual* Individual hardware administration guides
Concepts	*Oracle Solaris Cluster Concepts Guide*
Software installation	*Oracle Solaris Cluster Software Installation Guide*
Data service installation and administration	*Oracle Solaris Cluster Data Services Planning and Administration Guide* and individual data service guides
Data service development	*Oracle Solaris Cluster Data Services Developer's Guide*
System administration	*Oracle Solaris Cluster System Administration Guide* *Oracle Solaris Cluster Quick Reference*
Software upgrade	*Oracle Solaris Cluster Upgrade Guide*
Error messages	*Oracle Solaris Cluster Error Messages Guide*
Command and function references	*Oracle Solaris Cluster Reference Manual* *Oracle Solaris Cluster Data Services Reference Manual* *Oracle Solaris Cluster Geographic Edition Reference Manual* *Oracle Solaris Cluster Quorum Server Reference Manual*

Access to Oracle Support

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

Getting Help

If you have problems installing or using Oracle Solaris Cluster, contact your service provider and provide the following information.

- Your name and email address (if available)
- Your company name, address, and phone number
- The model number and serial number of your systems
- The release number of the operating environment (for example, Oracle Solaris 11)
- The release number of Oracle Solaris Cluster (for example, Oracle Solaris Cluster 4.0)

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

Command	Function
`prtconf -v`	Displays the size of the system memory and reports information about peripheral devices
`psrinfo -v`	Displays information about processors
`pkg list`	Reports which packages are installed
`prtdiag -v`	Displays system diagnostic information
`/usr/cluster/bin/clnode show-rev`	Displays Oracle Solaris Cluster release and package version information for each node

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

◆ ◆ ◆ CHAPTER 1

Preparing to Upgrade Oracle Solaris Cluster Software

This chapter provides the following information and procedures to prepare to upgrade or update the Oracle Solaris Cluster 4.0 software:

An upgrade moves the cluster to the latest major or minor Oracle Solaris Cluster release by updating all packages. The upgrade also upgrades the Oracle Solaris OS to the latest compatible version.

A software update changes a specific Oracle Solaris Cluster package to a newer Support Repository Update (SRU) level. For more information on using the `pkg` commands to update packages, see "Overview of Updating Oracle Solaris Cluster Software" in *Oracle Solaris Cluster System Administration Guide*.

Upgrade Requirements and Software Support Guidelines

Observe the following requirements and software-support guidelines when you upgrade the Oracle Solaris Cluster 4.0 software:

- **Minimum Oracle Solaris OS** - The cluster must run on Oracle Solaris 11 software.
- **Supported hardware** - The cluster hardware must be a supported configuration for Oracle Solaris Cluster 4.0 software. Contact your Oracle representative for information about current supported Oracle Solaris Cluster configurations.
- **Architecture changes during upgrade** - Oracle Solaris Cluster 4.0 software does not support upgrade between architectures.
- **Software migration** - Do not migrate from one type of software product to another product during Oracle Solaris Cluster upgrade. Perform only software configuration changes that are specified by upgrade procedures of an installed software product.

- **Data services** - You must upgrade data-service software to the latest Oracle Solaris Cluster version.
- **Upgrading to compatible versions** - You must upgrade all software on the cluster nodes to a version that is supported by Oracle Solaris Cluster 4.0 software. See "Supported Products" in *Oracle Solaris Cluster 4.0 Release Notes* for information about supported products.
- **Downgrade** - Oracle Solaris Cluster 4.0 software does not support any downgrade of Oracle Solaris Cluster software.
- **Limitation of scinstall for data-service upgrades** - The `scinstall` upgrade utility only upgrades those data services that are provided with Oracle Solaris Cluster 4.0 software. You must manually upgrade any custom or third-party data services.

Choosing an Oracle Solaris Cluster Upgrade Method

The following matrixes summarize the supported upgrade methods for each Oracle Solaris OS version and platform, *provided that all other requirements for any supported method are met*. Check the documentation for other products in the cluster, such as volume management software and other applications, for any additional upgrade requirements or restrictions.

TABLE 1–1 Upgrade an Oracle Solaris Cluster 4.0 Cluster to a New Release

Method	Oracle Solaris 11 SPARC	Oracle Solaris 11 x86
Standard upgrade	X	X
Rolling upgrade	X	X

Choose from the following methods to upgrade your Oracle Solaris Cluster 4.0 cluster software:

- "Standard Upgrade" on page 10 – Upgrade the entire cluster to a new SRU release and upgrade the Oracle Solaris OS to the latest compatible version.
- "Rolling Upgrade" on page 11 – Upgrade one cluster node at a time to a new SRU release and upgrade the Oracle Solaris OS to the latest compatible version. Cluster nodes not being upgraded remain in production.

For overview information about planning your Oracle Solaris Cluster 4.0 configuration, see Chapter 1, "Planning the Oracle Solaris Cluster Configuration," in *Oracle Solaris Cluster Software Installation Guide*.

Standard Upgrade

A standard upgrade upgrades a cluster to a new release and upgrades the Oracles Solaris OS to the latest compatible version. You do not need to place the cluster in noncluster mode before

performing this upgrade because the upgrade always occurs in the new boot environment and the existing boot environment remains unchanged. You can specify a name for the new boot environment or you can use the auto-generated name.

Any time you upgrade the Oracle Solaris Cluster software, you should also upgrade the data services and Geographic Edition software. However, if you want to upgrade the data services separately, see "Overview of the Installation and Configuration Process" in *Oracle Solaris Cluster Data Services Planning and Administration Guide*. If you want to upgrade Oracle Solaris Cluster Geographic Edition separately, see the *Oracle Solaris Cluster Geographic Edition Installation Guide*.

Note – An SRU can be installed with either the `pkg` commands or the `scinstall -u` command. For instructions on using the `pkg` command to upgrade single or multiple packages, see "Updating a Specific Package" in *Oracle Solaris Cluster System Administration Guide*.

For complete instructions on performing a standard upgrade, see Chapter 3, "Performing a Standard Upgrade."

The cluster outage is limited to the amount of time that is needed to reboot the cluster nodes into the upgraded boot environment.

