Oracle® Solaris Cluster Upgrade Guide

Part No: 821–1256–11
July 2012, Revision 11

ORACLE

130201@25097

Contents

Preface

The *Oracle Solaris Cluster Upgrade Guide for Solaris OS* 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: UltraSPARC, SPARC64, AMD64, and Intel 64. In this document, x86 refers to the larger family of 64-bit 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.

Using UNIX Commands

This document contains information about commands that are used to install, configure, or upgrade an Oracle Solaris 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 Oracle Solaris OS
- Other software documentation that you received with your system
- Oracle 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 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://docs.sun.com.

Topic	Documentation
Overview	*Oracle Solaris Cluster Overview*
	Oracle Solaris Cluster 3.3 Documentation Center
Concepts	*Oracle Solaris Cluster Concepts Guide*
Hardware installation and administration	*Oracle Solaris Cluster 3.3 Hardware Administration Manual*
	Individual hardware administration guides
Software installation	*Oracle Solaris Cluster Software Installation Guide*
Data service installation and administration	*Oracle Solaris Cluster Data Services Planning and Administration Guide*
	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 Quorum Server Reference Manual

For a complete list of Oracle Solaris Cluster documentation, see the release notes for your release of Oracle Solaris Cluster software at http://wikis.sun.com/display/SunCluster/Home/.

Documentation, Support, and Training

See the following web sites for additional resources:

- Documentation (`http://www.oracle.com/technetwork/indexes/documentation/index.html`)
- Support (`http://www.oracle.com/us/support/systems/index.html`)
- Training (`http://education.oracle.com`) – Click the Sun link in the left navigation bar.

Oracle Welcomes Your Comments

Oracle welcomes your comments and suggestions on the quality and usefulness of its documentation. If you find any errors or have any other suggestions for improvement, go to `http://www.oracle.com/technetwork/indexes/documentation/index.html` and click Feedback. Indicate the title and part number of the documentation along with the chapter, section, and page number, if available. Please let us know if you want a reply.

Oracle Technology Network (`http://www.oracle.com/technetwork/index.html`) offers a range of resources related to Oracle software:

- Discuss technical problems and solutions on the Discussion Forums (`http://forums.oracle.com`).
- Get hands-on step-by-step tutorials with Oracle By Example (`http://www.oracle.com/technology/obe/start/index.html`).
- Download Sample Code (`http://www.oracle.com/technology/sample_code/index.html`).

Getting Help

If you have problems installing or using Oracle Solaris Cluster software, contact your service provider and supply 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 Oracle Solaris OS (for example, Oracle Solaris 10)
- The release number of Oracle Solaris Cluster (for example, Oracle Solaris Cluster 3.3)

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
`showrev -p`	Reports which patches are installed
SPARC: `prtdiag -v`	Displays system diagnostic information
`/usr/cluster/bin/clnode show-rev`	Displays Oracle Solaris Cluster release and package version information

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 to Oracle Solaris Cluster 3.3 software:

Upgrade Requirements and Software Support Guidelines

Observe the following requirements and software-support guidelines when you upgrade to Oracle Solaris Cluster 3.3 software:

- **Upgrade of x86 based systems** - On x86 based systems, you cannot upgrade from the Solaris 9 OS to the Oracle Solaris 10 OS. You must reinstall the cluster with a fresh installation of the Oracle Solaris 10 OS and Oracle Solaris Cluster 3.3 software for x86 based systems. Follow procedures in Chapter 2, "Installing Software on Global-Cluster Nodes," in *Oracle Solaris Cluster Software Installation Guide*.
- **Minimum Oracle Solaris Cluster software version** - Oracle Solaris Cluster 3.3 software supports the following direct upgrade paths:
  - **SPARC: From version 3.1 8/05 through version 3.2 11/09** - Use the standard, dual-partition, or live upgrade method.
  - **From version 3.2 including update releases through version 3.2 11/09** - Use the standard, dual-partition, or live upgrade method.
  - **On version 3.3 to an Oracle Solaris Cluster 3.3 update release** *with no Oracle Solaris upgrade except to an Oracle Solaris update release, or to upgrade only Oracle Solaris to an update release* – You can also use the rolling upgrade method.

  See "Choosing an Oracle Solaris Cluster Upgrade Method" on page 13 for additional requirements and restrictions for each upgrade method.

- **Minimum Solaris OS** - The cluster must run on or be upgraded to at least Solaris 10 10/09 software, including the most current required patches, before you upgrade the Oracle Solaris Cluster software.
- **Supported hardware** - The cluster hardware must be a supported configuration for Oracle Solaris Cluster 3.3 software. Contact your Sun representative for information about current supported Oracle Solaris Cluster configurations.
- **Architecture changes during upgrade** - Oracle Solaris Cluster 3.3 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. For example, migration from Solaris Volume Manager disk sets to VxVM disk groups or from UFS file systems to VxFS file systems is not supported during Oracle Solaris Cluster upgrade. Perform only software configuration changes that are specified by upgrade procedures of an installed software product.
- **Global-devices partition size** - If a node uses a dedicated partition for the `/global/.devices/node@`*nodeid* file system and the partition is less than 512 Mbytes but it provides sufficient space for existing device nodes, you do not need to change the file-system size. The 512-Mbyte minimum applies to new installations of Oracle Solaris Cluster software.

  However, you must still ensure that the global-devices file system has ample space and ample inode capacity for existing devices and for any new devices that you intend to configure. Certain configuration changes, such as adding disks, disk volumes, or metadevices, might require increasing the partition size to provide sufficient additional inodes.
- **Data services** - You must upgrade data-service software to the Oracle Solaris Cluster 3.3 version.
- **Upgrading to compatible versions** - You must upgrade all software on the cluster nodes to a version that is supported by Oracle Solaris Cluster 3.3 software. For example, if a version of an application is supported on Sun Cluster 3.2 software but is not supported on Oracle Solaris Cluster 3.3 software, you must upgrade that application to the version that is supported on Oracle Solaris Cluster 3.3 software, if such a version exists. See "Supported Products" in Oracle Solaris Cluster 3.3 Release Notes for information about supported products.
- **Downgrade** - Oracle Solaris Cluster 3.3 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 3.3 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.

Note – If your cluster uses a ZFS root file system, you can upgrade the Oracle Solaris OS only by using the live upgrade method. See Oracle Solaris upgrade documentation for more information.

This limitation does not apply if you are not upgrading the Oracle Solaris OS.

TABLE 1–1 Upgrade From Oracle Solaris Cluster 3.1 8/05 Through 3.2 11/09 Software, Including Oracle Solaris OS Upgrade

	Oracle Solaris 10	
Method	**SPARC**	**x86**
Standard upgrade	X	X
Dual-partition upgrade	X	X
Live upgrade	X	X
Rolling upgrade	-	-

TABLE 1–2 Upgrade on Oracle Solaris Cluster 3.3 Software of Oracle Solaris OS Update Only

	Oracle Solaris 10	
Method	**SPARC**	**x86**
Standard upgrade	X	X
Dual-partition upgrade	X	X
Live upgrade	X	X
Rolling upgrade	X	X

Choose from the following methods to upgrade your cluster to Oracle Solaris Cluster 3.3 software:

- “Standard Upgrade” on page 14
- “Dual-Partition Upgrade” on page 14
- “Live Upgrade” on page 15
- “Rolling Upgrade” on page 16

For overview information about planning your Oracle Solaris Cluster 3.3 configuration, see Chapter 1, "Planning the Oracle Solaris Cluster Configuration," in *Oracle Solaris Cluster Software Installation Guide*.

Standard Upgrade

In a standard upgrade, you shut down the cluster before you upgrade the cluster nodes. You return the cluster to production after all nodes are fully upgraded.

- **ZFS root file systems** - If your cluster uses a ZFS root file system, you cannot use standard upgrade to upgrade the Solaris OS. You must use only the live upgrade method to upgrade the Solaris OS. But you can use standard upgrade to separately upgrade Oracle Solaris Cluster and other software.

Dual-Partition Upgrade

In a *dual-partition* upgrade, you divide the cluster into two groups of nodes. You bring down one group of nodes and upgrade those nodes. The other group of nodes continues to provide services. After you complete upgrade of the first group of nodes, you switch services to those upgraded nodes. You then upgrade the remaining nodes and boot them back into the rest of the cluster.

The cluster outage time is limited to the amount of time that is needed for the cluster to switch over services to the upgraded partition, with one exception. If you upgrade from the Sun Cluster 3.1 8/05 release and you intend to configure zone clusters, you must temporarily take the upgraded first partition out of cluster mode to set new private-network settings that were introduced in the Sun Cluster 3.2 release.

Observe the following additional restrictions and requirements for the dual–partition upgrade method:

- **ZFS root file systems** - If your cluster uses a ZFS root file system, you cannot use dual-partition upgrade to upgrade the Solaris OS. You must use only the live upgrade method to upgrade the Solaris OS. But you can use dual-partition upgrade to separately upgrade Oracle Solaris Cluster and other software.
- **HA for Sun Java System Application Server EE (HADB)** - If you are running the HA for Sun Java System Application Server EE (HADB) data service with Sun Java System Application Server EE (HADB) software as of version 4.4, you must shut down the database before you begin the dual-partition upgrade. The HADB database does not tolerate the loss of membership that would occur when a partition of nodes is shut down for upgrade. This requirement does not apply to versions before version 4.4.

- **Data format changes** - Do not use the dual-partition upgrade method if you intend to upgrade an application that requires that you change its data format during the application upgrade. The dual-partition upgrade method is not compatible with the extended downtime that is needed to perform data transformation.
- **Location of application software** - Applications must be installed on nonshared storage. Shared storage is not accessible to a partition that is in noncluster mode. Therefore, it is not possible to upgrade application software that is located on shared storage.
- **Division of storage** - Each shared storage device must be connected to a node in each group.
- **Single-node clusters** - Dual-partition upgrade is not available to upgrade a single-node cluster. Use the standard upgrade or live upgrade method instead.
- **Configuration changes** - Do not make cluster configuration changes that are not documented in the upgrade procedures. Such changes might not be propagated to the final cluster configuration. Also, validation attempts of such changes would fail because not all nodes are reachable during a dual-partition upgrade.

Live Upgrade

A live upgrade maintains your previous cluster configuration until you have upgraded all nodes and you commit to the upgrade. If the upgraded configuration causes a problem, you can revert to your previous cluster configuration until you can rectify the problem.

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

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

- **ZFS root file systems** - If your cluster configuration uses a ZFS root file system, you must use only live upgrade to upgrade the Solaris OS. See Solaris documentation for more information.
- **Dual-partition upgrade** - The live upgrade method cannot be used in conjunction with a dual-partition upgrade.
- **Non-global zones** - Unless the cluster is already running on at least Solaris 10 11/06, the live upgrade method does not support the upgrade of clusters that have non-global zones that are configured on any of the cluster nodes. Instead, use the standard upgrade or dual-partition upgrade method.
- **Disk space** - To use the live upgrade method, you must have enough spare disk space available to make a copy of each node's boot environment. You reclaim this disk space after the upgrade is complete and you have verified and committed the upgrade. For information about space requirements for an inactive boot environment, refer to or "Allocating Disk and Swap Space" in *Solaris 10 10/09 Installation Guide: Planning for Installation and Upgrade*.

Rolling Upgrade

In a rolling upgrade, you upgrade software to an update release 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 3.3 release.
- **Solaris upgrade paths** - You can upgrade the Solaris OS only to an update version of the same release. For example, you can perform a rolling upgrade from Solaris 10 5/08 to Solaris 10 10/09. But you cannot perform a rolling upgrade from a version of Solaris 9 to a version of Oracle Solaris 10.
- **ZFS root file systems** - If your cluster configuration uses a ZFS root file system, you cannot use rolling upgrade to upgrade the Solaris OS. You must use only live upgrade to upgrade the Solaris OS. See Solaris documentation for more information.
- **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

Performing a Standard Upgrade to Oracle Solaris Cluster 3.3 Software

This chapter provides the following information to upgrade to Oracle Solaris Cluster 3.3 software by using the standard nonrolling upgrade method:

- "How to Prepare the Cluster for Upgrade (Standard)" on page 20
- "How to Upgrade the Solaris OS and Volume Manager Software (Standard)" on page 27
- "How to Upgrade Oracle Solaris Cluster 3.3 Software (Standard)" on page 32

Note – If your cluster uses a ZFS root file system *and* is configured with zone clusters, you *cannot* use standard upgrade to upgrade the Solaris OS. Use only the live upgrade method to upgrade the Solaris OS. After Solaris is upgraded, you can use standard upgrade to upgrade the Oracle Solaris Cluster software.

Performing a Standard Upgrade of a Cluster

The following table lists the tasks to perform to upgrade to Oracle Solaris Cluster 3.3 software. You also perform these tasks to upgrade only the Solaris OS.

Note – If you upgrade the Solaris OS to a new marketing release, such as from Solaris 9 to Oracle Solaris 10 software, you must also upgrade the Oracle Solaris Cluster software and dependency software to the version that is compatible with the new OS version.

TABLE 2–1 Task Map: Performing a Standard Upgrade to Oracle Solaris Cluster 3.3 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 11 "Choosing an Oracle Solaris Cluster Upgrade Method" on page 13

TABLE 2–1 Task Map: Performing a Standard Upgrade to Oracle Solaris Cluster 3.3 Software *(Continued)*

Task	Instructions
2. If a quorum server is used, upgrade the Quorum Server software.	"How to Upgrade Quorum Server Software" on page 18
3. Remove the cluster from production and back up shared data. If Oracle Solaris Cluster Geographic Edition software is installed, uninstall it.	"How to Prepare the Cluster for Upgrade (Standard)" on page 20
4. Upgrade the Solaris software, if necessary, to a supported Solaris update. If the cluster uses dual-string mediators for Solaris Volume Manager software, unconfigure the mediators. As needed, upgrade Veritas Volume Manager (VxVM) and Veritas File System (VxFS). Solaris Volume Manager software is automatically upgraded with the Solaris OS.	"How to Upgrade the Solaris OS and Volume Manager Software (Standard)" on page 27
5. Upgrade to Oracle Solaris Cluster 3.3 framework and data-service software. If necessary, upgrade applications. If the cluster uses dual-string mediators and you upgraded the Solaris OS, reconfigure the mediators. If you upgraded VxVM, upgrade disk groups.	"How to Upgrade Oracle Solaris Cluster 3.3 Software (Standard)" on page 32
6. Use the `scversions` command to commit the cluster to the upgrade.	"How to Commit the Upgraded Cluster to Oracle Solaris Cluster 3.3 Software" on page 99
7. Verify successful completion of upgrade to Oracle Solaris Cluster 3.3 software.	"How to Verify Upgrade of Oracle Solaris Cluster 3.3 Software" on page 100
8. Enable resources and bring resource groups online. Migrate existing resources to new resource types. Upgrade to Oracle Solaris Cluster Geographic Edition 3.3 software, if used.	"How to Finish Upgrade to Oracle Solaris Cluster 3.3 Software" on page 101
9. (*Optional*) SPARC: Upgrade the Oracle Solaris Cluster module for Sun Management Center, if needed.	"SPARC: How to Upgrade Oracle Solaris Cluster Module Software for Sun Management Center" on page 121

▼ 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 3.3 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.**

   a. **Navigate to the directory where the uninstaller is located.**

      ```
      quorumserver# cd /var/sadm/prod/SUNWentsysver
      ```

      ver The version that is installed on your system.

   b. **Start the uninstallation wizard.**

      ```
      quorumserver# ./uninstall
      ```

   c. **Follow instructions on the screen to uninstall the Quorum Server software from the quorum-server host computer.**

      After removal is finished, you can view any available log. See Chapter 8, "Uninstalling," in *Sun Java Enterprise System 5 Update 1 Installation Guide for UNIX* for additional information about using the `uninstall` program.

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

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

6 **Install the Oracle Solaris Cluster 3.3 Quorum Server software, reconfigure the quorum server, and start the quorum server daemon.**

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

7 **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*.

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

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

▼ How to Prepare the Cluster for Upgrade (Standard)

Perform this procedure to remove the cluster from production before you perform a standard upgrade. 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 11.
- Have available the installation media, documentation, and patches for all software products that you are upgrading, including the following software:
  - Solaris OS
  - Oracle Solaris Cluster 3.3 framework
  - Oracle Solaris Cluster 3.3 patches
  - Oracle Solaris Cluster 3.3 data services (agents)
  - Applications that are managed by Oracle Solaris Cluster 3.3 data services
  - Veritas Volume Manager, if applicable

  See "Patches and Required Firmware Levels" in the Oracle Solaris Cluster 3.3 Release Notes for the location of patches and installation instructions.
- 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 *System Administration Guide: 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.

- On Sun Cluster 3.1 8/05 software, use the following command:

```
phys-schost% scstat
```

- On Sun Cluster 3.2 or Oracle Solaris Cluster 3.3 software, use the following command:

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

See the `scstat(1M)` or `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 Notify users that cluster services will be unavailable during the upgrade.

3 If Geographic Edition software is installed, uninstall it.

For uninstallation procedures, see the documentation for your version of Geographic Edition software.

