

Oracle® Solaris Cluster Upgrade Guide

ORACLE®

Part No: E39644
July 2014, E39644-01

Copyright © 2000, 2014, Oracle and/or its affiliates. All rights reserved.

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

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

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

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

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

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

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

This software or hardware and documentation may provide access to or information on content, products, and services from third parties. Oracle Corporation and its affiliates are not responsible for and expressly disclaim all warranties of any kind with respect to third-party content, products, and services. Oracle Corporation and its affiliates will not be responsible for any loss, costs, or damages incurred due to your access to or use of third-party content, products, or services.

Copyright © 2000, 2014, Oracle et/ou ses affiliés. Tous droits réservés.

Ce logiciel et la documentation qui l'accompagne sont protégés par les lois sur la propriété intellectuelle. Ils sont concédés sous licence et soumis à des restrictions d'utilisation et de divulgation. Sauf disposition de votre contrat de licence ou de la loi, vous ne pouvez pas copier, reproduire, traduire, diffuser, modifier, breveter, transmettre, distribuer, exposer, exécuter, publier ou afficher le logiciel, même partiellement, sous quelque forme et par quelque procédé que ce soit. Par ailleurs, il est interdit de procéder à toute ingénierie inverse du logiciel, de le désassembler ou de le décompiler, excepté à des fins d'interopérabilité avec des logiciels tiers ou tel que prescrit par la loi.

Les informations fournies dans ce document sont susceptibles de modification sans préavis. Par ailleurs, Oracle Corporation ne garantit pas qu'elles soient exemptes d'erreurs et vous invite, le cas échéant, à lui en faire part par écrit.

Si ce logiciel, ou la documentation qui l'accompagne, est concédé sous licence au Gouvernement des Etats-Unis, ou à toute entité qui délivre la licence de ce logiciel ou l'utilise pour le compte du Gouvernement des Etats-Unis, la notice suivante s'applique:

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

Ce logiciel ou matériel a été développé pour un usage général dans le cadre d'applications de gestion des informations. Ce logiciel ou matériel n'est pas conçu ni n'est destiné à être utilisé dans des applications à risque, notamment dans des applications pouvant causer des dommages corporels. Si vous utilisez ce logiciel ou matériel dans le cadre d'applications dangereuses, il est de votre responsabilité de prendre toutes les mesures de secours, de sauvegarde, de redondance et autres mesures nécessaires à son utilisation dans des conditions optimales de sécurité. Oracle Corporation et ses affiliés déclinent toute responsabilité quant aux dommages causés par l'utilisation de ce logiciel ou matériel pour ce type d'applications.

Oracle et Java sont des marques déposées d'Oracle Corporation et/ou de ses affiliés. Tout autre nom mentionné peut correspondre à des marques appartenant à d'autres propriétaires qu'Oracle.

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

Ce logiciel ou matériel et la documentation qui l'accompagne peuvent fournir des informations ou des liens donnant accès à des contenus, des produits et des services émanant de tiers. Oracle Corporation et ses affiliés déclinent toute responsabilité ou garantie expresse quant aux contenus, produits ou services émanant de tiers. En aucun cas, Oracle Corporation et ses affiliés ne sauraient être tenus pour responsables des pertes subies, des coûts occasionnés ou des dommages causés par l'accès à des contenus, produits ou services tiers, ou à leur utilisation.