Rolling Upgrade

In a rolling upgrade, you upgrade software to an update release or a new SRU on one node at a time. Services continue on the other nodes except for the time it takes to switch services from a node to be upgraded to a node that will remain in service.

Observe the following additional restrictions and requirements for the rolling upgrade method:

- **Minimum Oracle Solaris Cluster version** - The cluster must be running an Oracle Solaris Cluster 4.0 release.
- **Oracle Solaris upgrade paths** - You can upgrade the Oracle Solaris OS only to a new SRU or an update version of the same release.
- **Hardware configuration changes** - Do *not* change the cluster configuration during a rolling upgrade. For example, do not add to or change the cluster interconnect or quorum devices. If you need to make such a change, do so before you start the rolling upgrade procedure or wait until after all nodes are upgraded and the cluster is committed to the new software version.
- **Duration of the upgrade** - Limit the amount of time that you take to complete a rolling upgrade of all cluster nodes. After a node is upgraded, begin the upgrade of the next cluster node as soon as possible. You can experience performance penalties and other penalties when you run a mixed-version cluster for an extended period of time.

- **Software configuration changes** - Avoid installing new data services or issuing any administrative configuration commands during the upgrade.
- **New-feature availability** - Until all nodes of the cluster are successfully upgraded and the upgrade is committed, new features that are introduced by the new release might not be available.

◆ ◆ ◆ CHAPTER 2

Upgrading Zones Managed by HA for Oracle Solaris Zones

This chapter provides the following information to upgrade Oracle Solaris non-global zones that are managed by the Oracle Solaris Cluster software.

Upgrading Failover Zones

A failover zone is a non-global zone that is configured with the HA for Oracle Solaris Zones data service so that it can be managed by the Oracle Solaris Cluster software. If you have failover zones of brand type `solaris` configured on the cluster, perform this procedure in conjunction with the procedures for the cluster upgrade method that you use.

▼ How to Upgrade Failover Zones

1 **Determine the nodes where the resource group that manages the failover zone of brand type `solaris` is offline.**

```
# clresourcegroup status resourcegroup
```

2 **Suspend the resource group that manages the failover zone of brand type `solaris`.**

```
# clresourcegroup suspend resourcegroup
```

3 **Perform the standard or rolling upgrade on the nodes where the resource group that manages the failover zones is offline.**

Perform all steps, including booting into the new boot environment, in "How to Perform a Standard Upgrade" on page 21 or "How to Perform a Rolling Upgrade" on page 30. Then return to this procedure.

4 **Perform the standard or rolling upgrade on the nodes where the resource group that manages the failover zones of brand type `solaris` is online.**

Perform Steps 1–3 in "How to Perform a Standard Upgrade" on page 21 or "How to Perform a Rolling Upgrade" on page 30. Do *not* perform Step 4 to boot the machine into the new boot environment. Return to this procedure after you complete this task.

5 **Determine the universally unique ID (UUID) of the updated boot-environment root dataset of the node where the resource group that manages the failover zone of brand type `solaris` is online.**

You will see output similar to the following:

```
phys-schost-1# beadm list -H
...
b175b-SC;8fe53702-16c3-eb21-ed85-d19af92c6bbd;R;/;756
...
```

In this example, the UUID is `8fe53702-16c3-eb21-ed85-d19af92c6bbd` and the updated boot environment is `b175b-SC`.

6 **Set the same UUID on the updated boot environment for the other nodes where the resource group that manages the failover zone of brand type `solaris` is offline.**

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

For example:

```
phys-schost-2# zfs set org.opensolaris.libbe:uuid=8fe53702-16c3-eb21-ed85-d19af92c6bbd rpool/ROOT/b175b-SC
```

7 **Resume the resource group that manages the failover zone of brand type `solaris`.**

```
# clresourcegroup resume resourcegroup
```

8 **Switch the resource group that manages the failover zone of brand type `solaris` to a node that was already booted into the updated boot environment and verify that the zone started correctly.**

```
# clresourcegroup switch -n phys-schost-2 resourcegroup
```

9 **Boot the node from Step 4, where the resource group that manages the failover zone of brand type `solaris` had initially been online, into its updated boot environment.**

Next Steps Go to Chapter 5, "Completing the Upgrade."

◆ ◆ ◆ CHAPTER 3

Performing a Standard Upgrade

This chapter provides the following information to upgrade a Software Repository Update (SRU) for Oracle Solaris Cluster 4.0 software by using the standard nonrolling upgrade method:

- "How to Upgrade Quorum Server Software" on page 16
- "How to Prepare the Cluster for a Standard Upgrade" on page 17
- "How to Perform a Standard Upgrade" on page 21

Note – Upgrading Oracle Solaris Cluster 4.0 software also upgrades the Oracle Solaris Operating System to the latest compatible version.

Performing a Standard Upgrade of a Cluster

The following table lists the tasks to upgrade from Oracle Solaris Cluster 4.0 to a 4.0 SRU. Performing these tasks also upgrades the Oracle Solaris OS to the latest compatible version.

TABLE 3–1 Task Map: Performing a Standard Upgrade for Oracle Solaris Cluster 4.0 Software

Task	Instructions
1. Read the upgrade requirements and restrictions. Determine the proper upgrade method for your configuration and needs.	"Upgrade Requirements and Software Support Guidelines" on page 9 "Choosing an Oracle Solaris Cluster Upgrade Method" on page 10
2. If failover zones of brand type `solaris` are configured in the cluster, upgrade the failover zones.	"Upgrading Failover Zones" on page 13
2. If a quorum server is used, upgrade the Quorum Server software.	"How to Upgrade Quorum Server Software" on page 16
3. Remove the cluster from production and back up shared data.	"How to Prepare the Cluster for a Standard Upgrade" on page 17

TABLE 3–1 Task Map: Performing a Standard Upgrade for Oracle Solaris Cluster 4.0 Software *(Continued)*