4 Become superuser on a node of the cluster.

5 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 `scsetup` or `clsetup` utility, perform the following steps:**

a. Start the utility.

- On Sun Cluster 3.1 8/05 software, use the following command:

```
phys-schost# scsetup
```

- On Sun Cluster 3.2 or Oracle Solaris Cluster 3.3 software, use the following command:

```
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 `scsetup` 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.

- On Sun Cluster 3.1 8/05 software, use the following command:

```
phys-schost# scswitch -F -g resource-group
```

`-F`	Switches a resource group offline.
`-g` *resource-group*	Specifies the name of the resource group to take offline.

- On Sun Cluster 3.2 or Oracle Solaris Cluster 3.3 software, use the following command:

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

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

- On Sun Cluster 3.1 8/05 software, use the following command:

```
phys-schost# scrgadm -pv | grep "Res enabled"
(resource-group:resource) Res enabled: True
```

- On Sun Cluster 3.2 or Oracle Solaris Cluster 3.3 software, use the following command:

```
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.

- On Sun Cluster 3.1 8/05 software, use the following command:

  ```
  phys-schost# scswitch -n -j resource
  ```

  `-n` Disables.

  `-j` *resource* Specifies the resource.

- On Sun Cluster 3.2 or Oracle Solaris Cluster 3.3 software, use the following command:

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

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

e. Verify that all resources are disabled.

- On Sun Cluster 3.1 8/05 software, use the following command:

  ```
  phys-schost# scrgadm -pv | grep "Res enabled"
  (resource-group:resource) Res enabled: False
  ```

- On Sun Cluster 3.2 or Oracle Solaris Cluster 3.3 software, use the following command:

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

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

f. Move each resource group to the unmanaged state.

- On Sun Cluster 3.1 8/05 software, use the following command:

  ```
  phys-schost# scswitch -u -g resource-group
  ```

  `-u` Moves the specified resource group to the unmanaged state.

  `-g` *resource-group* Specifies the name of the resource group to move into the unmanaged state.

- On Sun Cluster 3.2 or Oracle Solaris Cluster 3.3 software, use the following command:

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

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

- On Sun Cluster 3.1 8/05 software, use the following command:

  ```
  phys-schost# scstat -g
  ```

- On Sun Cluster 3.2 or Oracle Solaris Cluster 3.3 software, use the following command:

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

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

8 Ensure that all shared data is backed up.

9 If you will upgrade the 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 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 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 Upgrade to Oracle Solaris Cluster 3.3 Software" on page 101.

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

- On Sun Cluster 3.1 8/05 software, use the following command:

  ```
  phys-schost# scswitch -z -D setname -h node
  ```

  `-z` Changes mastery.

  `-D` *devicegroup* Specifies the name of the disk set.

  `-h` *node* Specifies the name of the node to become primary of the disk set.

- On Sun Cluster 3.2 or Oracle Solaris Cluster 3.3 software, use the following command:

```
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.

10 From one node, shut down the cluster.

- On Sun Cluster 3.1 8/05 software, use the following command:

```
phys-schost# scshutdown -g0 -y
```

- On Sun Cluster 3.2 or Oracle Solaris Cluster 3.3 software, use the following command:

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

See the `scshutdown(1M)`man page for more information.

11 Boot each 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)
+-------------------------------------------------------------------------+
| Solaris 10 /sol_10_x86                                                  |
| Solaris failsafe                                                        |
|                                                                         |
+-------------------------------------------------------------------------+
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)
+-------------------------------------------------------------------------+
| root (hd0,0,a)                                                          |
| kernel /platform/i86pc/multiboot                                        |
| module /platform/i86pc/boot_archive                                     |
|+------------------------------------------------------------------------+
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/multiboot -x
```

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

The screen displays the edited command.

```
GNU GRUB version 0.97 (639K lower / 1047488K upper memory)
+-------------------------------------------------------------------------+
| root (hd0,0,a)                                                          |
| kernel /platform/i86pc/multiboot -x                                     |
| module /platform/i86pc/boot_archive                                     |
+-------------------------------------------------------------------------+
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.-
```

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.

---

12 Ensure that each system disk is backed up.

Next Steps Upgrade software on each node.

- To upgrade Solaris software before you perform Oracle Solaris Cluster software upgrade, go to "How to Upgrade the Solaris OS and Volume Manager Software (Standard)" on page 27.

- You must upgrade the Solaris software to a supported release if Oracle Solaris Cluster 3.3 software does not support the release of the Solaris OS that your cluster currently runs . See "Supported Products" in Oracle Solaris Cluster 3.3 Release Notes for more information.
- If Oracle Solaris Cluster 3.3 software supports the release of the Solaris OS that you currently run on your cluster, further Solaris software upgrade is optional.

- Otherwise, upgrade to Oracle Solaris Cluster 3.3 software. Go to "How to Upgrade Oracle Solaris Cluster 3.3 Software (Standard)" on page 32.

How to Upgrade the Solaris OS and Volume Manager Software (Standard)

Perform this procedure on each node in the cluster to upgrade the Solaris OS and optionally also VxVM, if used. Perform all steps from the global zone only. If the cluster already runs on a version of the Solaris OS that supports Oracle Solaris Cluster 3.3 software, further upgrade of the Solaris OS is optional.

If you do not intend to upgrade the Solaris OS or volume management software, proceed to "How to Upgrade Oracle Solaris Cluster 3.3 Software (Standard)" on page 32.

Note – The cluster must already run on, or be upgraded to, at least the minimum required level of the Oracle Solaris OS to support upgrade to Oracle Solaris Cluster 3.3 software. See "Supported Products" in Oracle Solaris Cluster 3.3 Release Notes for more information.

Before You Begin Ensure that all steps in "How to Prepare the Cluster for Upgrade (Standard)" on page 20 are completed.

1. **Become superuser on the cluster node to upgrade.**

   If you are performing a dual-partition upgrade, the node must be a member of the partition that is in noncluster mode.

2. **Determine whether the following Apache run-control scripts exist and are enabled or disabled:**

   ```
   /etc/rc0.d/K16apache
   /etc/rc1.d/K16apache
   /etc/rc2.d/K16apache
   /etc/rc3.d/S50apache
   /etc/rcS.d/K16apache
   ```

   Some applications, such as Oracle Solaris Cluster HA for Apache, require that Apache run control scripts be disabled.

- If these scripts exist and contain an uppercase `K` or `S` in the file name, the scripts are enabled. No further action is necessary for these scripts.
- If these scripts do not exist, in Step 7 you must ensure that any Apache run control scripts that are installed during the Solaris OS upgrade are disabled.
- If these scripts exist but the file names contain a lowercase `k` or `s`, the scripts are disabled. In Step 7 you must ensure that any Apache run control scripts that are installed during the Solaris OS upgrade are disabled.

3 Comment out all entries for globally mounted file systems in the node's `/etc/vfstab` file.

a. For later reference, make a record of all entries that are already commented out.

b. Temporarily comment out all entries for globally mounted file systems in the `/etc/vfstab` file.

Entries for globally mounted file systems contain the `global` mount option. Comment out these entries to prevent the Solaris upgrade from attempting to mount the global devices.

4 Determine which procedure to follow to upgrade the Solaris OS.

- To use Live Upgrade, go instead to Chapter 4, "Performing a Live Upgrade to Oracle Solaris Cluster 3.3 Software."
- To upgrade a cluster that uses Solaris Volume Manager by a method other than Live Upgrade, follow upgrade procedures in Solaris installation documentation.
- To upgrade a cluster that uses Veritas Volume Manager by a method other than Live Upgrade, follow upgrade procedures in Veritas Storage Foundation installation documentation.

Note – If your cluster has VxVM installed and you are upgrading the Solaris OS, you must reinstall or upgrade to VxVM software that is compatible with the version of Oracle Solaris 10 that you upgrade to.

---

5 Upgrade the Solaris software, following the procedure that you selected in Step 4.

Note – *Do not* perform the final reboot instruction in the Solaris software upgrade. Instead, do the following:

a. Return to this procedure to perform Step 6 and Step 7.
b. Reboot into noncluster mode in Step 8 to complete Solaris software upgrade.

- When prompted, choose the manual reboot option.

- When you are instructed to reboot a node during the upgrade process, *always* reboot into noncluster mode. For the `boot` and `reboot` commands, add the `-x` option to the command. The `-x` option ensures that the node reboots into noncluster mode. For example, either of the following two commands boot a node into single-user noncluster mode:

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

```
phys-schost# reboot -- -xs
    or
ok boot -xs
```

If the instruction says to run the `init S` command, use the `reboot -- -xs` command instead.

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

```
phys-schost# shutdown -g -y -i0
Press any key to continue
```

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 0.95 (631K lower / 2095488K upper memory)
+-------------------------------------------------------------------------+
| Solaris 10 /sol_10_x86                                                  |
| Solaris failsafe                                                        |
|                                                                         |
+-------------------------------------------------------------------------+
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 0.95 (615K lower / 2095552K upper memory)
+-------------------------------------------------------------------------+
| root (hd0,0,a)                                                          |
| kernel /platform/i86pc/multiboot                                        |
| module /platform/i86pc/boot_archive                                     |
+-------------------------------------------------------------------------+
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/multiboot -x
```

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

The screen displays the edited command.

```
GNU GRUB version 0.95 (615K lower / 2095552K upper memory)
+----------------------------------------------------------------------+
| root (hd0,0,a)                                                       |
| kernel /platform/i86pc/multiboot -x                                  |
| module /platform/i86pc/boot_archive                                  |
+----------------------------------------------------------------------+
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.-
```

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.

If the instruction says to run the `init S` command, shut down the system then change the GRUB kernel boot command to `/platform/i86pc/multiboot -sx` instead.

6 In the `/a/etc/vfstab` file, uncomment those entries for globally mounted file systems that you commented out in Step 3.

7 If Apache run control scripts were disabled or did not exist before you upgraded the Solaris OS, ensure that any scripts that were installed during Solaris upgrade are disabled.

To disable Apache run control scripts, use the following commands to rename the files with a lowercase `k` or `s`.

```
phys-schost# mv /a/etc/rc0.d/K16apache /a/etc/rc0.d/k16apache
phys-schost# mv /a/etc/rc1.d/K16apache /a/etc/rc1.d/k16apache
phys-schost# mv /a/etc/rc2.d/K16apache /a/etc/rc2.d/k16apache
phys-schost# mv /a/etc/rc3.d/S50apache /a/etc/rc3.d/s50apache
phys-schost# mv /a/etc/rcS.d/K16apache /a/etc/rcS.d/k16apache
```

Alternatively, you can rename the scripts to be consistent with your normal administration practices.

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

Include the double dashes (--) in the following command:

```
phys-schost# reboot -- -x
```

9 **If your cluster runs VxVM and you are upgrading it as well as upgrading the Solaris OS, perform the remaining steps in the procedure to reinstall or upgrade VxVM.**

Make the following changes to the procedure:

- **After VxVM upgrade is complete but *before* you reboot, verify the entries in the `/etc/vfstab` file.**

  If any of the entries that you uncommented in Step 6 were commented out, make those entries uncommented again.

- **If the VxVM procedures instruct you to perform a final reconfiguration reboot, *do not* use the `-r` option alone. Instead, reboot into noncluster mode by using the `-rx` options.**

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

    ```
    phys-schost# reboot -- -rx
    ```

  - **On x86 based systems, perform the shutdown and boot procedures that are described in Step 5 except add `-rx` to the kernel boot command instead of `-sx`.**

Note – If you see a message similar to the following, type the root password to continue upgrade processing. Do *not* run the `fsck` command nor type Ctrl-D.

```
WARNING - Unable to repair the /global/.devices/node@1 filesystem.
Run fsck manually (fsck -F ufs /dev/vx/rdsk/rootdisk_13vol). Exit the
shell when done to continue the boot process.

Type control-d to proceed with normal startup,
(or give root password for system maintenance):       Type the root password
```

---

10 **(Optional) SPARC: Upgrade VxFS.**

Follow procedures that are provided in your VxFS documentation.

11 **Install any required Solaris software patches and hardware-related patches, and download any needed firmware that is contained in the hardware patches.**

Note – Do not reboot after you add patches. Wait to reboot the node until after you upgrade the Oracle Solaris Cluster software.

See "Patches and Required Firmware Levels" in the Oracle Solaris Cluster 3.3 Release Notes for the location of patches and installation instructions.

Next Steps If you are only upgrading the Solaris OS to a Solaris update release and are not upgrading the Oracle Solaris Cluster software, skip to Chapter 6, "Completing the Upgrade."

Otherwise, upgrade to Oracle Solaris Cluster 3.3 software. Go to "How to Upgrade Oracle Solaris Cluster 3.3 Software (Standard)" on page 32.

Note – To complete the upgrade to a new marketing release of the Solaris OS, such as from Solaris 9 to Oracle Solaris 10 software, you must also upgrade the Oracle Solaris Cluster software and dependency software to the version that is compatible with the new version of the OS.

▼ How to Upgrade Oracle Solaris Cluster 3.3 Software (Standard)

Perform this procedure to upgrade each node of the cluster to Oracle Solaris Cluster 3.3 software. You must also perform this procedure after you upgrade to a different marketing release of the Solaris OS, such as from Solaris 9 to Oracle Solaris 10 software.

Perform all steps from the global zone only.

Tip – You can use the `cconsole` utility to perform this procedure on multiple nodes simultaneously. See "How to Install Cluster Control Panel Software on an Administrative Console" in *Oracle Solaris Cluster Software Installation Guide* for more information.

Before You Begin Perform the following tasks:

- Ensure that all steps in "How to Prepare the Cluster for Upgrade (Standard)" on page 20 are completed.
- If you upgraded to a new marketing release of the Solaris OS, such as from Solaris 9 to Oracle Solaris 10 software, ensure that all steps in "How to Upgrade the Solaris OS and Volume Manager Software (Standard)" on page 27 are completed.
- Ensure that you have installed all required Solaris software patches and hardware-related patches.

1 Become superuser on a node of the cluster.

2 Load the installation DVD-ROM into the DVD-ROM drive.

If the volume management daemon `vold(1M)` is running and is configured to manage CD-ROM or DVD devices, the daemon automatically mounts the media on the `/cdrom/cdrom0` directory.

3 **Change to the `/Solaris_`*arch*`/Product/sun_cluster/Solaris_`*ver*`/Tools/` directory, where *arch* is `sparc` or `x86` and where *ver* is `10` for Oracle Solaris 10.**

```
phys-schost# cd /cdrom/cdrom0/Solaris_arch/Product/sun_cluster/Solaris_ver/Tools
```

4 **Start the `scinstall` utility.**

```
phys-schost# ./scinstall
```

Note – Do not use the `/usr/cluster/bin/scinstall` command that is already installed on the node. You must use the `scinstall` command that is located on the installation DVD-ROM.

The `scinstall` Main Menu is displayed.

5 **Choose the menu item, Upgrade This Cluster Node.**

```
*** Main Menu ***

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

       1) Create a new cluster or add a cluster node
       2) Configure a cluster to be JumpStarted from this install server
     * 3) Manage a dual-partition upgrade
     * 4) Upgrade this cluster node
     * 5) Print release information for this cluster node

     * ?) Help with menu options
     * q) Quit

   Option:  4
```

The Upgrade Menu is displayed.

6 **Choose the menu item, Upgrade Oracle Solaris Cluster Framework on This Node.**

7 **Follow the menu prompts to upgrade the cluster framework.**

During the Oracle Solaris Cluster upgrade, `scinstall` might make one or more of the following configuration changes:

- Rename the `ntp.conf` file to `ntp.conf.cluster`, if `ntp.conf.cluster` does not already exist on the node.
- Set the `local-mac-address?` variable to `true`, if the variable is not already set to that value.

Upgrade processing is finished when the system displays the message `Completed Oracle Solaris Cluster framework upgrade` and prompts you to press Enter to continue.

8 **Quit the `scinstall` utility.**

9 **Upgrade data service packages.**

You must upgrade all data services to the Oracle Solaris Cluster 3.3 version.

Note – For HA for SAP Web Application Server, if you are using a J2EE engine resource or a web application server component resource or both, you must delete the resource and recreate it with the new web application server component resource. Changes in the new web application server component resource includes integration of the J2EE functionality. For more information, see *Oracle Solaris Cluster Data Service for SAP Web Application Server Guide*.

a. Start the upgraded interactive `scinstall` utility.

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

Note – Do not use the `scinstall` utility that is on the installation media to upgrade data service packages.

The `scinstall` Main Menu is displayed.

b. Choose the menu item, Upgrade This Cluster Node.

The Upgrade Menu is displayed.

c. Choose the menu item, Upgrade Oracle Solaris Cluster Data Service Agents on This Node.

d. Follow the menu prompts to upgrade Oracle Solaris Cluster data service agents that are installed on the node.

You can choose from the list of data services that are available to upgrade or choose to upgrade all installed data services.

e. When the system displays the message `Completed upgrade of Oracle Solaris Cluster data services agents`, press Enter.

The Upgrade Menu is displayed.

10 Quit the `scinstall` utility.

11 Unload the installation DVD-ROM from the DVD-ROM drive.

a. To ensure that the DVD-ROM is not being used, change to a directory that does *not* reside on the DVD-ROM.

b. Eject the DVD-ROM.