Contents

Using This Documentation	7
1 Preparing to Upgrade Oracle Solaris Cluster Software	9
Upgrade Requirements and Software Support Guidelines	9
Choosing an Oracle Solaris Cluster Upgrade Method	11
Standard Upgrade	11
Dual-Partition Upgrade	12
Rolling Upgrade	13
2 Upgrading Zones Managed by Oracle Solaris Cluster Software	15
Upgrading a solaris10 Brand Zone in a Zone Cluster	15
▼ How to Upgrade a solaris10 Brand Zone in a Zone Cluster (czonecluster)	16
▼ How to Upgrade a solaris10 Brand Zone in a Zone Cluster (patchadd)	16
▼ How to Upgrade a solaris10 Brand Zone in a Zone Cluster (scinstall)	16
▼ How to Upgrade a solaris10 Brand Zone in a Zone Cluster (Interactive scinstall)	17
Upgrading Failover Zones	18
▼ How to Upgrade a Failover Zone	18
3 Performing a Standard Upgrade	21
Performing a Standard Upgrade of a Cluster	21
▼ How to Upgrade Quorum Server Software	22
▼ How to Prepare the Cluster for Upgrade (Standard Upgrade)	24
▼ How to Upgrade the Software (Standard Upgrade)	25
4 Performing a Dual-Partition Upgrade	29
Performing a Dual-Partition Upgrade of a Cluster	29
▼ How to Upgrade Quorum Server Software	30

▼ How to Prepare the Cluster for Upgrade (Dual-Partition)	31
▼ How to Upgrade the Software (Dual-Partition)	36
5 Performing a Rolling Upgrade	41
Performing a Rolling Upgrade of a Cluster	41
▼ How to Upgrade Quorum Server Software	42
▼ How to Prepare a Cluster Node for Upgrade (Rolling Upgrade)	43
▼ How to Upgrade the Software (Rolling Upgrade)	46
6 Completing the Upgrade	49
Completing a Cluster Upgrade	49
▼ How to Commit the Upgraded Cluster	49
▼ How to Verify the Upgrade	50
▼ How to Finish the Upgrade	51
7 Recovering From an Incomplete Upgrade	57
Cluster Recovery After an Incomplete Upgrade	57
▼ How to Recover from a Failed Dual-Partition Upgrade	57
▼ x86: How to Recover From a Partially Completed Dual-Partition Upgrade	59
Recovering From Storage Configuration Changes During Upgrade	61
▼ How to Handle Storage Reconfiguration During an Upgrade	61
▼ How to Resolve Mistaken Storage Changes During an Upgrade	62
Index	63

Using This Documentation

- **Overview** – Contains guidelines and procedures to upgrade Oracle Solaris Cluster software.
- **Audience** – Experienced system administrators with extensive knowledge of Oracle software and hardware.
- **Required knowledge** – Knowledge of the Oracle Solaris operating system, of Oracle Solaris Cluster software, and expertise with the volume manager software that is used with Oracle Solaris Cluster software.

This document is not to be used as a planning or presales guide.

Product Documentation Library

Late-breaking information and known issues for this product are included in the documentation library at <http://www.oracle.com/pls/topic/lookup?ctx=E39579>.

Access to Oracle Support

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

Documentation Accessibility

For information about Oracle's commitment to accessibility, visit the Oracle Accessibility Program web site at <http://www.oracle.com/pls/topic/lookup?ctx=acc&id=docacc>.

Feedback

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

Preparing to Upgrade Oracle Solaris Cluster Software

This chapter provides the following information and procedures to prepare to upgrade or update the Oracle Solaris Cluster 4.2 software:

- [“Upgrade Requirements and Software Support Guidelines” on page 9](#)
- [“Choosing an Oracle Solaris Cluster Upgrade Method” on page 11](#)

An upgrade moves the cluster to the latest Oracle Solaris Cluster release by updating all packages. By default, the upgrade also upgrades the Oracle Solaris OS and any other packages to the latest compatible version.

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

Note - Applying an Oracle Solaris Cluster 4.2 Core SRU does not provide the same result as upgrading the software to the Oracle Solaris Cluster 4.2 release.

Upgrade Requirements and Software Support Guidelines

Observe the following requirements and software-support guidelines when you upgrade to the Oracle Solaris Cluster 4.2 software:

- **Minimum Oracle Solaris OS** - Oracle Solaris 11.1 is the minimum required to support Oracle Solaris Cluster 4.2 software.

Note - To upgrade only the Oracle Solaris OS from version 11.1 to 11.2, the cluster must be installed with at least Oracle Solaris Cluster 4.1 SRU8.