Task	Instructions
4. If the cluster uses dual-string mediators for Solaris Volume Manager software, unconfigure the mediators. Solaris Volume Manager software is automatically upgraded with the Oracle Solaris OS.	Follow upgrade procedures in Oracle Solaris installation documentation.
5. Upgrade to Oracle Solaris Cluster 4.0 framework and data-service software. If necessary, upgrade applications. If the cluster uses dual-string mediators and you upgraded the Oracle Solaris OS, reconfigure the mediators.	"How to Perform a Standard Upgrade" on page 21
7. Use the `scversions` command to commit the cluster to the upgrade.	"How to Commit the Upgraded Cluster" on page 31
8. Verify successful completion of upgrade to Oracle Solaris Cluster 4.0 software.	"How to Verify the Upgrade" on page 32
9. Enable resources and bring resource groups online. Migrate existing resources to new resource types. Upgrade to Oracle Solaris Cluster Geographic Edition 4.0 software, if used.	"How to Finish the Upgrade" on page 33

▼ How to Upgrade Quorum Server Software

If the cluster uses a quorum server, upgrade the Quorum Server software on the quorum server *before* you upgrade the cluster.

Note – If more than one cluster uses the quorum server, perform on each cluster the steps to remove the quorum server and later the steps to add back the quorum server.

Perform all steps as superuser on the cluster and on the quorum server.

1 If the cluster has two nodes and the quorum server is the cluster's only quorum device, temporarily add a second quorum device.

See "Adding a Quorum Device" in *Oracle Solaris Cluster System Administration Guide*.

If you add another quorum server as a temporary quorum device, the quorum server can run the same software version as the quorum server that you are upgrading, or it can run the 4.0 version of Quorum Server software.

2 Unconfigure the quorum server from each cluster that uses the quorum server.

```
phys-schost# clquorum remove quorumserver
```

3 From the quorum server to upgrade, verify that the quorum server no longer serves any cluster.

```
quorumserver# clquorumserver show +
```

If the output shows any cluster is still served by the quorum server, unconfigure the quorum server from that cluster. Then repeat this step to confirm that the quorum server is no longer configured with any cluster.

Note – If you have unconfigured the quorum server from a cluster but the `clquorumserver show` command still reports that the quorum server is serving that cluster, the command might be reporting stale configuration information. See "Cleaning Up Stale Quorum Server Cluster Information" in *Oracle Solaris Cluster System Administration Guide*.

---

4 **From the quorum server to upgrade, halt all quorum server instances.**

```
quorumserver# clquorumserver stop +
```

5 **Uninstall the Quorum Server software from the quorum server to upgrade.**

```
quorumserver# pkg uninstall ha-cluster/*
```

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

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

7 **Install the Oracle Solaris Cluster 4.0 Quorum Server software, reconfigure the quorum server, and start the quorum server daemon.**

Follow the steps in "How to Install and Configure Oracle Solaris Cluster Quorum Server Software" in *Oracle Solaris Cluster Software Installation Guide* for installing the Quorum Server software.

8 **From a cluster node, configure the upgraded quorum server as a quorum device.**

Follow the steps in "How to Configure Quorum Devices" in *Oracle Solaris Cluster Software Installation Guide*.

9 **If you configured a temporary quorum device, unconfigure it.**

```
phys-schost# clquorum remove tempquorum
```

▼ How to Prepare the Cluster for a Standard Upgrade

Perform this procedure to remove the cluster from production before you perform a standard upgrade. Performing a standard upgrade also upgrades the Oracle Solaris OS to the latest compatible version. Perform all steps from the global zone only.

Before You Begin Perform the following tasks:

- Ensure that the configuration meets the requirements for upgrade. See "Upgrade Requirements and Software Support Guidelines" on page 9.

- Have available the installation media, documentation, and software updates for all software products that you are upgrading, including the following software:
  - Oracle Solaris OS
  - Oracle Solaris Cluster
  - Applications that are managed by Oracle Solaris Cluster data services
  - Any other third-party applications to upgrade

  For instructions on updating single or multiple packages, see Chapter 11, "Updating Your Software," in *Oracle Solaris Cluster System Administration Guide*.
- If you use role-based access control (RBAC) instead of superuser to access the cluster nodes, ensure that you can assume an RBAC role that provides authorization for all Oracle Solaris Cluster commands. This series of upgrade procedures requires the following Oracle Solaris Cluster RBAC authorizations if the user is not superuser:
  - `solaris.cluster.modify`
  - `solaris.cluster.admin`
  - `solaris.cluster.read`

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

1 Ensure that the cluster is functioning normally.

a. View the current status of the cluster by running the following command from any node.

```
phys-schost% cluster status
```

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

b. Search the `/var/adm/messages` log on the same node for unresolved error messages or warning messages.

c. Check the volume-manager status.

2 Become superuser on a node of the cluster.

3 Take each resource group offline and disable all resources.

Take offline all resource groups in the cluster, including those that are in non-global zones. Then disable all resources, to prevent the cluster from bringing the resources online automatically if a node is mistakenly rebooted into cluster mode.

- **If you want to use the `clsetup` utility, perform the following steps:**

  a. Start the utility.

  ```
  phys-schost# clsetup
  ```

The Main Menu is displayed.

b. Choose the menu item, Resource Groups.

The Resource Group Menu is displayed.

c. Choose the menu item, Online/Offline or Switchover a Resource Group.

d. Follow the prompts to take offline all resource groups and to put them in the unmanaged state.

e. When all resource groups are offline, type `q` to return to the Resource Group Menu.

f. Exit the `clsetup` utility.

Type **`q`** to back out of each submenu or press **`Ctrl-C`**.