```
phys-schost# eject cdrom
```

12 If you have HA for NFS configured on a highly available local file system, ensure that the loopback file system (LOFS) is disabled.

Note – If you have non-global zones configured, LOFS must remain enabled. For guidelines about using LOFS and alternatives to disabling it, see "Cluster File Systems" in *Oracle Solaris Cluster Software Installation Guide*.

To disable LOFS, ensure that the `/etc/system` file contains the following entry:

```
exclude:lofs
```

This change becomes effective at the next system reboot.

13 As needed, manually upgrade any custom data services that are not supplied on the product media.

14 Verify that each data-service update is installed successfully.

View the upgrade log file that is referenced at the end of the upgrade output messages.

15 Upgrade software applications that are installed on the cluster.

If you want to upgrade VxVM and did not upgrade the Solaris OS, follow procedures in Veritas Storage Foundation installation documentation to upgrade VxVM without upgrading the operating system.

Note – If any upgrade procedure instruct you to perform a reboot, you must add the `-x` option to the boot command. This option boots the cluster into noncluster mode.

Ensure that application levels are compatible with the current versions of Oracle Solaris Cluster and Solaris software. See your application documentation for installation instructions.

16 If you upgraded from Sun Cluster 3.1 8/05 software, reconfigure the private-network address range.

Perform this step if you want to increase or decrease the size of the IP address range that is used by the private interconnect. The IP address range that you configure must minimally support the number of nodes and private networks in the cluster. See "Private Network" in *Oracle Solaris Cluster Software Installation Guide* for more information.

If you also expect to configure zone clusters, you specify that number in "How to Finish Upgrade to Oracle Solaris Cluster 3.3 Software" on page 101, after all nodes are back in cluster mode.

a. From one node, start the `clsetup` utility.

When run in noncluster mode, the `clsetup` utility displays the Main Menu for noncluster-mode operations.

b. Choose the menu item, Change IP Address Range.

The `clsetup` utility displays the current private-network configuration, then asks if you would like to change this configuration.

c. To change either the private-network IP address or the IP address range, type `yes` and press the Return key.

The `clsetup` utility displays the default private-network IP address, `172.16.0.0`, and asks if it is okay to accept this default.

d. Change or accept the private-network IP address.

- **To accept the default private-network IP address and proceed to changing the IP address range, type `yes` and press the Return key.**

  The `clsetup` utility will ask if it is okay to accept the default netmask. Skip to the next step to enter your response.

- **To change the default private-network IP address, perform the following substeps.**

  i. **Type `no` in response to the `clsetup` utility question about whether it is okay to accept the default address, then press the Return key.**

  The `clsetup` utility will prompt for the new private-network IP address.

  ii. **Type the new IP address and press the Return key.**

  The `clsetup` utility displays the default netmask and then asks if it is okay to accept the default netmask.

e. Change or accept the default private-network IP address netmask and range.

The default netmask is `255.255.240.0`. This default IP address range supports up to 64 nodes, up to 10 private networks, and up to 12 zone clusters in the cluster. If you choose to change the netmask, you specify in the following substeps the number of nodes and private networks that you expect in the cluster.

If you also expect to configure zone clusters, you specify that number in "How to Finish Upgrade to Oracle Solaris Cluster 3.3 Software" on page 101, after all nodes are back in cluster mode.

- **To accept the default IP address netmask and range, type `yes` and press the Return key.**

  Then skip to the next step.

- **To change the IP address netmask and range, perform the following substeps.**

  i. **Type `no` in response to the `clsetup` utility's question about whether it is okay to accept the default address range, then press the Return key.**

     When you decline the default netmask, the `clsetup` utility prompts you for the number of nodes and private networks that you expect to configure in the cluster.

  ii. **Enter the number of nodes and private networks that you expect to configure in the cluster.**

     From these numbers, the `clsetup` utility calculates two proposed netmasks:

     - The first netmask is the minimum netmask to support the number of nodes and private networks that you specified.
     - The second netmask supports twice the number of nodes and private networks that you specified, to accommodate possible future growth.

  iii. **Specify either of the calculated netmasks, or specify a different netmask that supports the expected number of nodes and private networks.**

f. **Type `yes` in response to the `clsetup` utility's question about proceeding with the update.**

g. **When finished, exit the `clsetup` utility.**

17 After all nodes in the cluster are upgraded, reboot the upgraded nodes.

a. **Shut down each node.**

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

b. **Boot each node into cluster mode.**

- **On SPARC based systems, do the following:**

```
ok boot
```

- **On x86 based systems, do the following:**

  When the GRUB menu is displayed, select the appropriate Solaris entry and press Enter. The GRUB menu appears similar to the following:

```
GNU GRUB version 0.95 (631K lower / 2095488K upper memory)
+-------------------------------------------------------------------------+
| Solaris 10 /sol_10_x86                                                  |
| Solaris failsafe                                                        |
|                                                                         |
+-------------------------------------------------------------------------+
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*.

Next Steps Go to Chapter 6, "Completing the Upgrade."

◆ ◆ ◆ CHAPTER 3

Performing a Dual-Partition Upgrade to Oracle Solaris Cluster 3.3 Software

This chapter provides the following information to upgrade a multiple-node cluster to Oracle Solaris Cluster 3.3 software by using the dual-partition upgrade method:

- "How to Prepare the Cluster for Upgrade (Dual-Partition)" on page 42
- "How to Upgrade the Solaris OS and Volume Manager Software (Dual-Partition)" on page 50
- "How to Upgrade Oracle Solaris Cluster 3.3 Software (Dual-Partition)" on page 56

Note – If your cluster uses a ZFS root file system *and* is configured with zone clusters, you *cannot* use dual-partition upgrade to upgrade the Solaris OS. Use only the live upgrade method to upgrade the Solaris OS. After Solaris is upgraded, you can use dual-partition upgrade to upgrade the Oracle Solaris Cluster software.

Performing a Dual-Partition Upgrade of a Cluster

The following table lists the tasks to perform to upgrade to Oracle Solaris Cluster 3.3 software. You also perform these tasks to upgrade only the Solaris OS.

Note – If you upgrade the Solaris OS to a new marketing release, such as from Solaris 9 to Oracle Solaris 10 software, you must also upgrade the Oracle Solaris Cluster software and dependency software to the version that is compatible with the new OS version.

TABLE 3–1 Task Map: Performing a Dual-Partition Upgrade to Oracle Solaris Cluster 3.3 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 11 "Choosing an Oracle Solaris Cluster Upgrade Method" on page 13
2. If a quorum server is used, upgrade the Quorum Server software.	"How to Upgrade Quorum Server Software" on page 40
3. If Oracle Solaris Cluster Geographic Edition software is installed, uninstall it. Partition the cluster into two groups of nodes.	"How to Prepare the Cluster for Upgrade (Dual-Partition)" on page 42
4. Upgrade the Solaris software, if necessary, to a supported Solaris update. If the cluster uses dual-string mediators for Solaris Volume Manager software, unconfigure the mediators. As needed, upgrade Veritas Volume Manager (VxVM) and Veritas File System (VxFS). Solaris Volume Manager software is automatically upgraded with the Solaris OS.	"How to Upgrade the Solaris OS and Volume Manager Software (Dual-Partition)" on page 50
5. Upgrade to Oracle Solaris Cluster 3.3 framework and data-service software. If necessary, upgrade applications. If the cluster uses dual-string mediators and you upgraded the Solaris OS, reconfigure the mediators. If you upgraded VxVM, upgrade disk groups.	"How to Upgrade Oracle Solaris Cluster 3.3 Software (Dual-Partition)" on page 56
6. Use the `scversions` command to commit the cluster to the upgrade.	"How to Commit the Upgraded Cluster to Oracle Solaris Cluster 3.3 Software" on page 99
7. Verify successful completion of upgrade to Oracle Solaris Cluster 3.3 software.	"How to Verify Upgrade of Oracle Solaris Cluster 3.3 Software" on page 100
8. Enable resources and bring resource groups online. Optionally, migrate existing resources to new resource types. Upgrade to Oracle Solaris Cluster Geographic Edition 3.3 software, if used.	"How to Finish Upgrade to Oracle Solaris Cluster 3.3 Software" on page 101
9. (*Optional*) SPARC: Upgrade the Oracle Solaris Cluster module for Sun Management Center, if needed.	"SPARC: How to Upgrade Oracle Solaris Cluster Module Software for Sun Management Center" on page 121

▼ 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 3.3 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.**

a. **Navigate to the directory where the uninstaller is located.**

```
quorumserver# cd /var/sadm/prod/SUNWentsysver
```

ver The version that is installed on your system.

b. **Start the uninstallation wizard.**

```
quorumserver# ./uninstall
```

c. Follow instructions on the screen to uninstall the Quorum Server software from the quorum-server host computer.

After removal is finished, you can view any available log. See Chapter 8, "Uninstalling," in *Sun Java Enterprise System 5 Update 1 Installation Guide for UNIX* for additional information about using the `uninstall` program.

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

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

6 Install the Oracle Solaris Cluster 3.3 Quorum Server software, reconfigure the quorum server, and start the quorum server daemon.

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

7 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*.

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

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

▼ How to Prepare the Cluster for Upgrade (Dual-Partition)

Perform this procedure to prepare a multiple-node cluster for a dual-partition upgrade. These procedures will refer to the two groups of nodes as the first partition and the second partition. The nodes that you assign to the second partition will continue cluster services while you upgrade the nodes in the first partition. After all nodes in the first partition are upgraded, you switch cluster services to the first partition and upgrade the second partition. After all nodes in the second partition are upgraded, you boot the nodes into cluster mode to rejoin the nodes from the first partition.

Note – If you are upgrading a single-node cluster, do not use this upgrade method. Instead, go to "How to Prepare the Cluster for Upgrade (Standard)" on page 20 or "How to Prepare the Cluster for Upgrade (Live Upgrade)" on page 70.

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 11.
- Have available the installation media, documentation, and patches for all software products that you are upgrading, including the following software:
  - Solaris OS
  - Oracle Solaris Cluster 3.3 framework
  - Oracle Solaris Cluster 3.3 patches
  - Oracle Solaris Cluster 3.3 data services (agents)
  - Applications that are managed by Oracle Solaris Cluster 3.3 data services
  - Veritas Volume Manager, if applicable

  See "Patches and Required Firmware Levels" in the Oracle Solaris Cluster 3.3 Release Notes for the location of patches and installation instructions.
- 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 *System Administration Guide: 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.

- On Sun Cluster 3.1 8/05 software, use the following command:

```
phys-schost% scstat
```

- On Sun Cluster 3.2 or Oracle Solaris Cluster 3.3 software, use the following command:

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

See the `scstat(1M)` or `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.

4 Ensure that the `RG_system` property of all resource groups in the cluster is set to `FALSE`.

A setting of `RG_system=TRUE` would restrict certain operations that the dual-partition software must perform.

a. On each node, determine whether any resource groups are set to `RG_system=TRUE`.

```
phys-schost# clresourcegroup show -p RG_system
```

Make note of which resource groups to change. Save this list to use when you restore the setting after upgrade is completed.

b. For each resource group that is set to `RG_system=TRUE`, change the setting to `FALSE`.

```
phys-schost# clresourcegroup set -p RG_system=FALSE resourcegroup
```

5 If Geographic Edition software is installed, uninstall it.

For uninstallation procedures, see the documentation for your version of Geographic Edition software.

6 If you will upgrade the 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 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 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 Upgrade to Oracle Solaris Cluster 3.3 Software" on page 101.

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

- On Sun Cluster 3.1 8/05 software, use the following command:

  ```
  phys-schost# scswitch -z -D setname -h node
  ```

`-z`	Changes mastery.
`-D` *devicegroup*	Specifies the name of the disk set.
`-h` *node*	Specifies the name of the node to become primary of the disk set.

- On Sun Cluster 3.2 or Oracle Solaris Cluster 3.3 software, use the following command:

  ```
  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.

7 If you are upgrading a two-node cluster, skip to Step 17.

Otherwise, proceed to Step 8 to determine the partitioning scheme to use. You will determine which nodes each partition will contain, but interrupt the partitioning process. You will then compare the node lists of all resource groups against the node members of each partition in the scheme that you will use. If any resource group does not contain a member of each partition, you must change the node list.

8 Load the installation DVD-ROM into the DVD-ROM drive.

If the volume management daemon `vold(1M)` is running and is configured to manage CD-ROM or DVD devices, the daemon automatically mounts the media on the `/cdrom/cdrom0` directory.

9 Become superuser on a node of the cluster.

10 Change to the `/cdrom/cdrom0/Solaris_`*arch*`/Product/sun_cluster/Solaris_`*ver*`/Tools/` directory, where *arch* is `sparc` or `x86` and where *ver* is `10` for Oracle Solaris 10 .

```
phys-schost# cd /cdrom/cdrom0/Solaris_arch/Product/sun_cluster/Solaris_ver/Tools
```

11 Start the `scinstall` utility in interactive mode.

```
phys-schost# ./scinstall
```

Note – Do not use the `/usr/cluster/bin/scinstall` command that is already installed on the node. You must use the `scinstall` command on the installation DVD-ROM.

The `scinstall` Main Menu is displayed.

12 Choose the menu item, Manage a Dual-Partition Upgrade.

```
*** Main Menu ***

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

        1) Create a new cluster or add a cluster node
        2) Configure a cluster to be JumpStarted from this install server
      * 3) Manage a dual-partition upgrade
      * 4) Upgrade this cluster node
      * 5) Print release information for this cluster node

      * ?) Help with menu options
      * q) Quit

    Option:  3
```

The Manage a Dual-Partition Upgrade Menu is displayed.

13 Choose the menu item, Display and Select Possible Partitioning Schemes.

14 Follow the prompts to perform the following tasks:

a. Display the possible partitioning schemes for your cluster.

b. Choose a partitioning scheme.

c. Choose which partition to upgrade first.

Note – Stop and do *not* respond yet when prompted, `Do you want to begin the dual-partition upgrade?`, but do not exit the `scinstall` utility. You will respond to this prompt in Step 19 of this procedure.

---

15 Make note of which nodes belong to each partition in the partition scheme.

16 On another node of the cluster, become superuser.

17 Ensure that any critical data services can switch over between partitions.

For a two-node cluster, each node will be the only node in its partition.

When the nodes of a partition are shut down in preparation for dual-partition upgrade, the resource groups that are hosted on those nodes switch over to a node in the other partition. If a resource group does not contain a node from each partition in its node list, the resource group cannot switch over. To ensure successful switchover of all critical data services, verify that the node list of the related resource groups contains a member of each upgrade partition.

a. Display the node list of each resource group that you require to remain in service during the entire upgrade.

- On Sun Cluster 3.1 8/05 software, use the following command:

  ```
  phys-schost# scrgadm -pv -g resourcegroup | grep "Res Group Nodelist"
  ```

  `-p` Displays configuration information.

  `-v` Displays in verbose mode.

  `-g` *resourcegroup* Specifies the name of the resource group.

- On Sun Cluster 3.2 or Oracle Solaris Cluster 3.3 software, use the following command:

  ```
  phys-schost# clresourcegroup show -p nodelist
  === Resource Groups and Resources ===

  Resource Group:                                 resourcegroup
    Nodelist:                                        node1 node2
  ...
  ```

b. If the node list of a resource group does not contain at least one member of each partition, redefine the node list to include a member of each partition as a potential primary node.

- On Sun Cluster 3.1 8/05 software, use the following command:

  ```
  phys-schost# scrgadm -a -g resourcegroup -h nodelist
  ```

  `-a` Adds a new configuration.

  `-h` Specifies a comma-separated list of node names.

- On Sun Cluster 3.2 or Oracle Solaris Cluster 3.3 software, use the following command:

  ```
  phys-schost# clresourcegroup add-node -n node resourcegroup
  ```

18 Determine your next step.

- **If you are upgrading a two-node cluster, return to Step 8 through Step 14 to designate your partitioning scheme and upgrade order.**

  When you reach the prompt `Do you want to begin the dual-partition upgrade?`, skip to Step 19.

- **If you are upgrading a cluster with three or more nodes, return to the node that is running the interactive `scinstall` utility.**

  Proceed to Step 19.

19 At the interactive `scinstall` prompt `Do you want to begin the dual-partition upgrade?`, type `Yes`.

The command verifies that a remote installation method is available.

20 When prompted, press Enter to continue each stage of preparation for dual-partition upgrade.

The command switches resource groups to nodes in the second partition, and then shuts down each node in the first partition.

21 After all nodes in the first partition are shut down, boot each node in that partition 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 0.95 (631K lower / 2095488K upper memory)
  +-------------------------------------------------------------------------+
  | Solaris 10 /sol_10_x86                                                  |
  | Solaris failsafe                                                        |
  |                                                                         |
  +-------------------------------------------------------------------------+
  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 0.95 (615K lower / 2095552K upper memory)
  +-------------------------------------------------------------------------+
  | root (hd0,0,a)                                                          |
  | kernel /platform/i86pc/multiboot                                        |
  | module /platform/i86pc/boot_archive                                     |
  +-------------------------------------------------------------------------+
  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/multiboot -x
```

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

The screen displays the edited command.

```
GNU GRUB version 0.95 (615K lower / 2095552K upper memory)
+----------------------------------------------------------------------+
| root (hd0,0,a)                                                       |
| kernel /platform/i86pc/multiboot -x                                  |
| module /platform/i86pc/boot_archive                                  |
+----------------------------------------------------------------------+
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.-
```

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.