- **Upgrade paths** – Oracle Solaris Cluster 4.2 software supports the following direct upgrade paths:
 - **Upgrade from at least version 4.0 SRU 2** - Use the standard or dual-partition method.
 - **Upgrade from at least version 4.1** – You can also use the rolling upgrade method.
- **Supported hardware** - The cluster hardware must be a supported configuration for Oracle Solaris Cluster 4.2 software. Contact your Oracle representative for information about current supported Oracle Solaris Cluster configurations.
- **Architecture changes during upgrade** - Oracle Solaris Cluster 4.2 software does not support upgrade between architectures.
- **Software migration** - Do not migrate from one type of software product to another product during Oracle Solaris Cluster upgrade. Perform only software configuration changes that are specified by upgrade procedures of an installed software product.
- **Upgrading to compatible versions** - You must upgrade all software on the cluster nodes to a version that is supported by Oracle Solaris Cluster 4.2 software. *For example, if an application version is supported on Oracle Solaris Cluster 4.1 software but is not supported on Oracle Solaris Cluster 4.2 software, you must upgrade to an application version that is supported on Oracle Solaris Cluster 4.2 software, if such a version of the application exists.* See the [Oracle Solaris Cluster 4 Compatibility Guide \(http://www.oracle.com/technetwork/server-storage/solaris-cluster/overview/solariscluster4-compatibilityguide-1429037.pdf\)](http://www.oracle.com/technetwork/server-storage/solaris-cluster/overview/solariscluster4-compatibilityguide-1429037.pdf) for information about supported products (requires My Oracle Support login).
- **Migration from Oracle Solaris 10** - You can alternatively migrate Oracle Solaris Cluster 3.3 software running on the Oracle Solaris 10 OS to a solaris10 brand zone cluster. On a global cluster that already runs Oracle Solaris Cluster 4.2 software on Oracle Solaris 11 software, create a new solaris10 brand zone cluster. Then migrate to the zone cluster your Oracle Solaris Cluster 3.3 configuration and applications. See [Chapter 6, “Creating Zone Clusters,”](#) in [“Oracle Solaris Cluster Software Installation Guide”](#) and [“Creating the Image for Directly Migrating Oracle Solaris 10 Systems Into Zones”](#) in [“Creating and using Oracle Solaris 10 Zones”](#).
- **Zone cluster upgrade** – To upgrade the Oracle Solaris OS or Oracle Solaris Cluster software that is running in a zone cluster, you must upgrade the underlying global cluster. Do not attempt to upgrade the OS or cluster software directly in the zone cluster. The exception is upgrade of Geographic Edition software, which you must upgrade directly in a zone cluster.

Note - This requirement does not apply to a solaris10 brand zone cluster, which is handled differently.

- **Failover zone upgrade** – You can upgrade failover zones only when using the standard upgrade or rolling upgrade method. You cannot upgrade failover zones if you are using the dual-partition upgrade method.
- **Downgrade** - Oracle Solaris Cluster 4.2 software does not support any downgrade of Oracle Solaris Cluster software.

- **Limitation of scinstall for data-service upgrades** - The `scinstall` upgrade utility only upgrades those data services that are provided with Oracle Solaris Cluster 4.2 software. You must manually upgrade any custom or third-party data services, unless they also have IPS packages and later versions are available in their publisher.

Choosing an Oracle Solaris Cluster Upgrade Method

The following matrix summarizes 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 of other products in the cluster, such as volume management software and other applications, for any additional upgrade requirements or restrictions.

Choose from the following methods to upgrade your Oracle Solaris Cluster 4.2 cluster software:

- [“Standard Upgrade” on page 11](#)
- [“Dual-Partition Upgrade” on page 12](#)
- [“Rolling Upgrade” on page 13](#)

For overview information about planning your Oracle Solaris Cluster 4.2 configuration, see [Chapter 1, “Planning the Oracle Solaris Cluster Configuration,”](#) in [“Oracle Solaris Cluster Software Installation Guide ”](#).