- **To use the command line, perform the following steps:**

a. Take each resource offline.

```
phys-schost# clresource offline resource-group
```

b. From any node, list all enabled resources in the cluster.

```
phys-schost# clresource show -p Enabled
=== Resources ===

Resource:                                       resource
  Enabled{nodename1}:                              True
  Enabled{nodename2}:                              True
...
```

c. Identify those resources that depend on other resources.

```
phys-schost# clresource show -p resource_dependencies
=== Resources ===

Resource:                                       node
  Resource_dependencies:                           node
...
```

You must disable dependent resources first before you disable the resources that they depend on.

d. Disable each enabled resource in the cluster.

```
phys-schost# clresource disable resource
```

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

e. Verify that all resources are disabled.

```
phys-schost# clresource show -p Enabled
=== Resources ===
```

```
Resource:                                          resource
  Enabled{nodename1}:                                   False
  Enabled{nodename2}:                                   False
...
```

f. Move each resource group to the unmanaged state.

```
phys-schost# clresourcegroup unmanage resource-group
```

4 Verify that all resources on all nodes are `Offline` and that all resource groups are in the `Unmanaged` state.

```
phys-schost# cluster status -t resource,resourcegroup
```

5 Stop all applications that are running on each node of the cluster.

6 Ensure that all shared data is backed up.

7 If your cluster uses dual-string mediators for Solaris Volume Manager software, unconfigure your mediators.

See "Configuring Dual-String Mediators" in *Oracle Solaris Cluster Software Installation Guide* for more information about mediators.

a. Run the following command to verify that no mediator data problems exist.

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

`-s setname`
Specifies the disk set name.

If the value in the Status field is `Bad`, repair the affected mediator host. Follow the procedure "How to Check For and Fix Bad Mediator Data" in *Oracle Solaris Cluster Software Installation Guide*.

b. List all mediators.

Save this information for when you restore the mediators during the procedure "How to Finish the Upgrade" on page 33.

c. For a disk set that uses mediators, take ownership of the disk set if no node already has ownership.

```
phys-schost# cldevicegroup switch -n node devicegroup
```

d. Unconfigure all mediators for the disk set.

```
phys-schost# metaset -s setname -d -m mediator-host-list
```

`-s` *setname*
Specifies the disk set name.

`-d`
Deletes from the disk set.

`-m` *mediator-host-list*
Specifies the name of the node to remove as a mediator host for the disk set.

See the `mediator(7D)` man page for further information about mediator-specific options to the `metaset` command.

e. Repeat Step c through Step d for each remaining disk set that uses mediators.

8 Ensure that each system disk is backed up.

Next Steps Upgrade the cluster and OS software. Go to "How to Perform a Standard Upgrade" on page 21.

▼ How to Perform a Standard Upgrade

A standard upgrade upgrades a cluster to a new release, including installed data service software, and upgrades the Oracles Solaris OS to the latest compatible version. You do not need to place the cluster in noncluster mode before performing this upgrade because the upgrade always occurs in the new boot environment and the existing boot environment remains unchanged. You can specify a name for the new boot environment or you can use the auto-generated name.

Depending on the content in the repositories, you might have to unset the undesired publishers or freeze the undesired incorporations if you do not want to upgrade the cluster software and the OS all at once. For instructions, see Chapter 5, "Configuring Installed Images," in *Adding and Updating Oracle Solaris 11 Software Packages* or the `pkg(1)` man page.

Anytime you upgrade the Oracle Solaris Cluster software, you should also upgrade the data services and Geographic Edition software.

Note – If you want to install an SRU by updating a single or multiple package, see "How to Update a Specific Package" in *Oracle Solaris Cluster System Administration Guide*. An SRU can be installed with either the `pkg` commands or the `scinstall -u` command.

---

1 Become superuser or assume a role that provides `solaris.cluster.admin` RBAC authorization.

2 Subscribe to the `ha-cluster` publisher which you want to upgrade.

```
# pkg unset-publisher ha-cluster
# pkg set-publisher -g FMRI for cluster ha-cluster
```

3 Run the upgrade.

```
# scinstall - u update [-b bename]
```

You can choose to specify a name for the new boot environment with the **-b** *bename* option.

If you are upgrading a failover zone, follow the instructions in "How to Upgrade Failover Zones" on page 13.

4 Follow the prompts to boot the machine into the new boot environment.

Next Steps Go to Chapter 5, "Completing the Upgrade."

◆ ◆ ◆ CHAPTER 4

Performing a Rolling Upgrade

This chapter provides procedures to perform a rolling upgrade of an Oracle Solaris Cluster 4.0 release to a newer Oracle Solaris Cluster 4.0 SRU, or to upgrade Oracle Solaris 11 to a newer Oracle Solaris 11 SRU. In a rolling upgrade, you upgrade one cluster node at a time, while the other cluster nodes remain in production. After all nodes are upgraded and have rejoined the cluster, you must commit the cluster to the new software version before you can use any new features.

Note – Oracle Solaris Cluster software does not support rolling upgrade to Oracle Solaris Cluster software from a previous marketing release. You can only perform an Oracle Solaris Cluster rolling upgrade of Oracle Solaris Cluster or Oracle Solaris software to an update of the same marketing release.

To upgrade an Oracle Solaris Cluster configuration from an earlier marketing release of Oracle Solaris software, use another upgrade method. See "Choosing an Oracle Solaris Cluster Upgrade Method" on page 10 to determine the best upgrade method for your configuration.

This chapter provides the following information to upgrade an Oracle Solaris Cluster 4.0 configuration to a new SRU of the Oracle Solaris Cluster 4.0 software or to an SRU of the Oracle Solaris OS by using the rolling upgrade method:

Performing a Rolling Upgrade of a Cluster

TABLE 4–1 Task Map: Performing a Rolling Upgrade to Oracle Solaris Cluster 4.0 Software

Task	Instructions
1. Read the upgrade requirements and restrictions.	"Upgrade Requirements and Software Support Guidelines" on page 9
2. If failover zones of brand type `solaris` are configured in the cluster, upgrade the failover zones.	"Upgrading Failover Zones" on page 13
2. If a quorum server is used, upgrade the Quorum Server software.	"How to Upgrade Quorum Server Software" on page 24
3. On one node of the cluster, move resource groups and device groups to another cluster node, and ensure that shared data and system disks are backed up. If the cluster uses dual-string mediators for Solaris Volume Manager software, unconfigure the mediators. Then reboot the node into noncluster mode.	"How to Prepare a Cluster Node for a Rolling Upgrade" on page 26
4. Upgrade the Oracle Solaris OS on the cluster node, if necessary, to a supported Oracle Solaris update release.	"How to Perform a Rolling Upgrade of the Solaris OS" on page 29
5. Upgrade the cluster node to Oracle Solaris Cluster 4.0 framework software. Optionally, upgrade data-service software. If necessary, upgrade applications.	"How to Perform a Rolling Upgrade" on page 30
6. Repeat Tasks 3 through 4 on each remaining node to upgrade.	
7. Use the `scversions` command to commit the cluster to the upgrade.	"How to Commit the Upgraded Cluster" on page 31
8. Verify successful completion of upgrade to Oracle Solaris Cluster 4.0 software.	"How to Verify the Upgrade" on page 32
9. Enable resources and bring resource groups online. Migrate existing resources to new resource types. Upgrade to the Oracle Solaris Cluster Geographic Edition 4.0 software, if used.	"How to Finish the Upgrade" on page 33