---

22 Ensure that each system disk is backed up.

23 If any applications that are running in the second partition are not under control of the Resource Group Manager (RGM), create scripts to halt the applications before you begin to upgrade those nodes.

During dual-partition upgrade processing, these scripts would be called to stop applications such as Oracle Real Application Clusters before the nodes in the second partition are halted.

a. Create the scripts that you need to stop applications that are not under RGM control.

- Create separate scripts for those applications that you want stopped before applications under RGM control are stopped and for those applications that you want stop afterwards.

- To stop applications that are running on more than one node in the partition, write the scripts accordingly.
- Use any name and directory path for your scripts that you prefer.

b. Ensure that each node in the cluster has its own copy of your scripts.

c. On each node, modify the following Oracle Solaris Cluster scripts to call the scripts that you placed on that node.

- `/etc/cluster/ql/cluster_pre_halt_apps` - Use this file to call those scripts that you want to run *before* applications that are under RGM control are shut down.
- `/etc/cluster/ql/cluster_post_halt_apps` - Use this file to call those scripts that you want to run *after* applications that are under RGM control are shut down.

The Oracle Solaris Cluster scripts are issued from one arbitrary node in the partition during post-upgrade processing of the partition. Therefore, ensure that the scripts on any node of the partition will perform the necessary actions for all nodes in the partition.

Next Steps Upgrade software on each node in the first partition.

- To upgrade Solaris software before you perform Oracle Solaris Cluster software upgrade, go to "How to Upgrade the Solaris OS and Volume Manager Software (Dual-Partition)" on page 50.
  - If Oracle Solaris Cluster 3.3 software does not support the release of the Solaris OS that you currently run on your cluster, you must upgrade the Solaris software to a supported release. See "Supported Products" in Oracle Solaris Cluster 3.3 Release Notes for more information.
  - If Oracle Solaris Cluster 3.3 software supports the release of the Solaris OS that you currently run on your cluster, further Solaris software upgrade is optional.
- Otherwise, upgrade to Oracle Solaris Cluster 3.3 software. Go to "How to Upgrade Oracle Solaris Cluster 3.3 Software (Dual-Partition)" on page 56.

▼ How to Upgrade the Solaris OS and Volume Manager Software (Dual-Partition)

Perform this procedure on each node in the cluster to upgrade the Solaris OS and optionally VxVM, if used. Perform all steps from the global zone only.

If the cluster already runs on a version of the Solaris OS that supports Oracle Solaris Cluster 3.3 software, further upgrade of the Solaris OS is optional. If you do not intend to upgrade the Solaris OS or VxVM, proceed to "How to Upgrade Oracle Solaris Cluster 3.3 Software (Standard)" on page 32.

Note – The cluster must already run on, or be upgraded to, at least the minimum required level of the Solaris OS to support upgrade to Oracle Solaris Cluster 3.3 software. See "Supported Products" in Oracle Solaris Cluster 3.3 Release Notes for more information.

Before You Begin Ensure that all steps in "How to Prepare the Cluster for Upgrade (Standard)" on page 20 are completed.

1 **Become superuser on the cluster node to upgrade.**

The node must be a member of the partition that is in noncluster mode.

2 **Determine whether the following Apache run-control scripts exist and are enabled or disabled:**

```
/etc/rc0.d/K16apache
/etc/rc1.d/K16apache
/etc/rc2.d/K16apache
/etc/rc3.d/S50apache
/etc/rcS.d/K16apache
```

Some applications, such as Oracle Solaris Cluster HA for Apache, require that Apache run control scripts be disabled.

- If these scripts exist and contain an uppercase `K` or `S` in the file name, the scripts are enabled. No further action is necessary for these scripts.
- If these scripts do not exist, in Step 7 you must ensure that any Apache run control scripts that are installed during the Solaris OS upgrade are disabled.
- If these scripts exist but the file names contain a lowercase `k` or `s`, the scripts are disabled. In Step 7 you must ensure that any Apache run control scripts that are installed during the Solaris OS upgrade are disabled.

3 **Comment out all entries for globally mounted file systems in the node's `/etc/vfstab` file.**

a. **For later reference, make a record of all entries that are already commented out.**

b. **Temporarily comment out all entries for globally mounted file systems in the `/etc/vfstab` file.**

Entries for globally mounted file systems contain the `global` mount option. Comment out these entries to prevent the Solaris upgrade from attempting to mount the global devices.

4 **Determine which procedure to follow to upgrade the Solaris OS.**

- To use Live Upgrade, go instead to Chapter 4, "Performing a Live Upgrade to Oracle Solaris Cluster 3.3 Software."
- To upgrade a cluster that uses Solaris Volume Manager by a method other than Live Upgrade, follow upgrade procedures in Solaris installation documentation.

- To upgrade a cluster that uses Veritas Volume Manager by a method other than Live Upgrade, follow upgrade procedures in Veritas Storage Foundation installation documentation.

Note – If your cluster has VxVM installed and you are upgrading the Solaris OS, you must reinstall or upgrade to VxVM software that is compatible with the version of Oracle Solaris 10 you upgraded to.

---

5 Upgrade the Solaris software, following the procedure that you selected in Step 4.

a. When prompted, choose the manual reboot option.

b. When prompted to reboot, always reboot into noncluster mode.

Note – *Do not* perform the final reboot instruction in the Solaris software upgrade. Instead, do the following:

a. Return to this procedure to perform Step 6 and Step 7.
b. Reboot into noncluster mode in Step 8 to complete Solaris software upgrade.

Execute the following commands to boot a node into noncluster mode during Solaris upgrade:

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

```
phys-schost# reboot -- -x
     or
ok boot -x
```

If the instruction says to run the `init S` command, use the `reboot -- -xs` command instead.

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

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

Press any key to continue
```

i. 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 0.95 (631K lower / 2095488K upper memory)
+-------------------------------------------------------------------------+
| Solaris 10 /sol_10_x86                                                  |
| Solaris failsafe                                                        |
|                                                                         |
+-------------------------------------------------------------------------+
```

```
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*.

ii. 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 0.95 (615K lower / 2095552K upper memory)
+-------------------------------------------------------------------------+
| root (hd0,0,a)                                                          |
| kernel /platform/i86pc/multiboot                                        |
| module /platform/i86pc/boot_archive                                     |
+-------------------------------------------------------------------------+
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.
```

iii. 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/multiboot -x
```

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

The screen displays the edited command.

```
GNU GRUB version 0.95 (615K lower / 2095552K upper memory)
+-------------------------------------------------------------------------+
| root (hd0,0,a)                                                          |
| kernel /platform/i86pc/multiboot -x                                     |
| module /platform/i86pc/boot_archive                                     |
+-------------------------------------------------------------------------+
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.-
```

v. Type `b` to boot the node into noncluster mode.

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

If the instruction says to run the `init S` command, shut down the system then change the GRUB kernel boot command to `/platform/i86pc/multiboot -sx` instead.

6 **In the `/a/etc/vfstab` file, uncomment those entries for globally mounted file systems that you commented out in Step 3.**

7 **If Apache run control scripts were disabled or did not exist before you upgraded the Solaris OS, ensure that any scripts that were installed during Solaris upgrade are disabled.**

To disable Apache run control scripts, use the following commands to rename the files with a lowercase `k` or `s`.

```
phys-schost# mv /a/etc/rc0.d/K16apache /a/etc/rc0.d/k16apache
phys-schost# mv /a/etc/rc1.d/K16apache /a/etc/rc1.d/k16apache
phys-schost# mv /a/etc/rc2.d/K16apache /a/etc/rc2.d/k16apache
phys-schost# mv /a/etc/rc3.d/S50apache /a/etc/rc3.d/s50apache
phys-schost# mv /a/etc/rcS.d/K16apache /a/etc/rcS.d/k16apache
```

Alternatively, you can rename the scripts to be consistent with your normal administration practices.

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

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

  Include the double dashes (--) in the command:

  ```
  phys-schost# reboot -- -x
  ```

- **On x86 based systems, perform the shutdown and boot procedures that are described in Step 5 except add `-x` to the kernel boot command instead of `-sx`.**

9 **If your cluster runs VxVM and you are upgrading it as well as upgrading the Solaris OS, perform the remaining steps in the procedure to reinstall or upgrade VxVM.**

Make the following changes to the procedure:

- **After VxVM upgrade is complete but before you reboot, verify the entries in the `/etc/vfstab` file.**

  If any of the entries that you uncommented in Step 6 were commented out, make those entries uncommented again.

- **If the VxVM procedures instruct you to perform a final reconfiguration reboot, *do not* use the `-r` option alone. Instead, reboot into noncluster mode by using the `-rx` options.**

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

    ```
    phys-schost# reboot -- -rx
    ```

- **On x86 based systems, perform the shutdown and boot procedures that are described in Step 5 except add `-rx` to the kernel boot command instead of `-sx`.**

Note – If you see a message similar to the following, type the root password to continue upgrade processing. Do *not* run the `fsck` command nor type Ctrl-D.

```
WARNING - Unable to repair the /global/.devices/node@1 filesystem.
Run fsck manually (fsck -F ufs /dev/vx/rdsk/rootdisk_13vol). Exit the
shell when done to continue the boot process.

Type control-d to proceed with normal startup,
(or give root password for system maintenance):       Type the root password
```

---

10 (Optional) SPARC: Upgrade VxFS.

Follow procedures that are provided in your VxFS documentation.

11 Install any required Solaris software patches and hardware-related patches, and download any needed firmware that is contained in the hardware patches.

Note – Do not reboot after you add patches. Wait to reboot the node until after you upgrade the Oracle Solaris Cluster software.

See "Patches and Required Firmware Levels" in the Oracle Solaris Cluster 3.3 Release Notes for the location of patches and installation instructions.

Next Steps If you are already running Oracle Solaris Cluster 3.3 software and only upgrading the Oracle Solaris 10 OS to an Oracle Solaris 10 update release, you do not need to upgrade the Oracle Solaris Cluster software. Go to Chapter 6, "Completing the Upgrade."

Otherwise, upgrade to Oracle Solaris Cluster 3.3 software. Go to "How to Upgrade Oracle Solaris Cluster 3.3 Software (Dual-Partition)" on page 56.

Note – To complete the upgrade to a new marketing release of the Solaris OS, such as from Solaris 9 to Oracle Solaris 10 software, you must also upgrade the Oracle Solaris Cluster software and dependency software to the version that is compatible with the new version of the Solaris OS.

▼ How to Upgrade Oracle Solaris Cluster 3.3 Software (Dual-Partition)

Perform this procedure to upgrade each node of the cluster to Oracle Solaris Cluster 3.3 software. You must also perform this procedure after you upgrade to a different marketing release of the Solaris OS, such as from Solaris 9 to Oracle Solaris 10 software.

Perform all steps from the global zone only.

Tip – You can use the `cconsole` utility to perform this procedure on multiple nodes simultaneously. See "How to Install Cluster Control Panel Software on an Administrative Console" in *Oracle Solaris Cluster Software Installation Guide* for more information.

Before You Begin Perform the following tasks:

- Ensure that all steps in "How to Prepare the Cluster for Upgrade (Dual-Partition)" on page 42 are completed.
- Ensure that the node you are upgrading belongs to the partition that is not active in the cluster and that the node is in noncluster mode.
- If you upgraded to a new marketing release of the Solaris OS, such as from Solaris 9 to Oracle Solaris 10 software, ensure that all steps in "How to Upgrade the Solaris OS and Volume Manager Software (Dual-Partition)" on page 50 are completed.
- Ensure that you have installed all required Solaris software patches and hardware-related patches.

1 Become superuser on a node that is a member of the partition that is in noncluster mode.

2 Load the installation DVD-ROM into the DVD-ROM drive.

If the volume management daemon `vold(1M)` is running and is configured to manage CD-ROM or DVD devices, the daemon automatically mounts the media on the `/cdrom/cdrom0` directory.

3 Change to the `/Solaris_`*arch*`/Product/sun_cluster/Solaris_`*ver*`/Tools/` directory, where *arch* is `sparc` or `x86` and where *ver* is `10` for Oracle Solaris 10.

```
phys-schost# cd /cdrom/cdrom0/Solaris_arch/Product/sun_cluster/Solaris_ver/Tools
```

4 Start the `scinstall` utility.

```
phys-schost# ./scinstall
```

Note – Do not use the `/usr/cluster/bin/scinstall` command that is already installed on the node. You must use the `scinstall` command that is located on the installation DVD-ROM.

The `scinstall` Main Menu is displayed.

5 Choose the menu item, Upgrade This Cluster Node.

```
*** Main Menu ***

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

       1) Create a new cluster or add a cluster node
       2) Configure a cluster to be JumpStarted from this install server
     * 3) Manage a dual-partition upgrade
     * 4) Upgrade this cluster node
     * 5) Print release information for this cluster node

     * ?) Help with menu options
     * q) Quit

   Option:  4
```

The Upgrade Menu is displayed.

6 Choose the menu item, Upgrade Oracle Solaris Cluster Framework on This Node.

7 Follow the menu prompts to upgrade the cluster framework.

During the Oracle Solaris Cluster upgrade, `scinstall` might make one or more of the following configuration changes:

- Rename the `ntp.conf` file to `ntp.conf.cluster`, if `ntp.conf.cluster` does not already exist on the node.
- Set the `local-mac-address?` variable to `true`, if the variable is not already set to that value.

Upgrade processing is finished when the system displays the message `Completed Oracle Solaris Cluster framework upgrade` and prompts you to press Enter to continue.

8 Quit the `scinstall` utility.

9 Upgrade data service packages.

You must upgrade all data services to the Oracle Solaris Cluster 3.3 version.

Note – For HA for SAP Web Application Server, if you are using a J2EE engine resource or a web application server component resource or both, you must delete the resource and recreate it with the new web application server component resource. Changes in the new web application server component resource includes integration of the J2EE functionality. For more information, see *Oracle Solaris Cluster Data Service for SAP Web Application Server Guide*.

a. Start the upgraded interactive `scinstall` utility.

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

Note – Do not use the `scinstall` utility that is on the installation media to upgrade data service packages.

The `scinstall` Main Menu is displayed.

b. Choose the menu item, Upgrade This Cluster Node.

The Upgrade Menu is displayed.

c. Choose the menu item, Upgrade Oracle Solaris Cluster Data Service Agents on This Node.

d. Follow the menu prompts to upgrade Oracle Solaris Cluster data service agents that are installed on the node.

You can choose from the list of data services that are available to upgrade or choose to upgrade all installed data services.

e. When the system displays the message `Completed upgrade of Oracle Solaris Cluster data services agents`, press Enter.

The Upgrade Menu is displayed.

10 Quit the `scinstall` utility.

11 Unload the installation DVD-ROM from the DVD-ROM drive.

a. To ensure that the DVD-ROM is not being used, change to a directory that does *not* reside on the DVD-ROM.

b. Eject the DVD-ROM.

```
phys-schost# eject cdrom
```

12 If you have Oracle Solaris Cluster HA for NFS configured on a highly available local file system, ensure that the loopback file system (LOFS) is disabled.

Note – If you have non-global zones configured, LOFS must remain enabled. For guidelines about using LOFS and alternatives to disabling it, see "Cluster File Systems" in *Oracle Solaris Cluster Software Installation Guide*.

To disable LOFS, ensure that the `/etc/system` file contains the following entry:

```
exclude:lofs
```

This change becomes effective at the next system reboot.

13 As needed, manually upgrade any custom data services that are not supplied on the product media.

14 Verify that each data-service update is installed successfully.

View the upgrade log file that is referenced at the end of the upgrade output messages.

15 Upgrade software applications that are installed on the cluster.

Ensure that application levels are compatible with the current versions of Oracle Solaris Cluster and Solaris software. See your application documentation for installation instructions.

If you want to upgrade VxVM and did not upgrade the Solaris OS, follow procedures in Veritas Storage Foundation installation documentation to upgrade VxVM without upgrading the operating system.

Note – If any upgrade procedure instruct you to perform a reboot, you must add the `-x` option to the boot command. This option boots the cluster into noncluster mode.

---

16 Repeat all steps in this procedure up to this point on all remaining nodes that you need to upgrade in the partition.

17 After all nodes in a partition are upgraded, apply the upgrade changes.

a. From one node in the partition that you are upgrading, start the interactive `scinstall` utility.

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

Note – Do not use the `scinstall` command that is located on the installation media. Only use the `scinstall` command that is located on the cluster node.

The `scinstall` Main Menu is displayed.

b. Type option number for Apply Dual-Partition Upgrade Changes to the Partition.

c. Follow the prompts to continue each stage of the upgrade processing.

The command performs the following tasks, depending on which partition the command is run from:

- **First partition** - The command halts each node in the second partition, one node at a time. When a node is halted, any services on that node are automatically switched over to a node in the first partition, provided that the node list of the related resource group

contains a node in the first partition. After all nodes in the second partition are halted, the nodes in the first partition are booted into cluster mode and take over providing cluster services.

Caution – Do not reboot any node of the first partition again until after the upgrade is completed on all nodes. If you again reboot a node of the first partition before the second partition is upgraded and rebooted into the cluster, the upgrade might fail in an unrecoverable state.

- **Second partition** - The command boots the nodes in the second partition into cluster mode, to join the active cluster that was formed by the first partition. After all nodes have rejoined the cluster, the command performs final processing and reports on the status of the upgrade.

d. Exit the `scinstall` utility, if it is still running.

e. If you are finishing upgrade from Sun Cluster 3.1 8/05 software of the first partition and you want to configure zone clusters, set the expected number of nodes and private networks in the cluster.

If you upgraded from Sun Cluster 3.1 8/05 software and do not want to configure zone clusters, or if you upgraded from Sun Cluster 3.2 software, this task is optional.

i. Boot all nodes in the first partition into noncluster mode.

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

```
ok boot -x
```

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

In the GRUB menu, use the arrow keys to select the appropriate 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)
+-------------------------------------------------------------------------+
| Solaris 10 /sol_10_x86                                                  |
| Solaris failsafe                                                        |
|                                                                         |
+-------------------------------------------------------------------------+
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*.

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)
+-------------------------------------------------------------------------+
| root (hd0,0,a)                                                          |
| kernel /platform/i86pc/multiboot                                        |
| module /platform/i86pc/boot_archive                                     |
|+------------------------------------------------------------------------+
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.
```

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/multiboot -x
```

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

The screen displays the edited command.