Standard Upgrade

A standard upgrade upgrades a cluster to a new release and upgrades the Oracle Solaris OS to the latest compatible version. You do not need to place the cluster in noncluster mode before performing this upgrade because the upgrade always occurs in the new boot environment (BE) and the existing BE remains unchanged. You can specify a name for the new BE or you can use the auto-generated name. All cluster nodes must be rebooted into the upgraded BE at the same time.

Any time you upgrade the Oracle Solaris Cluster software, by default you also upgrade the data services and Geographic Edition software. However, if you want to upgrade the data services separately, see [“Overview of the Installation and Configuration Process”](#) in [“Oracle Solaris Cluster Data Services Planning and Administration Guide ”](#). If you want to upgrade Oracle Solaris Cluster Geographic Edition separately, see the [“Oracle Solaris Cluster Geographic Edition Installation and Configuration Guide ”](#).

Note - If you want to update individual packages, use the `pkg` command. See [“Updating a Specific Package”](#) in [“Oracle Solaris Cluster System Administration Guide ”](#).

For complete instructions on performing a standard upgrade, see [Chapter 3, “Performing a Standard Upgrade”](#).

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

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.

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

- **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 method instead.
- **Failover zone upgrade** – You cannot upgrade failover zones if you are using the dual-partition upgrade method. You can only upgrade failover zones if you are using the standard upgrade or rolling upgrade method.
- **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.

Rolling Upgrade

In a rolling upgrade, you upgrade software to an update release, such as from Oracle Solaris version 11.1 to version 11.2, or from Oracle Solaris Cluster version 4.1 to version 4.2 or to a version 4.1 SRU. You perform the upgrade 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:

- **Oracle Solaris upgrade paths** - You can upgrade the Oracle Solaris OS only to a new SRU or an update version of the same release. For example, you can perform a rolling upgrade from Oracle Solaris 11.1 to a later compatible Oracle Solaris 11 release, but you cannot perform a rolling upgrade from a version of Oracle Solaris 10.
- **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 booted into its upgraded boot environment (BE), boot the next cluster node into its upgraded BE 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.

Upgrading Zones Managed by Oracle Solaris Cluster Software

This chapter provides the following information to upgrade Oracle Solaris non-global zones of `solaris` and `solaris10` brand types that are managed by the Oracle Solaris Cluster software. For more information about making a non-global zone highly available, see [“Oracle Solaris Cluster Data Service for Oracle Solaris Zones Guide”](#).

- [“Upgrading a `solaris10` Brand Zone in a Zone Cluster” on page 15](#)
- [“Upgrading Failover Zones” on page 18](#)

Upgrading a `solaris10` Brand Zone in a Zone Cluster

Before you upgrade a `solaris10` brand zone cluster, ensure that the following conditions are met:

- The version of Oracle Solaris Cluster software that runs in the zone cluster is supported for upgrade to the new version you want to run in the zone cluster.
- The version of Oracle Solaris Cluster software that runs in the global cluster supports the version of Oracle Solaris Cluster software that the zone cluster will upgrade to.

Use the `clzonecluster show-rev` subcommand to check the cluster release information for each node of the zone cluster. For more information, see the [`clzonecluster\(1CL\)` man page](#).

This section contains the following procedures:

- [“How to Upgrade a `solaris10` Brand Zone in a Zone Cluster \(`clzonecluster`\)” on page 16](#)
- [“How to Upgrade a `solaris10` Brand Zone in a Zone Cluster \(`patchadd`\)” on page 16](#)
- [“How to Upgrade a `solaris10` Brand Zone in a Zone Cluster \(`scinstall`\)” on page 16](#)
- [“How to Upgrade a `solaris10` Brand Zone in a Zone Cluster \(Interactive `scinstall`\)” on page 17](#)

▼ How to Upgrade a solaris10 Brand Zone in a Zone Cluster (clzonecluster)

1. **Bring the zone cluster into an offline-running state.**

```
phys-schost# clzonecluster reboot -o zoneclustername
```

2. **Issue the following command from the command line of *one* node of the global cluster:**

```
phys-schost#clzonecluster install-cluster -p patchdir=patchdir \  
[,patchlistfile=filename] zoneclustername
```

Use the new release patches for the patchlistfile property. For more information about the clzonecluster install-cluster subcommand, see the [clzonecluster\(1CL\)](#) man page.