▼ How to Upgrade Quorum Server Software

If the cluster uses a quorum server, upgrade the Quorum Server software on the quorum server *before* you upgrade the cluster.

Note – If more than one cluster uses the quorum server, perform these steps for each of those clusters.

Perform all steps as superuser on the cluster and on the quorum server.

1. **If the cluster has two nodes and the quorum server is the cluster's only quorum device, temporarily add a second quorum device.**

   See "Adding a Quorum Device" in *Oracle Solaris Cluster System Administration Guide*.

   If you add another quorum server as a temporary quorum device, the quorum server can run the same software version as the quorum server that you are upgrading, or it can run the 4.0 version of Quorum Server software.

2. **Unconfigure the quorum server from each cluster that uses the quorum server.**

   ```
   phys-schost# clquorum remove quorumserver
   ```

3. **From the quorum server to upgrade, verify that the quorum server no longer serves any cluster.**

   ```
   quorumserver# clquorumserver show +
   ```

   If the output shows any cluster is still served by the quorum server, unconfigure the quorum server from that cluster. Then repeat this step to confirm that the quorum server is no longer configured with any cluster.

   Note – If you have unconfigured the quorum server from a cluster but the `clquorumserver show` command still reports that the quorum server is serving that cluster, the command might be reporting stale configuration information. See "Cleaning Up Stale Quorum Server Cluster Information" in *Oracle Solaris Cluster System Administration Guide*.

4. **From the quorum server to upgrade, halt all quorum server instances.**

   ```
   quorumserver# clquorumserver stop +
   ```

5. **Uninstall the Quorum Server software from the quorum server to upgrade.**

   ```
   quorumserver# pkg uninstall ha-cluster/*
   ```

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

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

7. **Install the Oracle Solaris Cluster 4.0 Quorum Server software, reconfigure the quorum server, and start the quorum server daemon.**

   Follow the steps in "How to Install and Configure Oracle Solaris Cluster Quorum Server Software" in *Oracle Solaris Cluster Software Installation Guide* for installing the Quorum Server software.

8. **From a cluster node, configure the upgraded quorum server as a quorum device.**

   Follow the steps in "How to Configure Quorum Devices" in *Oracle Solaris Cluster Software Installation Guide*.

9 If you configured a temporary quorum device, unconfigure it.

```
phys-schost# clquorum remove tempquorum
```

▼ How to Prepare a Cluster Node for a Rolling Upgrade

Perform this procedure on one node at a time. You will take the upgraded node out of the cluster while the remaining nodes continue to function as active cluster members.

Before You Begin Perform the following tasks:

- Ensure that the configuration meets requirements for upgrade. See "Upgrade Requirements and Software Support Guidelines" on page 9.
- Have available the installation media, documentation, and upgrades for all the software products that you are upgrading, including the following software:
  - Oracle Solaris OS
  - Oracle Solaris Cluster 4.0 framework
  - Oracle Solaris Cluster 4.0 required software updates
  - Oracle Solaris Cluster 4.0 data services (agents)
  - Applications that are managed by Oracle Solaris Cluster 4.0 data service agents

  For instructions on updating single or multiple packages, see Chapter 11, "Updating Your Software," in *Oracle Solaris Cluster System Administration Guide*.

1 Ensure that the cluster is functioning normally.

a. View the current status of the cluster by running the following command from any node.

```
phys-schost% cluster status
```

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

b. Search the `/var/adm/messages` log on the same node for unresolved error messages or warning messages.

c. Check the volume-manager status.

2 If necessary, notify users that cluster services might be temporarily interrupted during the upgrade.

Service interruption will be approximately the amount of time that your cluster normally takes to switch services to another node.

3 Become superuser on a node of the cluster.

4 Move all resource groups and device groups that are running on the node to upgrade.

```
phys-schost# clnode evacuate node-to-evacuate
```

See the clnode(1CL) man page for more information.

5 **Move any resource groups that are running in a zone cluster node on the node to upgrade.**

```
phys-schost# clresourcegroup evacuate -n zone-cluster-node \
-Z zone-cluster-name resource-group
```

6 **Verify that the move was completed successfully.**

```
phys-schost# cluster status -t devicegroup,resourcegroup
```

7 **Ensure that the system disk, applications, and all data are backed up.**

8 **If you will upgrade the Oracle Solaris OS and your cluster uses dual-string mediators for Solaris Volume Manager software, unconfigure your mediators.**

See "Configuring Dual-String Mediators" in *Oracle Solaris Cluster Software Installation Guide* for more information.

a. **Run the following command to verify that no mediator data problems exist.**

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

`-s setname`
Specifies the disk set name

If the value in the Status field is Bad, repair the affected mediator host. Follow the procedure "How to Check For and Fix Bad Mediator Data" in *Oracle Solaris Cluster Software Installation Guide*.

b. **List all mediators.**

Save this information for when you restore the mediators during the procedure "How to Commit the Upgraded Cluster" on page 31.

c. **For a disk set that uses mediators, take ownership of the disk set if no node already has ownership.**

```
phys-schost# cldevicegroup switch -n node devicegr
```

d. **Unconfigure all mediators for the disk set.**

```
phys-schost# metaset -s setname -d -m mediator-host-list
```

`-s setname`
Specifies the disk-set name

`-d`
Deletes from the disk set

`-m mediator-host-list`
Specifies the name of the node to remove as a mediator host for the disk set

See the `mediator(7D)` man page for further information about mediator-specific options to the `metaset` command.

e. Repeat these steps for each remaining disk set that uses mediators.