```
GNU GRUB version 0.97 (639K lower / 1047488K upper memory)
+-------------------------------------------------------------------------+
| root (hd0,0,a)                                                          |
| kernel /platform/i86pc/multiboot -x                                     |
| module /platform/i86pc/boot_archive                                     |
+-------------------------------------------------------------------------+
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.-
```

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.

---

ii. From one node, start the `clsetup` utility.

When run in noncluster mode, the `clsetup` utility displays the Main Menu for noncluster-mode operations.

iii. Choose the menu item, Change IP Address Range.

The `clsetup` utility displays the current private-network configuration, then asks if you would like to change this configuration.

iv. To change either the private-network IP address or the IP address range, type `yes` and press the Return key.

The `clsetup` utility displays the default private-network IP address, `172.16.0.0`, and asks if it is okay to accept this default.

v. Change or accept the private-network IP address.

- **To accept the default private-network IP address and proceed to changing the IP address range, type `yes` and press the Return key.**

  The `clsetup` utility will ask if it is okay to accept the default netmask. Skip to the next step to enter your response.

- **To change the default private-network IP address, perform the following substeps.**

  Type `no` in response to the `clsetup` utility question about whether it is okay to accept the default address, then press the Return key.

  The `clsetup` utility will prompt for the new private-network IP address.

  Type the new IP address and press the Return key.

  The `clsetup` utility displays the default netmask and then asks if it is okay to accept the default netmask.

vi. Change or accept the default private-network IP address netmask and range.

The default netmask is `255.255.240.0`. This default IP address range supports up to 64 nodes, up to 10 private networks, and up to 12 zone clusters in the cluster. If you choose to change the netmask, you specify in the following substeps the number of nodes and private networks that you expect in the cluster.

If you also expect to configure zone clusters, you specify that number in "How to Finish Upgrade to Oracle Solaris Cluster 3.3 Software" on page 101, after all nodes are back in cluster mode.

- **To accept the default IP address netmask and range, type `yes` and press the Return key.**

  Then skip to the next step.

- **To change the IP address netmask and range, perform the following substeps.**

  Type `no` in response to the `clsetup` utility's question about whether it is okay to accept the default address range, then press the Return key.

  When you decline the default netmask, the `clsetup` utility prompts you for the number of nodes and private networks that you expect to configure in the cluster.

  Enter the number of nodes and private networks that you expect to configure in the cluster.

  From these numbers, the `clsetup` utility calculates two proposed netmasks:

  - The first netmask is the minimum netmask to support the number of nodes and private networks that you specified.
  - The second netmask supports twice the number of nodes and private networks that you specified, to accommodate possible future growth.

  Specify either of the calculated netmasks, or specify a different netmask that supports the expected number of nodes and private networks.

vii. Type `yes` in response to the `clsetup` utility's question about proceeding with the update.

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

ix. Boot the nodes of the first partition into cluster mode.

f. If you are finishing upgrade of the first partition, perform the following substeps to prepare the second partition for upgrade.

Otherwise, if you are finishing upgrade of the second partition, proceed to "How to Verify Upgrade of Oracle Solaris Cluster 3.3 Software" on page 100.

i. Boot each node in the second partition into noncluster mode.

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

  ```
  ok boot -x
  ```

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

  In the GRUB menu, use the arrow keys to select the appropriate 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)
+-------------------------------------------------------------------------+
| Solaris 10 /sol_10_x86                                                  |
| Solaris failsafe                                                        |
|                                                                         |
+-------------------------------------------------------------------------+
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*.

  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)
+-------------------------------------------------------------------------+
| root (hd0,0,a)                                                          |
| kernel /platform/i86pc/multiboot                                        |
| module /platform/i86pc/boot_archive                                     |
|+------------------------------------------------------------------------+
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.
```

  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/multiboot -x
```

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

  The screen displays the edited command.

```
GNU GRUB version 0.97 (639K lower / 1047488K upper memory)
+-------------------------------------------------------------------------+
| root (hd0,0,a)                                                          |
| kernel /platform/i86pc/multiboot -x                                     |
| module /platform/i86pc/boot_archive                                     |
+-------------------------------------------------------------------------+
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.-
```

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.

---

ii. Upgrade the nodes in the second partition.

To upgrade Solaris software before you perform Oracle Solaris Cluster software upgrade, go to "How to Upgrade the Solaris OS and Volume Manager Software (Dual-Partition)" on page 50.

Otherwise, upgrade Oracle Solaris Cluster software on the second partition. Return to Step 1.

18 If you changed the `RG_system` property of any resource groups to `FALSE`, change the settings back to `TRUE`.

```
phys-schost# clresourcegroup set -p RG_system=TRUE resourcegroup
```

Next Steps Go to Chapter 6, "Completing the Upgrade."

Troubleshooting If you experience an unrecoverable error during dual-partition upgrade, perform recovery procedures in "How to Recover from a Failed Dual-Partition Upgrade" on page 111.

◆ ◆ ◆ CHAPTER 4

Performing a Live Upgrade to Oracle Solaris Cluster 3.3 Software

This chapter provides the following information to upgrade to Oracle Solaris Cluster 3.3 software by using the live upgrade method:

- "How to Prepare the Cluster for Upgrade (Live Upgrade)" on page 70
- "How to Upgrade the Solaris OS and Oracle Solaris Cluster 3.3 Software (Live Upgrade)" on page 71

If your cluster configuration uses a ZFS root file system *and* is configured with zone clusters, you can use live upgrade *only* to upgrade the Solaris OS. To upgrade Oracle Solaris Cluster software, after using live upgrade to upgrade Solaris software, use either standard upgrade or dual-partition upgrade to upgrade Oracle Solaris Cluster software.

Performing a Live Upgrade of a Cluster

The following table lists the tasks to perform to upgrade to Oracle Solaris Cluster 3.3 software. You also perform these tasks to upgrade only the Solaris OS.

Note – If you upgrade the Solaris OS to a new marketing release, such as from Solaris 9 to Oracle Solaris 10 software, you must also upgrade the Oracle Solaris Cluster software and dependency software to the version that is compatible with the new OS version.

TABLE 4–1 Task Map: Performing a Live Upgrade to Oracle Solaris Cluster 3.3 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 11 "Choosing an Oracle Solaris Cluster Upgrade Method" on page 13

TABLE 4–1 Task Map: Performing a Live Upgrade to Oracle Solaris Cluster 3.3 Software *(Continued)*

Task	Instructions
2. If a quorum server is used, upgrade the Quorum Server software.	"How to Upgrade Quorum Server Software" on page 68
3. If Oracle Solaris Cluster Geographic Edition software is installed, uninstall it.	"How to Prepare the Cluster for Upgrade (Live Upgrade)" on page 70
4. If the cluster uses dual-string mediators for Solaris Volume Manager software, unconfigure the mediators. Upgrade the Solaris software, if necessary, to a supported Solaris update. Upgrade to Oracle Solaris Cluster 3.3 framework and data-service software. If necessary, upgrade applications. If the cluster uses dual-string mediators, reconfigure the mediators. As needed, upgrade Veritas Volume Manager (VxVM)software and disk groups and Veritas File System (VxFS).	"How to Upgrade the Solaris OS and Oracle Solaris Cluster 3.3 Software (Live Upgrade)" on page 71
5. Use the `scversions` command to commit the cluster to the upgrade.	"How to Commit the Upgraded Cluster to Oracle Solaris Cluster 3.3 Software" on page 99
6. Verify successful completion of upgrade to Oracle Solaris Cluster 3.3 software.	"How to Verify Upgrade of Oracle Solaris Cluster 3.3 Software" on page 100
7. Enable resources and bring resource groups online. Migrate existing resources to new resource types. Upgrade to Oracle Solaris Cluster Geographic Edition 3.3 software, if used.	"How to Finish Upgrade to Oracle Solaris Cluster 3.3 Software" on page 101
8. (*Optional*) SPARC: Upgrade the Oracle Solaris Cluster module for Sun Management Center, if needed.	"SPARC: How to Upgrade Oracle Solaris Cluster Module Software for Sun Management Center" on page 121

▼ 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 3.3 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.

a. Navigate to the directory where the uninstaller is located.

```
quorumserver# cd /var/sadm/prod/SUNWentsysver
```

ver The version that is installed on your system.

b. Start the uninstallation wizard.

```
quorumserver# ./uninstall
```

c. Follow instructions on the screen to uninstall the Quorum Server software from the quorum-server host computer.

After removal is finished, you can view any available log. See Chapter 8, "Uninstalling," in *Sun Java Enterprise System 5 Update 1 Installation Guide for UNIX* for additional information about using the `uninstall` program.

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

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

6 Install the Oracle Solaris Cluster 3.3 Quorum Server software, reconfigure the quorum server, and start the quorum server daemon.

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

7 **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*.

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

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

▼ How to Prepare the Cluster for Upgrade (Live Upgrade)

Perform this procedure to prepare a cluster for live upgrade.

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 11.
- Have available the installation media, documentation, and patches for all software products that you are upgrading, including the following software:
  - Solaris OS
  - Oracle Solaris Cluster 3.3 framework
  - Oracle Solaris Cluster 3.3 patches
  - Oracle Solaris Cluster 3.3 data services (agents)
  - Applications that are managed by Oracle Solaris Cluster 3.3 data services
  - Veritas Volume Manager, if applicable

  See "Patches and Required Firmware Levels" in the Oracle Solaris Cluster 3.3 Release Notes for the location of patches and installation instructions.
- 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 *System Administration Guide: 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.**

      - On Sun Cluster 3.1 8/05 software, use the following command:

        ```
        phys-schost% scstat
        ```

      - On Sun Cluster 3.2 or Oracle Solaris Cluster 3.3 software, use the following command:

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

      See the `scstat(1M)` or `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 will 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. **If Geographic Edition software is installed, uninstall it.**

   For uninstallation procedures, see the documentation for your version of Geographic Edition software.

4. **Become superuser on a node of the cluster.**

5. **Ensure that all shared data is backed up.**

6. **Ensure that each system disk is backed up.**

Next Steps Perform a live upgrade of the Solaris OS, Oracle Solaris Cluster 3.3 software, and other software. Go to "How to Upgrade the Solaris OS and Oracle Solaris Cluster 3.3 Software (Live Upgrade)" on page 71.

▼ How to Upgrade the Solaris OS and Oracle Solaris Cluster 3.3 Software (Live Upgrade)

Perform this procedure to upgrade the Solaris OS, volume-manager software, and Oracle Solaris Cluster software by using the live upgrade method. The Oracle Solaris Cluster live

upgrade method uses the Solaris Live Upgrade feature. For information about live upgrade of the Solaris OS, refer to the following Solaris documentation:

- *Solaris 10 10/09 Installation Guide: Solaris Live Upgrade and Upgrade Planning*
- If non-global zones are installed on the cluster, see Chapter 8, "Upgrading the Solaris OS on a System With Non-Global Zones Installed," in *Solaris 10 10/09 Installation Guide: Solaris Live Upgrade and Upgrade Planning*.

Note – The cluster must already run on, or be upgraded to, at least the minimum required level of the Solaris OS to support upgrade to Oracle Solaris Cluster 3.3 software. See "Supported Products" in Oracle Solaris Cluster 3.3 Release Notes for more information.

Perform this procedure on each node in the cluster.

Tip – You can use the `cconsole` utility to perform this procedure on multiple nodes simultaneously. See "How to Install Cluster Control Panel Software on an Administrative Console" in *Oracle Solaris Cluster Software Installation Guide* for more information.

Before You Begin

- Ensure that all steps in "How to Prepare the Cluster for Upgrade (Live Upgrade)" on page 70 are completed.

1 Install a supported version of Solaris Live Upgrade software.

Follow instructions in "Solaris Live Upgrade System Requirements" in *Solaris 10 10/09 Installation Guide: Solaris Live Upgrade and Upgrade Planning* and "Installing Solaris Live Upgrade" in *Solaris 10 10/09 Installation Guide: Solaris Live Upgrade and Upgrade Planning*.

2 If you will upgrade the 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 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 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 Upgrade to Oracle Solaris Cluster 3.3 Software" on page 101.

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

- On Sun Cluster 3.1 8/05 software, use the following command:

  ```
  phys-schost# scswitch -z -D setname -h node
  ```

`-z`	Changes mastery.
`-D` *devicegroup*	Specifies the name of the disk set.
`-h` *node*	Specifies the name of the node to become primary of the disk set.

- On Sun Cluster 3.2 or Oracle Solaris Cluster 3.3 software, use the following command:

  ```
  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.

3 On each node that uses a UFS root file system, temporarily change the name of the global devices entry in the `/etc/vfstab` file from the DID name to the physical name.

This name change is necessary for live upgrade software to recognize the global-devices file system. You will restore the DID names after the live upgrade is completed.

a. Back up the `/etc/vfstab` file.

```
phys-schost# cp /etc/vfstab /etc/vfstab.old
```

b. Open the `/etc/vfstab` file for editing.

c. Locate and edit the line that corresponds to `/global/.device/`*node*`@N`.

- Change the DID names to the physical names by changing `/dev/did/{r}dsk/d`*Y*`s`*Z* to `/dev/{r}dsk/c`*N*`t`*X*`d`*Y*`s`*Z*.
- Remove `global` from the entry.

The following example shows the names of DID device `d3s3`, which corresponds to `/global/.devices/node@2`, changed to its physical device names and the global entry removed:

Original:

```
/dev/did/dsk/d3s3  /dev/did/rdsk/d3s3  /global/.devices/node@2  ufs  2  no  global
```

Changed:

```
dev/dsk/c0t0d0s3   /dev/rdsk/c0t0d0s3  /global/.devices/node@2  ufs  2  no  -
```

d. Temporarily comment out any entries for highly available local file systems that are managed by HAStoragePlus.

4 Build an inactive boot environment (BE).

```
phys-schost# lucreate options-n BE-name
```

`-n` *BE-name* Specifies the name of the boot environment that is to be upgraded.

For information about important options to the `lucreate` command, see *Solaris 10 10/09 Installation Guide: Solaris Live Upgrade and Upgrade Planning* and the `lucreate`(1M) man page.

5 If necessary, upgrade the Solaris OS software in your inactive BE.

If the cluster already runs on a properly patched version of the Solaris OS that supports Oracle Solaris Cluster 3.3 software, this step is optional.

- **If you use Solaris Volume Manager software, run the following command:**

  ```
  phys-schost# luupgrade -u -n BE-name -s os-image-path
  ```

  `-u` Upgrades an operating system image on a boot environment.

  `-s` *os-image-path* Specifies the path name of a directory that contains an operating system image.

- **If you use Veritas Volume Manager, follow live upgrade procedures in your Veritas Storage Foundation installation documentation for upgrading the operating system.**

6 Mount your inactive BE by using the `lumount` command.

```
phys-schost# lumount -n BE-name -m BE-mount-point
```

`-m` *BE-mount-point* Specifies the mount point of *BE-name*.

For more information, see *Solaris 10 10/09 Installation Guide: Solaris Live Upgrade and Upgrade Planning* and the `lumount`(1M) man page.

7 Apply any necessary Solaris patches.

You might need to patch your Solaris software to use Solaris Live Upgrade. For details about the patches that the Solaris OS requires and where to download them, see "Managing Packages and Patches With Solaris Live Upgrade" in *Solaris 9 9/04 Installation Guide* or "Upgrading a System With Packages or Patches" in *Solaris 10 10/09 Installation Guide: Solaris Live Upgrade and Upgrade Planning*.

8 If necessary, upgrade your VxVM software.

Refer to your Veritas Storage Foundation installation documentation for procedures to use the live upgrade method.

9 (Optional) SPARC: Upgrade VxFS.

Follow procedures that are provided in your VxFS documentation.

10 If your cluster hosts software applications that require an upgrade and that you can upgrade by using Solaris Live Upgrade, upgrade those software applications.

However, if some software applications to upgrade cannot use Solaris Live Upgrade, such as Sun QFS software, upgrade the applications in Step 25.

11 Load the installation DVD-ROM into the DVD-ROM drive.

If the volume management daemon `vold(1M)` is running and is configured to manage CD-ROM or DVD devices, the daemon automatically mounts the media on the `/cdrom/cdrom0` directory.

12 Change to the `/cdrom/cdrom0/Solaris_`*arch*`/Product/sun_cluster/Solaris_`*ver*`/Tools/` directory, where *arch* is `sparc` or `x86` and where *ver* is `10` for Oracle Solaris 10 .

```
phys-schost# cd /cdrom/cdrom0/Solaris_arch/Product/sun_cluster/Solaris_ver/Tools
```

13 Upgrade Oracle Solaris Cluster software.

```
phys-schost# ./scinstall -u update -R BE-mount-point
```

`-u update`	Specifies that you are performing an upgrade of Oracle Solaris Cluster software.
`-R` *BE-mount-point*	Specifies the mount point for your alternate boot environment.

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

14 Upgrade data services.

```
phys-schost# BE-mount-point/usr/cluster/bin/scinstall -u update -s all  \
-d /cdrom/cdrom0/Solaris_arch/Product/sun_cluster_agents -R BE-mount-point
```

15 Unload the installation DVD-ROM from the DVD-ROM drive.

a. To ensure that the DVD-ROM is not being used, change to a directory that does *not* reside on the DVD-ROM.

b. Eject the DVD-ROM.

```
phys-schost# eject cdrom
```

16 Repeat all steps, starting from Step 1, on each node in the cluster.

Note – Do not reboot any node until all nodes in the cluster are upgraded on their inactive BE.

---

17 On each cluster node that uses a UFS root file system, restore the DID names of the global-devices entry in the `/etc/vfstab` file.

a. On the current, unupgraded BE, restore the original `/etc/vfstab` file.

```
phys-schost# cp /etc/vstab.old /etc/vfstab
```

b. In the alternate BE, open the `/etc/vfstab` file for editing.

c. Locate the line that corresponds to `/global/.devices/`*node*`@`*N* and replace the dash (-) at the end of the entry with the word `global`.

```
/dev/dsk/cNtXdYsZ /dev/rdsk/cNtXdYsZ /global/.devices/node@N ufs 2 no global
```

When the node is rebooted into the upgraded alternate BE, the DID names are substituted in the `/etc/vfstab` file automatically.

d. Uncomment the entries for highly available local file systems that you commented out in Step 3.

18 On each node, unmount the inactive BE.

```
phys-schost# luumount -n BE-name
```

19 On each node, activate the upgraded inactive BE.

```
phys-schost# luactivate BE-name
```

BE-name The name of the alternate BE that you built in Step 4.

20 Shut down each node in the cluster.

Note – Do **not** use the `reboot` or `halt` command. These commands do not activate a new BE.

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

21 Determine your next step.

- **If one of the following conditions applies, skip to Step 23 to boot each node into *noncluster* mode.**
  - You upgraded from Sun Cluster 3.1 8/05 software and you want to configure zone clusters.
  - Your cluster hosts software applications that require upgrade and for which you cannot use Solaris Live Upgrade.
  - *(Optional)* You want to change the private-network IP address range.
- **If you have no additional software to upgrade, go to Step 22 to boot each node into cluster mode.**

22 To complete the upgrade, boot each node into cluster mode.

Ensure that all nodes in the cluster are shut down before you boot nodes into cluster mode.

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

```
ok boot
```

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

  When the GRUB menu is displayed, select the appropriate Solaris entry and press Enter. The GRUB menu appears similar to the following:

```
GNU GRUB version 0.95 (631K lower / 2095488K upper memory)
+-------------------------------------------------------------------------+
| Solaris 10 /sol_10_x86                                                  |
| Solaris failsafe                                                        |
|                                                                         |
+-------------------------------------------------------------------------+
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.
```

The nodes reboot into cluster mode using the new, upgraded BE. The cluster upgrade is completed.

23 To perform additional upgrade tasks, boot into noncluster mode.

Ensure that all nodes in the cluster are shut down before you boot nodes 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 0.95 (631K lower / 2095488K upper memory)
+-------------------------------------------------------------------------+
| Solaris 10 /sol_10_x86                                                  |
| Solaris failsafe                                                        |
|                                                                         |
+-------------------------------------------------------------------------+
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 0.95 (615K lower / 2095552K upper memory)
+-------------------------------------------------------------------------+
| root (hd0,0,a)                                                          |
| kernel /platform/i86pc/multiboot                                        |
| module /platform/i86pc/boot_archive                                     |
+-------------------------------------------------------------------------+
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/multiboot -x
```

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

     The screen displays the edited command.

```
GNU GRUB version 0.95 (615K lower / 2095552K upper memory)
+-------------------------------------------------------------------------+
| root (hd0,0,a)                                                          |
| kernel /platform/i86pc/multiboot -x                                     |
| module /platform/i86pc/boot_archive                                     |
+-------------------------------------------------------------------------+
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.-
```

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.

If the instruction says to run the `init S` command, shut down the system then change the GRUB kernel boot command to `/platform/i86pc/multiboot -sx` instead.

The upgraded BE now runs in noncluster mode.

24 If you upgraded from Sun Cluster 3.1 8/05 software, reconfigure the private-network address range.

Perform this step to increase or decrease the size of the IP address range that is used by the private interconnect. The IP address range that you configure must minimally support the number of nodes and private networks in the cluster. See "Private Network" in *Oracle Solaris Cluster Software Installation Guide* for more information.

If you also expect to configure zone clusters, you specify that number in "How to Finish Upgrade to Oracle Solaris Cluster 3.3 Software" on page 101, after all nodes are back in cluster mode.

a. From one node, start the `clsetup` utility.

When run in noncluster mode, the `clsetup` utility displays the Main Menu for noncluster-mode operations.

b. Choose the menu item, Change IP Address Range.

The `clsetup` utility displays the current private-network configuration, then asks if you would like to change this configuration.

c. To change either the private-network IP address or the IP address range, type `yes` and press the Return key.

The `clsetup` utility displays the default private-network IP address, `172.16.0.0`, and asks if it is okay to accept this default.

d. Change or accept the private-network IP address.

- **To accept the default private-network IP address and proceed to changing the IP address range, type `yes` and press the Return key.**

  The `clsetup` utility will ask if it is okay to accept the default netmask. Skip to the next step to enter your response.

- **To change the default private-network IP address, perform the following substeps.**

  i. **Type `no` in response to the `clsetup` utility question about whether it is okay to accept the default address, then press the Return key.**

  The `clsetup` utility will prompt for the new private-network IP address.

  ii. **Type the new IP address and press the Return key.**

  The `clsetup` utility displays the default netmask and then asks if it is okay to accept the default netmask.

e. Change or accept the default private-network IP address netmask and range.

The default netmask is `255.255.240.0`. This default IP address range supports up to 64 nodes, up to 10 private networks, and up to 12 zone clusters in the cluster. If you choose to change the netmask, you specify in the following substeps the number of nodes and private networks that you expect in the cluster.

If you also expect to configure zone clusters, you specify that number in "How to Finish Upgrade to Oracle Solaris Cluster 3.3 Software" on page 101, after all nodes are back in cluster mode.

- **To accept the default IP address netmask and range, type `yes` and press the Return key.**

  Then skip to the next step.

- **To change the IP address netmask and range, perform the following substeps.**

  i. **Type `no` in response to the `clsetup` utility's question about whether it is okay to accept the default address range, then press the Return key.**

  When you decline the default netmask, the `clsetup` utility prompts you for the number of nodes and private networks that you expect to configure in the cluster.

  ii. **Enter the number of nodes and private networks that you expect to configure in the cluster.**

  From these numbers, the `clsetup` utility calculates two proposed netmasks:

  - The first netmask is the minimum netmask to support the number of nodes and private networks that you specified.

- The second netmask supports twice the number of nodes and private networks that you specified, to accommodate possible future growth.

iii. Specify either of the calculated netmasks, or specify a different netmask that supports the expected number of nodes and private networks.

f. Type `yes` in response to the `clsetup` utility's question about proceeding with the update.

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

25 Upgrade any software applications that require an upgrade and for which you cannot use Solaris Live Upgrade.

Note – If an upgrade process directs you to reboot, always reboot into **noncluster** mode, as described in Step 23, until all upgrades are complete.

---

26 After all nodes are upgraded, boot the nodes into cluster mode.

a. Shut down each node.

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

b. When all nodes are shut down, boot each node into cluster mode.

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

  ```
  ok boot
  ```

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

  When the GRUB menu is displayed, select the appropriate Solaris entry and press Enter. The GRUB menu appears similar to the following:

  ```
  GNU GRUB version 0.95 (631K lower / 2095488K upper memory)
  +-------------------------------------------------------------------------+
  | Solaris 10 /sol_10_x86                                                  |
  | Solaris failsafe                                                        |
  |                                                                         |
  +-------------------------------------------------------------------------+
  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.
  ```

  The cluster upgrade is completed.

Example 4–1 Live Upgrade to Oracle Solaris Cluster 3.3 Software

This example shows a live upgrade of a cluster node. The example upgrades the SPARC based node to the Oracle Solaris 10 OS, Oracle Solaris Cluster 3.3 framework, and all Oracle Solaris

Cluster data services that support the live upgrade method. In this example, `sc31u4` is the original boot environment (BE). The new BE that is upgraded is named `sc33` and uses the mount point `/sc33`. The directory `/net/installmachine/export/solaris10/OS_image/` contains an image of the Oracle Solaris 10 OS. The `installer` state file is named `sc33state`.

The following commands typically produce copious output. This output is shown only where necessary for clarity.

```
phys-schost# lucreate sc31u4 -m /:/dev/dsk/c0t4d0s0:ufs -n sc33
...
lucreate: Creation of Boot Environment sc33 successful.

phys-schost# luupgrade -u -n sc33 -s /net/installmachine/export/solaris10/OS_image/
The Solaris upgrade of the boot environment sc33 is complete.
```

Apply Oracle Solaris patches

```
phys-schost# lumount sc33 /sc33
```

Insert the installation DVD-ROM.

```
phys-schost# cd /cdrom/cdrom0/Solaris_sparc
phys-schost# ./installer -no -saveState sc33state
phys-schost# ./installer -nodisplay -noconsole -state sc33state -altroot /sc33
phys-schost# cd /cdrom/cdrom0/Solaris_sparc/sun_cluster/Sol_9/Tools
phys-schost# ./scinstall -u update -R /sc33
phys-schost# /sc33/usr/cluster/bin/scinstall -u update -s all \
-d /cdrom/cdrom0 -R /sc33
phys-schost# cd /
phys-schost# eject cdrom

phys-schost# luumount sc33
phys-schost# luactivate sc33
Activation of boot environment sc33 successful.
```

Upgrade all other nodes

Shut down all nodes

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

When all nodes are shut down, boot each node into cluster mode

```
ok boot
```

At this point, you might upgrade data-service applications that cannot use the live upgrade method, before you reboot into cluster mode.

Troubleshooting

DID device name errors - During the creation of the inactive BE, if you receive an error that a file system that you specified with its DID device name, `/dev/dsk/did/d`*N*`s`*X*, does not exist, but the device name does exist, you must specify the device by its physical device name. Then change the `vfstab` entry on the alternate BE to use the DID device name instead. Perform the following steps:

1) For all unrecognized DID devices, specify the corresponding physical device names as arguments to the `-m` or `-M` option in the `lucreate` command. For example, if `/global/.devices/node@`*nodeid* is mounted on a DID device, use `lucreate -m /global/.devices/node@`*nodeid*`:/dev/dsk/c`*N*`t`*X*`d`*Y*`s`*Z*`:ufs [-m…] -n` *BE-name* to create the BE.

2) Mount the inactive BE by using the `lumount -n` *BE-name* `-m` *BE-mount-point* command.

3) Edit the `/`*BE-name*`/etc/vfstab` file to convert the physical device name, `/dev/dsk/c`*N*`t`*X*`d`*Y*`s`*Z*, to its DID device name, `/dev/dsk/did/d`*N*`s`*X*.

Mount point errors - During creation of the inactive boot environment, if you receive an error that the mount point that you supplied is not mounted, mount the mount point and rerun the `lucreate` command.

New BE boot errors - If you experience problems when you boot the newly upgraded environment, you can revert to your *original* BE. For specific information, see "Failure Recovery: Falling Back to the Original Boot Environment (Command-Line Interface)" in *Solaris 9 9/04 Installation Guide* or Chapter 6, "Failure Recovery: Falling Back to the Original Boot Environment (Tasks)," in *Solaris 10 10/09 Installation Guide: Solaris Live Upgrade and Upgrade Planning*.

Global-devices file-system errors - After you upgrade a cluster on which the root disk is encapsulated, you might see one of the following error messages on the cluster console during the first reboot of the upgraded BE:

```
mount: /dev/vx/dsk/bootdg/node@1 is already mounted or
/global/.devices/node@1 is busy
Trying to remount /global/.devices/node@1
mount: /dev/vx/dsk/bootdg/node@1 is already mounted or
/global/.devices/node@1 is busy

WARNING - Unable to mount one or more of the following filesystem(s):
/global/.devices/node@1
If this is not repaired, global devices will be unavailable.
Run mount manually (mount filesystem...).
After the problems are corrected, please clear the
maintenance flag on globaldevices by running the
following command:
/usr/sbin/svcadm clear svc:/system/cluster/globaldevices:default
```

```
Dec 6 12:17:23 svc.startd[8]:
svc:/system/cluster/globaldevices:default: Method
"/usr/cluster/lib/svc/method/globaldevices start" failed with exit
status 96.
[ system/cluster/globaldevices:default misconfigured (see 'svcs -x' for
details) ]
Dec 6 12:17:25 Cluster.CCR: /usr/cluster/bin/scgdevs: Filesystem
/global/.devices/node@1 is not available in /etc/mnttab.
Dec 6 12:17:25 Cluster.CCR: /usr/cluster/bin/scgdevs: Filesystem
/global/.devices/node@1 is not available in /etc/mnttab.
```

These messages indicate that the `vxio` minor number is the same on each cluster node. Reminor the root disk group on each node so that each number is unique in the cluster. See "How to Assign a New Minor Number to a Device Group" in *Oracle Solaris Cluster Software Installation Guide*.

Next Steps Go to Chapter 6, "Completing the Upgrade."

See Also You can choose to keep your original, and now inactive, boot environment for as long as you need to. When you are satisfied that your upgrade is acceptable, you can then choose to remove the old environment or to keep and maintain it.

- If you used an unmirrored volume for your inactive BE, delete the old BE files. For specific information, see the appropriate procedure for your *original* Solaris OS version.
  - "Deleting an Inactive Boot Environment" in *Solaris 9 9/04 Installation Guide*.
  - "Deleting an Inactive Boot Environment" in *Solaris 10 10/09 Installation Guide: Solaris Live Upgrade and Upgrade Planning*.
- If you detached a plex to use as the inactive BE, reattach the plex and synchronize the mirrors. For more information about working with a plex, see the appropriate procedure for your *original* Solaris OS versions.
  - "Example of Detaching and Upgrading One Side of a RAID 1 Volume (Mirror) (Command-Line Interface)" in *Solaris 9 9/04 Installation Guide*.
  - "Example of Detaching and Upgrading One Side of a RAID-1 Volume (Mirror)" in *Solaris 10 10/09 Installation Guide: Solaris Live Upgrade and Upgrade Planning*.

You can also maintain the inactive BE. For information about how to maintain the environment, see the appropriate procedure for your *original* Solaris OS versions.

- Chapter 37, "Maintaining Solaris Live Upgrade Boot Environments (Tasks)," in *Solaris 9 9/04 Installation Guide*.
- Chapter 7, "Maintaining Solaris Live Upgrade Boot Environments (Tasks)," in *Solaris 10 10/09 Installation Guide: Solaris Live Upgrade and Upgrade Planning*.

◆ ◆ ◆ CHAPTER 5

Performing a Rolling Upgrade

This chapter provides procedures to perform a rolling upgrade of an Oracle Solaris Cluster 3.3 release to an Oracle Solaris Cluster 3.3 update release, or to upgrade Oracle Solaris 10 to an Oracle Solaris 10 update release. 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 10 software from a previous marketing release such as Solaris 9 software, or to Oracle Solaris Cluster software from a previous marketing release such as Sun Cluster 3.2 software. You can only perform and Oracle Solaris Cluster rolling upgrade of Oracle Solaris Cluster or Solaris software to an update of the same marketing release.

To upgrade an Oracle Solaris Cluster configuration from an earlier marketing release of Solaris software, use another upgrade method. See "Choosing an Oracle Solaris Cluster Upgrade Method" on page 13 to determine the best upgrade method for your configuration.

This chapter provides the following information to upgrade an Oracle Solaris Cluster 3.3 configuration to an update release of the Oracle Solaris Cluster 3.3 software or to an update release of the Solaris OS by using the rolling upgrade method:

- "Performing a Rolling Upgrade of a Cluster" on page 86

Performing a Rolling Upgrade of a Cluster

TABLE 5–1 Task Map: Performing a Rolling Upgrade to Oracle Solaris Cluster 3.3 Software

Task	Instructions
1. Read the upgrade requirements and restrictions.	"Upgrade Requirements and Software Support Guidelines" on page 11
2. If a quorum server is used, upgrade the Quorum Server software.	"How to Upgrade Quorum Server Software" on page 86
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 Oracle Solaris Cluster Geographic Edition software is installed, uninstall it. 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 88
4. Upgrade the Solaris OS on the cluster node, if necessary, to a supported Solaris update release.	"How to Perform a Rolling Upgrade of a Solaris Maintenance Update" on page 92
5. Upgrade the cluster node to Oracle Solaris Cluster 3.3 framework software. Optionally, upgrade data-service software. If necessary, upgrade applications.	"How to Perform a Rolling Upgrade of Oracle Solaris Cluster 3.3 Software" on page 94
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 to Oracle Solaris Cluster 3.3 Software" on page 99
8. Verify successful completion of upgrade to Oracle Solaris Cluster 3.3 software.	"How to Verify Upgrade of Oracle Solaris Cluster 3.3 Software" on page 100
9. Enable resources and bring resource groups online. Migrate existing resources to new resource types. Upgrade to the Oracle Solaris Cluster Geographic Edition 3.3 software, if used.	"How to Finish Upgrade to Oracle Solaris Cluster 3.3 Software" on page 101
10. (*Optional*) SPARC: Upgrade the Oracle Solaris Cluster module to Sun Management Center.	"SPARC: How to Upgrade Oracle Solaris Cluster Module Software for Sun Management Center" on page 121

▼ 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 3.3 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.**

a. **Navigate to the directory where the uninstaller is located.**