9 Shut down the node that you want to upgrade and boot it into noncluster mode.

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

```
phys-schost# shutdown -y -g0
ok boot -x
```

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

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

  For more information about GRUB based booting, see *Booting and Shutting Down Oracle Solaris on x86 Platforms*.

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

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

```
grub edit> kernel /platform/i86pc/kernel/$ISADIR/unix -B $ZFS-BOOTFS -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 to add the `-x` option to the kernel boot parameter command.

The other nodes of the cluster continue to function as active cluster members.

Next Steps To upgrade the nodes of an Oracle Solaris Cluster cluster to a new SRU release, go to "How to Perform a Rolling Upgrade" on page 30. To upgrade your Oracle Solaris software, go to "How to Perform a Rolling Upgrade of the Solaris OS" on page 29.

Note – The cluster must already run on, or be upgraded to, at least the minimum required level of the Oracle Solaris OS to support Oracle Solaris Cluster 4.0 software. See "Supported Products" in *Oracle Solaris Cluster 4.0 Release Notes*.

▼ How to Perform a Rolling Upgrade of the Solaris OS

Perform this procedure to upgrade the Oracle Solaris OS to a supported SRU release.

Before You Begin Ensure that all steps in "How to Prepare a Cluster Node for a Rolling Upgrade" on page 26 are completed.

1 Follow the instructions in the Oracle Solaris installation guide to install the SRU.

Note – Do not reboot the node when prompted to reboot at the end of installation processing.

2 Install any required Oracle Solaris software updates and hardware-related updates, and download any needed firmware that is contained in the hardware updates.

Note – Do not reboot the node until Step 3.

3 Reboot the node into noncluster mode.

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

```
phys-schost# shutdown -y -g0
ok boot -x
```

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

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

     For more information about GRUB based booting, see *Booting and Shutting Down Oracle Solaris on x86 Platforms*.

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

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

```
grub edit> kernel /platform/i86pc/kernel/$ISADIR/unix -B $ZFS-BOOTFS -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 to add the `-x` option to the kernel boot parameter command.

Next Steps Go to "How to Perform a Rolling Upgrade" on page 30.

▼ How to Perform a Rolling Upgrade

Perform this procedure to upgrade an SRU on a node that runs Oracle Solaris Cluster 4.0 software while the remaining cluster nodes are in cluster mode.

If you performed a rolling upgrade and you have failover zones of brand type `solaris` configured on the cluster, you must perform additional steps. Follow the instructions in "How to Upgrade Failover Zones" on page 13.

Note – Until all nodes of the cluster are upgraded and the upgrade is committed, new features that are introduced by the new release might not be available.

---

1 Become superuser or assume a role that provides `solaris.cluster.admin` RBAC authorization on the node of the cluster you want to upgrade.

2 Subscribe to the `ha-cluster` publisher which you want to upgrade.

```
# pkg unset-publisher ha-cluster
# pkg set-publisher -g FMRI for cluster ha-cluster
```

3 Run the upgrade on the node you want to upgrade.

```
# scinstall - u update [-b bename]
```

You can choose to specify a name for the new boot environment with the **-b** *bename* option.

4 Follow the prompts to boot the node into the new boot environment.

Perform these stepson each remaining node, one at a time, that you want to upgrade. After you have completed the upgrade, you must perform additional steps to commit the upgrade. See Chapter 5, "Completing the Upgrade."

Next Steps When all nodes in the cluster are upgraded, go to Chapter 5, "Completing the Upgrade."

◆ ◆ ◆ CHAPTER 5

Completing the Upgrade

This chapter provides the following information to complete all Oracle Solaris Cluster 4.0 software upgrade methods:

Completing a Cluster Upgrade

▼ How to Commit the Upgraded Cluster

Before You Begin Ensure that all upgrade procedures are completed for all cluster nodes that you are upgrading.

1. **From one node, check the upgrade status of the cluster.**

```
phys-schost# scversions
```

2. **From the following table, perform the action that is listed for the output message from Step 1.**

Output Message	Action
`Upgrade commit is needed.`	Proceed to Step 3.
`Upgrade commit is NOT needed. All versions match.`	Go to "How to Verify the Upgrade" on page 32.
`Upgrade commit cannot be performed until all cluster nodes are upgraded. Please run scinstall(1m) on cluster nodes to identify older versions.`	Return to the Oracle Solaris Cluster upgrade procedures that you used and upgrade the remaining cluster nodes.

Output Message	Action
`Check upgrade cannot be performed until all cluster nodes are upgraded. Please run scinstall(1m) on cluster nodes to identify older versions.`	Return to the Oracle Solaris Cluster upgrade procedures that you used and upgrade the remaining cluster nodes.

3 After all nodes have rejoined the cluster, from one node commit the cluster to the upgrade.

```
phys-schost# scversions -c
```

Committing the upgrade enables the cluster to utilize all features in the newer software. New features are available only after you perform the upgrade commitment.

4 From one node, verify that the cluster upgrade commitment has succeeded.

```
phys-schost# scversions
Upgrade commit is NOT needed. All versions match.
```

Next Steps Go to "How to Verify the Upgrade" on page 32.

▼ How to Verify the Upgrade

Perform this procedure to verify that the cluster is successfully upgraded to Oracle Solaris Cluster 4.0 software. Perform all steps from the global zone only.

Before You Begin

- Ensure that all upgrade procedures are completed for all cluster nodes that you are upgrading.
- Ensure that all steps in "How to Commit the Upgraded Cluster" on page 31 are completed successfully.

1 On each node, become superuser.

2 On each upgraded node, view the installed levels of Oracle Solaris Cluster software.

```
phys-schost# clnode show-rev -v
```

The first line of output states which version of Oracle Solaris Cluster software the node is running. This version should match the version that you just upgraded to.

3 From any node, verify that all upgraded cluster nodes are running in cluster mode (`Online`).

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

See the `clnode(1CL)` man page for more information about displaying cluster status.