```
quorumserver# cd /var/sadm/prod/SUNWentsysver
```

ver The version that is installed on your system.

b. **Start the uninstallation wizard.**

```
quorumserver# ./uninstall
```

c. Follow instructions on the screen to uninstall the Quorum Server software from the quorum-server host computer.

After removal is finished, you can view any available log. See Chapter 8, "Uninstalling," in *Sun Java Enterprise System 5 Update 1 Installation Guide for UNIX* for additional information about using the `uninstall` program.

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

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

6 Install the Oracle Solaris Cluster 3.3 Quorum Server software, reconfigure the quorum server, and start the quorum server daemon.

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

7 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*.

8 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 11.
- Have available the installation media, documentation, and patches for all the software products that you are upgrading, including the following software:
  - Solaris OS
  - Oracle Solaris Cluster 3.3 framework
  - Oracle Solaris Cluster 3.3 required patches
  - Oracle Solaris Cluster 3.3 data services (agents)
  - Applications that are managed by Oracle Solaris Cluster 3.3 data service agents

  See "Patches and Required Firmware Levels" in the Oracle Solaris Cluster 3.3 Release Notes for the location of patches and installation instructions.

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. **If you are upgrading Oracle Solaris Cluster 3.3 software and Oracle Solaris Cluster Geographic Edition software is installed, uninstall it.**

   For uninstallation procedures, see the documentation for your version of Oracle Solaris Cluster Geographic Edition software.

4. **Become superuser on a node of the cluster.**

5. **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.

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 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 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 to Oracle Solaris Cluster 3.3 Software" on page 99.

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 Solaris entry and type e to edit its commands.

The GRUB menu appears similar to the following:

```
GNU GRUB version 0.95 (631K lower / 2095488K upper memory)
+-------------------------------------------------------------------------+
| Solaris 10 /sol_10_x86                                                  |
| Solaris failsafe                                                        |
|                                                                         |
+-------------------------------------------------------------------------+
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 0.95 (615K lower / 2095552K upper memory)
+-------------------------------------------------------------------------+
| root (hd0,0,a)                                                          |
| kernel /platform/i86pc/multiboot                                        |
| module /platform/i86pc/boot_archive                                     |
+-------------------------------------------------------------------------+
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/multiboot -x
```

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

The screen displays the edited command.

```
GNU GRUB version 0.95 (615K lower / 2095552K upper memory)
+-------------------------------------------------------------------------+
| root (hd0,0,a)                                                          |
| kernel /platform/i86pc/multiboot -x                                     |
| module /platform/i86pc/boot_archive                                     |
+-------------------------------------------------------------------------+
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.-
```

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 Solaris software to a Maintenance Update release, go to "How to Perform a Rolling Upgrade of a Solaris Maintenance Update" on page 92.

Note – The cluster must already run on, or be upgraded to, at least the minimum required level of the Solaris OS to support Oracle Solaris Cluster 3.3 software. See the Oracle Solaris Cluster 3.3 Release Notes for information about supported releases of the Solaris OS.

Otherwise, if you do not intend to upgrade the Solaris OS, go to "How to Perform a Rolling Upgrade of Oracle Solaris Cluster 3.3 Software" on page 94.

▼ How to Perform a Rolling Upgrade of a Solaris Maintenance Update

Perform this procedure to upgrade the Solaris OS to a supported Maintenance Update release.

Note – You cannot perform a rolling upgrade to upgrade a cluster from Solaris 9 to Oracle Solaris 10 software. Go to "Choosing an Oracle Solaris Cluster Upgrade Method" on page 13 to identify the appropriate upgrade method to use.

Before You Begin Ensure that all steps in "How to Prepare a Cluster Node for a Rolling Upgrade" on page 88 are completed.

1. **Temporarily comment out all entries for globally mounted file systems in the node's `/etc/vfstab` file.**

   Perform this step to prevent the Solaris upgrade from attempting to mount the global devices.

2. **Follow the instructions in the Solaris maintenance update installation guide to install the Maintenance Update release.**

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

   ---

3. **Uncomment all entries in the `/a/etc/vfstab` file for globally mounted file systems that you commented out in Step 1.**

4. **Install any required Solaris software patches and hardware-related patches, and download any needed firmware that is contained in the hardware patches.**

Note – Do not reboot the node until Step 5.

5 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 Solaris entry and type e to edit its commands.**

  The GRUB menu appears similar to the following:

```
GNU GRUB version 0.95 (631K lower / 2095488K upper memory)
+-------------------------------------------------------------------------+
| Solaris 10 /sol_10_x86                                                  |
| Solaris failsafe                                                        |
|                                                                         |
+-------------------------------------------------------------------------+
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 0.95 (615K lower / 2095552K upper memory)
+-------------------------------------------------------------------------+
| root (hd0,0,a)                                                          |
| kernel /platform/i86pc/multiboot                                        |
| module /platform/i86pc/boot_archive                                     |
+-------------------------------------------------------------------------+
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/multiboot -x
```

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

The screen displays the edited command.

```
GNU GRUB version 0.95 (615K lower / 2095552K upper memory)
+-------------------------------------------------------------------------+
| root (hd0,0,a)                                                          |
| kernel /platform/i86pc/multiboot -x                                     |
| module /platform/i86pc/boot_archive                                     |
+-------------------------------------------------------------------------+
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.-
```

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 of Oracle Solaris Cluster 3.3 Software" on page 94.

▼ How to Perform a Rolling Upgrade of Oracle Solaris Cluster 3.3 Software

Perform this procedure to upgrade a node that runs Oracle Solaris Cluster 3.3 software while the remaining cluster nodes are in cluster mode.

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 on the node of the cluster.

2 If you upgraded the Solaris OS but do not need to upgrade to an Oracle Solaris Cluster update release, skip to Step 13.

3 Load the installation DVD-ROM into the DVD-ROM drive.

If the volume management daemon `vold(1M)` is running and is configured to manage CD-ROM or DVD devices, the daemon automatically mounts the media on the `/cdrom/cdrom0` directory.

4. **Change to the `/cdrom/cdrom0/Solaris_`*arch*`/Product/sun_cluster/Solaris_`*ver*`/Tools/` directory, where *arch* is `sparc` or `x86` and where *ver* is `10` for Oracle Solaris 10 .**

   ```
   phys-schost# cd /cdrom/cdrom0/Solaris_arch/Product/sun_cluster/Solaris_ver/Tools
   ```

5. **Start the `scinstall` utility.**

   ```
   phys-schost# ./scinstall
   ```

   Note – Do not use the `/usr/cluster/bin/scinstall` command that is already installed on the node. You must use the `scinstall` command that is located on the installation DVD-ROM.

   The `scinstall` Main Menu is displayed.

6. **Choose the menu item, Upgrade This Cluster Node.**

   ```
   *** Main Menu ***

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

         1) Create a new cluster or add a cluster node
         2) Configure a cluster to be JumpStarted from this install server
       * 3) Manage a dual-partition upgrade
       * 4) Upgrade this cluster node
       * 5) Print release information for this cluster node

       * ?) Help with menu options
       * q) Quit

     Option:  4
   ```

   The Upgrade Menu is displayed.

7. **Choose the menu item, Upgrade Oracle Solaris Cluster Framework on This Node.**

8. **Follow the menu prompts to upgrade the cluster framework.**

   Upgrade processing is finished when the system displays the message `Completed Oracle Solaris Cluster framework upgrade` and prompts you to press Enter to continue.

9. **Quit the `scinstall` utility.**

10. **(Optional) Upgrade data service packages.**

    Note – For HA for SAP Web Application Server, if you are using a J2EE engine resource or a web application server component resource or both, you must delete the resource and recreate it with the new web application server component resource. Changes in the new web application server component resource includes integration of the J2EE functionality. For more information, see *Oracle Solaris Cluster Data Service for SAP Web Application Server Guide*.

a. Start the upgraded interactive `scinstall` utility.

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

Note – Do not use the `scinstall` utility that is on the installation media to upgrade data service packages.

The `scinstall` Main Menu is displayed.

b. Choose the menu item, Upgrade This Cluster Node.

The Upgrade Menu is displayed.

c. Choose the menu item, Upgrade Oracle Solaris Cluster Data Service Agents on This Node.

d. Follow the menu prompts to upgrade Oracle Solaris Cluster data service agents that are installed on the node.

You can choose from the list of data services that are available to upgrade or choose to upgrade all installed data services.

e. When the system displays the message `Completed upgrade of Oracle Solaris Cluster data services agents`, press Return.

The Upgrade Menu is displayed.

11 Quit the `scinstall` utility.

12 Unload the installation DVD-ROM from the DVD-ROM drive.

a. To ensure that the DVD-ROM is not being used, change to a directory that does *not* reside on the DVD-ROM.

b. Eject the DVD-ROM.

```
phys-schost# eject cdrom
```

13 If you have HA for NFS configured on a highly available local file system, ensure that the loopback file system (LOFS) is disabled.

Note – If you have non-global zones configured, LOFS must remain enabled. For guidelines about using LOFS and alternatives to disabling it, see "Cluster File Systems" in *Oracle Solaris Cluster Software Installation Guide*.

As of the Sun Cluster 3.2 release, LOFS is no longer disabled by default during Oracle Solaris Cluster software installation or upgrade. To disable LOFS, ensure that the `/etc/system` file contains the following entry:

```
exclude:lofs
```

This change becomes effective at the next system reboot.

14 As needed, manually upgrade any custom data services that are not supplied on the product media.

15 Verify that each data-service update is installed successfully.

View the upgrade log file that is referenced at the end of the upgrade output messages.

16 Upgrade software applications that are installed on the cluster.

If you want to upgrade VxVM and did not upgrade the Solaris OS, follow procedures in Veritas Storage Foundation installation documentation to upgrade VxVM without upgrading the operating system.

Note – If any upgrade procedure instruct you to perform a reboot, you must add the `-x` option to the boot command. This option boots the cluster into noncluster mode.

Ensure that application levels are compatible with the current versions of Oracle Solaris Cluster and Solaris software. See your application documentation for installation instructions.

17 Shut down the node.

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

18 Reboot the node into the cluster.

- On SPARC based systems, perform the following command:

  ```
  ok boot
  ```

- On x86 based systems, perform the following commands:

  When the GRUB menu is displayed, select the appropriate Solaris entry and press Enter. The GRUB menu appears similar to the following:

  ```
  GNU GRUB version 0.95 (631K lower / 2095488K upper memory)
  +-------------------------------------------------------------------------+
  | Solaris 10 /sol_10_x86                                                  |
  | Solaris failsafe                                                        |
  |                                                                         |
  +-------------------------------------------------------------------------+
  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.
  ```

19 Return to "How to Prepare a Cluster Node for a Rolling Upgrade" on page 88 and repeat all upgrade procedures on the next node to upgrade.

Repeat this process until all nodes in the cluster are upgraded.

Next Steps When all nodes in the cluster are upgraded, go to Chapter 6, "Completing the Upgrade."

◆ ◆ ◆ CHAPTER 6

Completing the Upgrade

This chapter provides the following information to complete all Oracle Solaris Cluster 3.3 software upgrade methods:

Completing a Cluster Upgrade

▼ How to Commit the Upgraded Cluster to Oracle Solaris Cluster 3.3 Software

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 Upgrade of Oracle Solaris Cluster 3.3 Software" on page 100.
`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 Upgrade of Oracle Solaris Cluster 3.3 Software" on page 100.

▼ How to Verify Upgrade of Oracle Solaris Cluster 3.3 Software

Perform this procedure to verify that the cluster is successfully upgraded to Oracle Solaris Cluster 3.3 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 to Oracle Solaris Cluster 3.3 Software" on page 99 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 6–1 Verifying Upgrade to Oracle Solaris Cluster 3.3 Software

The following example shows the commands used to verify upgrade of a two-node cluster to Oracle Solaris Cluster 3.3 software. The cluster node names are `phys-schost-1` and `phys-schost-2`.

```
phys-schost# clnode show-rev -v
3.3
...
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 Upgrade to Oracle Solaris Cluster 3.3 Software" on page 101.

▼ How to Finish Upgrade to Oracle Solaris Cluster 3.3 Software

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 Upgrade of Oracle Solaris Cluster 3.3 Software" on page 100 are completed.

1 Install any Oracle Solaris Cluster 3.3 framework and data-service software patches.

See "Patches and Required Firmware Levels" in the Oracle Solaris Cluster 3.3 Release Notes for the location of patches and installation instructions.

2 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 Sun Java Web Console agent.

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

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

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

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

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

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

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

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

f. 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
```

g. 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
```

3 If you upgraded Solaris software, ensure that external access to RPC communication and Sun Java Web Console is enabled.

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

During 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, and the Sun Java Web Console service, which is used by the Oracle Solaris Cluster Manager GUI.

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

a. Display the status of external access to RPC communication and Oracle Java Web Console.

- For 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
  ```

- For Oracle Java Web Console, the external access is available if the output of the following command returns an entry for 6789, which is the port number that is used to connect to Oracle Java Web Console.

  ```
  phys-schost# netstat -a | grep 6789
  ```

If external access to both services is enabled, skip to Step 4. 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. If external access to Oracle Java Web Console is restricted, perform the following commands.

```
phys-schost# svccfg
svc:> select system/webconsole
svc:/system/webconsole> setprop options/tcp_listen=true
svc:/system/webconsole> quit
phys-schost# /usr/sbin/smcwebserver restart
```

For more information about what services the restricted network profile restricts to local connections, see "Planning Network Security" in *Solaris 10 10/09 Installation Guide: Planning for Installation and Upgrade*.

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

4 On each node, start the security file agent and then start the Sun Java Web Console agent.

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

5 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.

6 If you upgraded HA for SAP liveCache from the Sun Cluster 3.1 8/05 version to the Oracle Solaris Cluster 3.3 version, modify the `/opt/SUNWsclc/livecache/bin/lccluster` configuration file.

a. Become superuser on a node that will host the liveCache resource.

b. Copy the new `/opt/SUNWsclc/livecache/bin/lccluster` file to the `/sapdb/`*LC_NAME*`/db/sap/` directory.

Overwrite the `lccluster` file that already exists from the previous configuration of the data service.

c. Configure this `/sapdb/`*LC_NAME*`/db/sap/lccluster` file as documented in "How to Register and Configure Solaris Cluster HA for SAP liveCache" in *Oracle Solaris Cluster Data Service for SAP liveCache Guide*.

7 If you upgraded the 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.

8 If you upgraded VxVM, upgrade all disk groups.

a. Bring online and take ownership of a disk group to upgrade.

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

b. Synchronize the disk group.

This step resolves any changes made to VxVM minor numbers during VxVM upgrade.

```
phys-schost# cldevicegroup sync devicegroup
```

c. Run the following command to upgrade a disk group to the highest version supported by the VxVM release you installed.

```
phys-schost# vxdg upgrade devicegroup
```

See your VxVM administration documentation for more information about upgrading disk groups.

d. On each node that is directly connected to the disk group, bring online and take ownership of the upgraded disk group.

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

This step is necessary is to update the VxVM device files major number with the latest `vxio` number that might have been assigned during the upgrade.

e. Repeat for each remaining VxVM disk group in the cluster.

9 Migrate resources to new resource type versions.

You must migrate all resources to the Oracle Solaris Cluster 3.3 resource-type version to use the new features and bug fixes that are provided in this release.

Note – For HA for SAP Web Application Server, if you are using a J2EE engine resource or a web application server component resource or both, you must delete the resource and recreate it with the new web application server component resource. Changes in the new web application server component resource includes integration of the J2EE functionality. For more information, see *Oracle Solaris Cluster Data Service for SAP Web Application Server Guide*.

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 3.3 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.

---

10 If your cluster runs the HA for Sun Java System Application Server EE (HADB) data service and you shut down the HADB database before you began a dual-partition upgrade, re-enable the resource and start the database.

```
phys-schost# clresource enable hadb-resource
phys-schost# hadbm start database-name
```

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

11 If you upgraded to the Oracle Solaris 10 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"
```

12 If you upgraded to the Oracle Solaris 10 OS and you intend to configure zone clusters, set the number of expected zone clusters.

Specify the number of zone clusters that you expect to configure in the cluster.

```
phys-schost# cluster set net-props -p num_zoneclusters=N
```

The command calculates the number of additional private-network IP addresses that are needed and automatically modifies the IP address range.

13 From any node, start the `clsetup` utility.

```
phys-schost# clsetup
```

The `clsetup` Main Menu is displayed.

14 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.

15 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.

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

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

17 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
...
```

18 (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.

19 (Optional) Install or complete upgrade of Oracle Solaris Cluster Geographic Edition 3.3 software.

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

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

An archived backup of your cluster configuration facilitates easier recovery of the your cluster configuration,

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

Troubleshooting

Setting the number of zone clusters fails – After a cluster is upgraded from the Sun Cluster 3.1 8/05 release to the Oracle Solaris Cluster 3.3 release, setting the expected number of zone clusters might fail if the expected number of nodes and private networks has not yet been set. These settings were introduced in the Sun Cluster 3.2 release. To correct this problem, do the following:

- 1) Set the number of nodes and private networks in the cluster. Follow instructions in "How to Change the Private Network Address or Address Range of an Existing Cluster" in *Oracle Solaris Cluster System Administration Guide*. This task requires putting all cluster nodes into noncluster mode.
- 2) After you set the expected number of nodes and reboot all nodes into cluster mode, rerun Step 12 to set the expected number of zone clusters.

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` or `smcwebserver start` commands:

```
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

phys-schost# smcwebserver start
/usr/sbin/smcwebserver: /usr/jdk/jdk1.5.0_06/bin/java: not found
```

These errors are generated because the start commands 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 files to use the current Java directory:

```
/etc/webconsole/console/config.properties/etc/opt/SUNWcacao/cacao.properties
```

Next Steps If you have a SPARC based system and use Sun Management Center to monitor the cluster, go to "SPARC: How to Upgrade Oracle Solaris Cluster Module Software for Sun Management Center" on page 121.

Otherwise, the cluster upgrade is complete.

◆ ◆ ◆ CHAPTER 7

Recovering From an Incomplete Upgrade

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

Cluster Recovery After an Incomplete Upgrade

This section provides information to recover from incomplete upgrades of an Oracle Solaris Cluster configuration.

▼ How to Recover from a Failed Dual-Partition Upgrade

If you experience an unrecoverable error during upgrade, perform this procedure to back out of the upgrade.

Note – You cannot restart a dual-partition upgrade after the upgrade has experienced an unrecoverable error.

---

1 Become superuser on each node of the cluster.

2 Boot each 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)
+-------------------------------------------------------------------------+
| Solaris 10 /sol_10_x86                                                  |
| Solaris failsafe                                                        |
|                                                                         |
+-------------------------------------------------------------------------+
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)
+-------------------------------------------------------------------------+
| root (hd0,0,a)                                                          |
| kernel /platform/i86pc/multiboot                                        |
| module /platform/i86pc/boot_archive                                     |
|+------------------------------------------------------------------------+
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/multiboot -x
```

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

     The screen displays the edited command.

```
GNU GRUB version 0.97 (639K lower / 1047488K upper memory)
+-------------------------------------------------------------------------+
| root (hd0,0,a)                                                          |
| kernel /platform/i86pc/multiboot -x                                     |
| module /platform/i86pc/boot_archive                                     |
+-------------------------------------------------------------------------+
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.-
```

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.

---

3 On each node, run the upgrade recovery script from the installation media.

If the node successfully upgraded to Oracle Solaris Cluster 3.3 software from a 3.2 release, you can alternatively run the `scinstall` command from the `/usr/cluster/bin` directory.

Note – If you upgraded from the Sun Cluster 3.1 8/05 release, run the `scinstall` command **only** from the installation media. Recovery capability for dual-partition upgrade is not available from the 3.1 version of the `scinstall` command.

```
phys-schost# cd /cdrom/cdrom0/Solaris_arch/Product/sun_cluster/Solaris_ver/Tools
phys-schost# ./scinstall -u recover
```

`-u`	Specifies upgrade.
`recover`	Restores the `/etc/vfstab` file and the Cluster Configuration Repository (CCR) database to their original state before the start of the dual-partition upgrade.

The recovery process leaves the cluster nodes in noncluster mode. Do **not** attempt to reboot the nodes into cluster mode.

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

4 Perform either of the following tasks.

- **Restore the old software from backup to return the cluster to its original state.**
- **Continue to upgrade software on the cluster by using the standard upgrade method.**

  This method requires that all cluster nodes remain in noncluster mode during the upgrade. See the task map for standard upgrade, Table 2–1. You can resume the upgrade at the last task or step in the standard upgrade procedures that you successfully completed before the dual-partition upgrade failed.

▼ SPARC: How to Recover From a Partially Completed Dual-Partition Upgrade

Perform this procedure if a dual-partition upgrade fails and the state of the cluster meets *all* of the following criteria:

- The nodes of the first partition are upgraded.
- None of the nodes of the second partition are yet upgraded.
- None of the nodes of the second partition are in cluster mode.

You can also perform this procedure if the upgrade has succeeded on the first partition but you want to back out of the upgrade.

Note – Do not perform this procedure after dual-partition upgrade processes have begun on the second partition. Instead, perform "How to Recover from a Failed Dual-Partition Upgrade" on page 111.

Before You Begin Before you begin, ensure that all second-partition nodes are halted. First-partition nodes can be either halted or running in noncluster mode.

Perform all steps as superuser.

1 **Boot each node in the second partition into noncluster mode.**

```
phys-schost# ok boot -x
```

2 **From the DVD image, run the `scinstall -u recover` command on each node in the second partition.**

Change to the `/Solaris_`*arch*`/Product/sun_cluster/Solaris_`*ver*`/Tools/` directory, where *arch* is `sparc` or `x86` and where `ver` is `10` for Oracle Solaris 10.

```
phys-schost# cd /cdrom/cdrom0/Solaris_arch/Product/sun_cluster/Solaris_ver/Tools/
phys-schost# ./scinstall -u recover
```

The command restores the original CCR information, restores the original `/etc/vfstab` file, and eliminates modifications for startup.

3 **Boot each node of the second partition into cluster mode.**

```
phys-schost# shutdown -g0 -y -i6
```

When the nodes of the second partition come up, the second partition resumes supporting cluster data services while running the old software with the original configuration.

4 **Restore the original software and configuration data from backup media to the nodes in the first partition.**

5 **Boot each node in the first partition into cluster mode.**

```
phys-schost# shutdown -g0 -y -i6
```

The nodes rejoin the cluster.

x86: How to Recover From a Partially Completed Dual-Partition Upgrade

Perform this procedure if a dual-partition upgrade fails and the state of the cluster meets *all* of the following criteria:

- The nodes of the first partition are upgraded.
- None of the nodes of the second partition are yet upgraded.
- None of the nodes of the second partition are in cluster mode.

You can also perform this procedures if the upgrade has succeeded on the first partition but you want to back out of the upgrade.

Note – Do not perform this procedure after dual-partition upgrade processes have begun on the second partition. Instead, perform "How to Recover from a Failed Dual-Partition Upgrade" on page 111.

Before You Begin Before you begin, ensure that all second-partition nodes are halted. First-partition nodes can be either halted or running in noncluster mode.

Perform all steps as superuser.

1 **Boot each node in the second partition into noncluster mode by completing the following steps.**

2 **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 0.95 (631K lower / 2095488K upper memory)
+-------------------------------------------------------------------------+
| Solaris 10 /sol_10_x86                                                  |
| Solaris failsafe                                                        |
|                                                                         |
+-------------------------------------------------------------------------+
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*.

3 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 0.95 (615K lower / 2095552K upper memory)
+-------------------------------------------------------------------------+
| root (hd0,0,a)                                                          |
| kernel /platform/i86pc/multiboot                                        |
| module /platform/i86pc/boot_archive                                     |
+-------------------------------------------------------------------------+
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.
```

4 Add the -x option 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.


phys-schost# grub edit> kernel /platform/i86pc/multiboot -x
```

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

The screen displays the edited command.

```
GNU GRUB version 0.95 (615K lower / 2095552K upper memory)
+-------------------------------------------------------------------------+
| root (hd0,0,a)                                                          |
| kernel /platform/i86pc/multiboot -x                                     |
| module /platform/i86pc/boot_archive                                     |
+-------------------------------------------------------------------------+
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.-
```

6 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.

---

7 On each node in the second partition, run the `scinstall -u recover` command.

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

The command restores the original CCR information, restores the original `/etc/vfstab` file, and eliminates modifications for startup.

8 **Boot each node of the second partition into cluster mode.**

```
phys-schost# shutdown -g0 -y -i6
```

When the nodes of the second partition come up, the second partition resumes supporting cluster data services while running the old software with the original configuration.

9 **Restore the original software and configuration data from backup media to the nodes in the first partition.**

10 **Boot each node in the first partition into cluster mode.**

```
phys-schost# shutdown -g0 -y -i6
```

The nodes rejoin the cluster.

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 117
- "How to Resolve Mistaken Storage Changes During an Upgrade" on page 118

▼ 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 Solaris 9 or Oracle Solaris 10 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.**

```
phys-schost# cldevice repair device
```

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

3 **Update the DID driver.**

```
phys-schost# scdidadm -ui
phys-schost# scdidadm -r
```

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

`-i` Initializes the DID driver.

`-r` Reconfigures the database.

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 Step 4 in "How to Upgrade Oracle Solaris Cluster 3.3 Software (Standard)" on page 32.

▼ 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 117.

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# scdidadm -r
```

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

`-i` Initializes the DID driver.

`-r` Reconfigures the database.

See the `scdidadm(1M)` man page 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 Step 4 in "How to Upgrade Oracle Solaris Cluster 3.3 Software (Standard)" on page 32.

◆ ◆ ◆ CHAPTER 8

SPARC: Upgrading Sun Management Center Software

This chapter provides the following procedures to upgrade the Oracle Solaris Cluster module for Sun Management Center:

Upgrading the Oracle Solaris Cluster Module for Sun Management Center

This section provides procedures to upgrade the Oracle Solaris Cluster module for Sun Management Center software.

▼ SPARC: How to Upgrade Oracle Solaris Cluster Module Software for Sun Management Center

Perform the following steps to upgrade Oracle Solaris Cluster module software on the Sun Management Center server machine and console machine.

Note – To upgrade the Sun Management Center software itself, do not perform this procedure. Instead, go to "SPARC: How to Upgrade Sun Management Center Software" on page 123 to upgrade the Sun Management Center software and the Oracle Solaris Cluster module.

Before You Begin Have available the installation DVD-ROM for the SPARC platform or the path to the DVD-ROM image.

1 **As superuser, remove the existing Oracle Solaris Cluster module packages from each machine.**

Use the `pkgrm(1M)` command to remove all Oracle Solaris Cluster module packages from all locations that are listed in the following table.

Location	Module Package to Remove
Sun Management Center console machine	`SUNWscscn`
Sun Management Center server machine	`SUNWscssv`, `SUNWscshl`

```
machine# pkgrm module-package
```

Note – Oracle Solaris Cluster module software on the cluster nodes was already upgraded during the cluster-framework upgrade.

2 **As superuser, reinstall Oracle Solaris Cluster module packages on each machine.**

a. **Insert the installation DVD-ROM for the SPARC platform into the DVD-ROM drive of the machine.**

b. **Change to the `Solaris_sparc/Product/sun_cluster/Solaris_`*ver*`/Packages/` directory, where *ver* is `10` for Oracle Solaris 10.**

```
machine# cd Solaris_sparc/Product/sun_cluster/Solaris_ver/Packages/
```

c. **Install the `SUNWscssv` package on the server machine.**

Note that you do not upgrade to a new `SUNWscscn` package on the console machine or a new `SUNWscshl` package on the server machine.

```
machine# pkgadd -d . SUNWscssv
```

d. **Unload the installation DVD-ROM from the DVD-ROM drive.**

i. **To ensure that the DVD-ROM is not being used, change to a directory that does *not* reside on the DVD-ROM.**

ii. **Eject the DVD-ROM.**

```
machine# eject cdrom
```

Upgrading Sun Management Center Software

This section provides procedures to upgrade Sun Management Center software.

▼ SPARC: How to Upgrade Sun Management Center Software

Perform the following steps to upgrade from Sun Management Center 3.6 or 3.6.1 software to Sun Management Center 3.6.1 or 4.0 software.

Before You Begin Have available the following items:

- Oracle Solaris Cluster installation DVD-ROM for the SPARC platform and, if applicable, for the x86 platform, or the paths to the DVD-ROM images. You use the DVD-ROM to reinstall the Oracle Solaris Cluster 3.3 version of the Oracle Solaris Cluster module packages after you upgrade Sun Management Center software.

  Note – The agent packages to install on the cluster nodes are available for both SPARC based systems and x86 based systems. The package for the server machine is available for SPARC based systems only.

- Sun Management Center documentation.
- Sun Management Center patches and Oracle Solaris Cluster module patches, if any.

  See "Patches and Required Firmware Levels" in the Oracle Solaris Cluster 3.3 Release Notes for the location of patches and installation instructions.

1 Stop any Sun Management Center processes.

a. If the Sun Management Center console is running, exit the console.

In the console window, choose File⇒Exit.

b. On each Sun Management Center agent machine (cluster node), stop the Sun Management Center agent process.

```
phys-schost# /opt/SUNWsymon/sbin/es-stop -a
```

c. On the Sun Management Center server machine, stop the Sun Management Center server process.

```
server# /opt/SUNWsymon/sbin/es-stop -S
```

2 As superuser, remove Oracle Solaris Cluster–module packages.

Use the `pkgrm(1M)` command to remove all Oracle Solaris Cluster module packages from all locations that are listed in the following table.

Location	Module Package to Remove
Each cluster node	`SUNWscsam`, `SUNWscsal`
Sun Management Center console machine	`SUNWscscn`
Sun Management Center server machine	`SUNWscssv`, `SUNWscshl`

```
machine# pkgrm module-package
```

If you do not remove the listed packages, the Sun Management Center software upgrade might fail because of package dependency problems. You reinstall these packages in Step 4, after you upgrade Sun Management Center software.

3 Upgrade the Sun Management Center software.

Follow the upgrade procedures in your Sun Management Center documentation.

4 As superuser, reinstall Oracle Solaris Cluster module packages from the installation DVD-ROM to the locations that are listed in the following table.

Location	Module Package to Install
Each cluster node	`SUNWscsam`, `SUNWscsal`
Sun Management Center server machine	`SUNWscssv`

a. Insert the installation DVD-ROM for the appropriate platform in the DVD-ROM drive of the machine.

b. Change to the `/Solaris_`*arch*`/Product/sun_cluster/Solaris_`*ver*`/Packages/` directory, where *arch* is `sparc` or `x86`, and *ver* is `10` for Oracle Solaris 10.

```
machine# cd /cdrom/cdrom0/Solaris_arch/Product/sun_cluster/Solaris_ver/Packages/
```

Note – The agent packages to install on the cluster nodes are available for both SPARC based systems and x86 based systems. The package for the server machine is available for SPARC based systems only.

c. Install the appropriate module package on the machine.

- **For cluster nodes, use the following command:**

  ```
  phys-schost# pkgadd -G -d .
  module-package
  ```

  The `-G` option adds packages to the current zone only. You must add these packages only to the global zone. Therefore, this option also specifies that the packages are *not* propagated to any existing non-global zone or to any non-global zone that is created later.

- **SPARC: For the server machine, use the following command:**

  ```
  phys-schost# pkgadd -d . module-package
  ```

5 Apply any Sun Management Center patches and any Oracle Solaris Cluster module patches to each node of the cluster.

6 Restart Sun Management Center agent, server, and console processes.

Follow procedures in "How to Start Sun Management Center" in *Oracle Solaris Cluster Software Installation Guide*.

7 Load the Oracle Solaris Cluster module.

Follow procedures in "How to Load the Oracle Solaris Cluster Module" in *Oracle Solaris Cluster Software Installation Guide*.

If the Oracle Solaris Cluster module was previously loaded, unload the module and then reload it to clear all cached alarm definitions on the server. To unload the module, choose Unload Module from the Module menu on the console's Details window.

Index

T

U

V

Z