Example 5–1 Verifying Upgrade to Oracle Solaris Cluster 4.0 Software

The following example shows the commands used to verify upgrade of a two-node cluster to Oracle Solaris Cluster 4.0 software. The cluster node names are `phys-schost-1` and `phys-schost-2`.

```
phys-schost# clnode show-rev -v
4.0
...
phys-schost# clnode status
=== Cluster Nodes ===

--- Node Status ---

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

Next Steps Go to "How to Finish the Upgrade" on page 33.

▼ How to Finish the Upgrade

Perform this procedure to finish Oracle Solaris Cluster upgrade. Perform all steps from the global zone only.

Before You Begin Ensure that all steps in "How to Verify the Upgrade" on page 32 are completed.

1 Copy the security files for the common agent container to all cluster nodes.

This step ensures that security files for the common agent container are identical on all cluster nodes and that the copied files retain the correct file permissions.

a. On each node, stop the security file agent.

```
phys-schost# /usr/sbin/cacaoadm stop
```

b. On one node, change to the `/etc/cacao/instances/default/` directory.

```
phys-schost-1# cd /etc/cacao/instances/default/
```

c. Create a tar file of the `/etc/cacao/instances/default/` directory.

```
phys-schost-1# tar cf /tmp/SECURITY.tar security
```

d. Copy the `/tmp/SECURITY.tar` file to each of the other cluster nodes.

e. On each node to which you copied the `/tmp/SECURITY.tar` file, extract the security files.

Any security files that already exist in the `/etc/cacao/instances/default/` directory are overwritten.

```
phys-schost-2# cd /etc/cacao/instances/default/
phys-schost-2# tar xf /tmp/SECURITY.tar
```

f. Delete the `/tmp/SECURITY.tar` file from each node in the cluster.

You must delete each copy of the tar file to avoid security risks.

```
phys-schost-1# rm /tmp/SECURITY.tar
phys-schost-2# rm /tmp/SECURITY.tar
```

2 If you upgraded Oracle Solaris software, ensure that external access to RPC communication is enabled.

If you did not upgrade the Oracle Solaris software, omit this step.

During Oracle Solaris OS upgrade, the default restricted network profile is applied, which disables external access for certain network services. This includes the RPC communication service, which is required for cluster communication.

The following steps check whether external access to these services are available and, if necessary, restore the Oracle Solaris functionality.

a. Display the status of external access to RPC communication.

External access is available if the output of the following command shows that the `local_only` property is set to false.

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

If external access is enabled, skip to Step 3. Otherwise, continue to Step b

b. If external access to RPC communication is restricted, perform 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
```

c. Repeat Step a to confirm that external access is restored.

3 On each node, start the security file agent.

```
phys-schost# /usr/sbin/cacaoadm start
```

4 If you upgraded any data services that are not supplied on the product media, register the new resource types for those data services.

Follow the documentation that accompanies the data services.

5 If you upgraded the Oracle Solaris OS and your configuration uses dual-string mediators for Solaris Volume Manager software, restore the mediator configurations.

a. Determine which node has ownership of a disk set to which you will add the mediator hosts.

```
phys-schost# metaset -s setname
```

`-s` *setname*
Specifies the disk set name.

b. On the node that masters or will master the disk set, become superuser.

c. If no node has ownership, take ownership of the disk set.

```
phys-schost# cldevicegroup switch -n node devicegroup
```

node
Specifies the name of the node to become primary of the disk set.

devicegroup
Specifies the name of the disk set.

d. Re-create the mediators.

```
phys-schost# metaset -s setname -a -m mediator-host-list
```

`-a`
Adds to the disk set.

`-m mediator-host-list`
Specifies the names of the nodes to add as mediator hosts for the disk set.

e. Repeat these steps for each disk set in the cluster that uses mediators.

6 Migrate resources to new resource type versions.

You must migrate all resources to the Oracle Solaris Cluster 4.0 resource-type version to use the new features and bug fixes that are provided in this release.

See "Upgrading a Resource Type" in *Oracle Solaris Cluster Data Services Planning and Administration Guide*, which contains procedures which use the command line. Alternatively, you can perform the same tasks by using the Resource Group menu of the `clsetup` utility. The process involves performing the following tasks:

- Registering the new resource type.
- Migrating the eligible resource to the new version of its resource type.
- Modifying the extension properties of the resource type.

Note – The Oracle Solaris Cluster 4.0 release might introduce new default values for some extension properties. These changes affect the behavior of any existing resource that uses the default values of such properties. If you require the previous default value for a resource, modify the migrated resource to set the property to the previous default value.

---

7 **If you upgraded to the Oracle Solaris 11 OS and the Apache `httpd.conf` file is located on a cluster file system, ensure that the `HTTPD` entry in the Apache control script still points to that location.**

   a. **View the `HTTPD` entry in the `/usr/apache/bin/apchectl` file.**

      The following example shows the `httpd.conf` file located on the `/global` cluster file system.

      ```
      phys-schost# cat /usr/apache/bin/apchectl | grep HTTPD=/usr
      HTTPD="/usr/apache/bin/httpd -f /global/web/conf/httpd.conf"
      ```

   b. **If the file does not show the correct `HTTPD` entry, update the file.**

      ```
      phys-schost# vi /usr/apache/bin/apchectl
      #HTTPD=/usr/apache/bin/httpd
      HTTPD="/usr/apache/bin/httpd -f /global/web/conf/httpd.conf"
      ```

8 **From any node, start the `clsetup` utility.**

   ```
   phys-schost# clsetup
   ```

   The `clsetup` Main Menu is displayed.

9 **Re-enable all disabled resources.**

   a. **Choose the menu item, Resource Groups.**

      The Resource Group Menu is displayed.

   b. **Choose the menu item, Enable/Disable a Resource.**

   c. **Choose a resource to enable and follow the prompts.**

   d. **Repeat Step c for each disabled resource.**

   e. **When all resources are re-enabled, type `q` to return to the Resource Group Menu.**

10 **Bring each resource group back online.**

   This step includes the bringing online of resource groups in non-global zones.

   a. **Choose the menu item, Online/Offline or Switchover a Resource Group.**

   b. **Follow the prompts to put each resource group into the managed state and then bring the resource group online.**

11 **When all resource groups are back online, exit the `clsetup` utility.**

   Type **`q`** to back out of each submenu, or press Ctrl-C.

12 If, before upgrade, you enabled automatic node reboot if all monitored shared-disk paths fail, ensure that the feature is still enabled.

Also perform this task if you want to configure automatic reboot for the first time.

a. Determine whether the automatic reboot feature is enabled or disabled.

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

- **If the `reboot_on_path_failure` property is set to `enabled`, no further action is necessary.**
- **If `reboot_on_path_failure` property is set to `disabled`, proceed to the next step to re-enable the property.**

b. Enable the automatic reboot feature.

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

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

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

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

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

13 If used, install or complete upgrade of Oracle Solaris Cluster Geographic Edition 4.0 software.

See Chapter 4, "Upgrading the Geographic Edition Software," in *Oracle Solaris Cluster Geographic Edition Installation Guide*.

14 Revalidate the upgraded cluster configuration.

See "How to Validate the Cluster" in *Oracle Solaris Cluster Software Installation Guide*.

15 (Optional) Capture the disk partitioning information for future reference.

```
phys-schost# prtvtoc /dev/rdsk/cNtXdYsZ > filename
```

Store the file in a location outside the cluster. If you make any disk configuration changes, run this command again to capture the changed configuration. If a disk fails and needs replacement, you can use this information to restore the disk partition configuration. For more information, see the `prtvtoc(1M)` man page.

16 (Optional) Make a backup of your cluster configuration.

An archived backup of your cluster configuration facilitates easier recovery of your cluster configuration.

For more information, see "How to Back Up the Cluster Configuration" in *Oracle Solaris Cluster System Administration Guide*.

Troubleshooting

Resource-type migration failure - Normally, you migrate resources to a new resource type while the resource is offline. However, some resources need to be online for a resource-type migration to succeed. If resource-type migration fails for this reason, error messages similar to the following are displayed:

```
phys-schost - Resource depends on a SUNW.HAStoragePlus type resource that is not
online anywhere.
(C189917) VALIDATE on resource nfsrs, resource group rg, exited with
non-zero exit status.
(C720144) Validation of resource nfsrs in resource group rg on node
phys-schost failed.
```

If resource-type migration fails because the resource is offline, use the `clsetup` utility to re-enable the resource and then bring its related resource group online. Then repeat migration procedures for the resource.

Java binaries location change - If the location of the Java binaries changed during the upgrade of Oracle Solaris software, you might see error messages similar to the following when you attempt to run the `cacaoadm start` command:

```
phys-schost# /usr/sbin/cacaoadm start
No suitable Java runtime found. Java 1.5.0_06 or higher is required.
Jan 3 17:10:26 ppups3 cacao: No suitable Java runtime found. Java 1.5.0_06 or
higher is required.
Cannot locate all the dependencies
```

This error is generated because the start command cannot locate the current location of the Java binaries. The `JAVA_HOME` property still points to the directory where the previous version of Java was located, but that previous version was removed during upgrade.

To correct this problem, change the setting of `JAVA_HOME` in the following configuration file to use the current Java directory:

```
/etc/opt/SUNWcacao/cacao.properties
```

Next Steps The cluster upgrade is complete.

◆ ◆ ◆ CHAPTER 6

Recovering From an Incomplete Upgrade

This chapter provides the following information to recover from certain kinds of incomplete upgrades:

- "Recovering From Storage Configuration Changes During Upgrade" on page 39

Recovering From Storage Configuration Changes During Upgrade

This section provides the following repair procedures to follow if changes were inadvertently made to the storage configuration during upgrade:

- "How to Handle Storage Reconfiguration During an Upgrade" on page 39
- "How to Resolve Mistaken Storage Changes During an Upgrade" on page 40

▼ How to Handle Storage Reconfiguration During an Upgrade

Any changes to the storage topology, including running Oracle Solaris Cluster commands, should be completed before you upgrade the cluster to Oracle Solaris 11 software. If, however, changes were made to the storage topology during the upgrade, perform the following procedure. This procedure ensures that the new storage configuration is correct and that existing storage that was not reconfigured is not mistakenly altered.

Before You Begin Ensure that the storage topology is correct. Check whether the devices that were flagged as possibly being replaced map to devices that actually were replaced. If the devices were not replaced, check for and correct possible accidental configuration changes, such as incorrect cabling.

1 On a node that is attached to the unverified device, become superuser.

2 **Manually update the unverified device for the node list.**

```
phys-schost# cldevice repair -n node[,...] device
```

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

3 **Update the DID driver for the node list.**

```
phys-schost# cldevice refresh -n node[,...]
```

4 **Repeat Step 2 through Step 3 on all other nodes that are attached to the unverified device.**

Next Steps Return to the remaining upgrade tasks. Go to "How to Perform a Standard Upgrade" on page 21.

▼ How to Resolve Mistaken Storage Changes During an Upgrade

If accidental changes are made to the storage cabling during the upgrade, perform the following procedure to return the storage configuration to the correct state.

Note – This procedure assumes that no physical storage was actually changed. If physical or logical storage devices were changed or replaced, instead follow the procedures in "How to Handle Storage Reconfiguration During an Upgrade" on page 39.

Before You Begin Return the storage topology to its original configuration. Check the configuration of the devices that were flagged as possibly being replaced, including the cabling.

1 **On each node of the cluster, become superuser.**

2 **Update the DID driver on each node of the cluster.**

```
phys-schost# scdidadm -ui
phys-schost# cldevice refresh
```

`-u`
Loads the device–ID configuration table into the kernel.

`-i`
Initializes the DID driver.

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

3 **If the `scdidadm` command returned any error messages in Step 2, make further modifications as needed to correct the storage configuration, then repeat Step 2.**

Next Steps Return to the remaining upgrade tasks. Go to "How to Perform a Standard Upgrade" on page 21.

Index

V