Troubleshooting If the scinstall -u update command fails with a message that the pkg command is out of date, follow the instructions in the message.

▼ How to Upgrade a solaris10 Brand Zone in a Zone Cluster (patchadd)

Perform the following steps on each configured zone-cluster node.

1. **Bring the zone cluster into an offline-running state.**

```
phys-schost# clzonecluster reboot -o zoneclustername
```

2. **Log in to the zonecluster node.**

```
phys-schost# zlogin zoneclustername
```

3. **From the command line, type patchadd.**

```
zchost# patchadd
```

4. **Install the patches that correspond to the new release.**

▼ How to Upgrade a solaris10 Brand Zone in a Zone Cluster (scinstall)

Perform the following steps on each zone-cluster node.

1. **Bring the zone cluster into an offline-running state.**

```
phys-schost# clzonecluster reboot -o zoneclustername
```

2. **Log in to the zone cluster.**

```
phys-schost# zlogin zoneclustername
```

3. **From the command line, type the following commands.**

```
zchost# cd new-release-dvd-path/Solaris_arch/Product/sun_cluster/Solaris_10/Tools/  
zchost# ./scinstall -u update
```

In the path, *arch* is sparc or x86.

▼ How to Upgrade a solaris10 Brand Zone in a Zone Cluster (Interactive scinstall)

Perform the following steps for each configured zone-cluster node.

1. **Bring the zone cluster into an offline-running state.**

```
phys-schost# clzonecluster reboot -o zoneclustername
```

2. **Log in to the zone cluster.**

```
phys-schost# zlogin zoneclustername
```

3. **Change to the DVD image directory inside the zone.**

You must run the `scinstall` utility from the latest DVD directory. For example:

```
zchost# cd /net/my-machine/export/dvd/Solaris_arch/Product/sun_cluster/Solaris_10/Tools/
```

In the path, *arch* is sparc or x86.

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

```
zchost# pwd  
/net/my-machine/export/dvd/Solaris_arch/Product/sun_cluster/Solaris_10/Tools/  
zchost# ./scinstall
```

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

The Upgrading Oracle Solaris Cluster on This Node menu is displayed.

6. **Follow the menu prompts to upgrade the solaris10 brand zone on this zone cluster node.**

This action upgrades the Oracle Solaris Cluster framework and data service agents on this node.

7. **Quit the `scinstall` utility.**

Upgrading Failover Zones

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

Note - You can upgrade a failover zone only when using the standard upgrade or rolling upgrade method. You cannot upgrade a failover zone when using the dual-partition upgrade method.

▼ How to Upgrade a Failover Zone

Note - This task does not apply to `solaris10` brand failover zones. To manage packages and patches in a `solaris10` brand failover zone, follow normal Oracle Solaris 10 Zone procedures. See “[Creating and using Oracle Solaris 10 Zones](#)”.

1. **Identify which resource groups that manage a failover zone of brand type `solaris` are online.**

```
# clresourcegroup status
```

Make a record for later reference of the online resource groups and on which cluster node each resource group is online.

2. **Identify the resource name for the `sczbt` components that manage the failover zone.**

Do this for each resource group that you identified in [Step 1](#).

```
# clresource status -g resource-group
```

Make a record for later reference of the resource names.

3. **Suspend each online resource group that manages a failover zone of brand type `solaris`.**

```
# clresourcegroup suspend resource-group
```

The resource groups remain online but they are prevented from failing the zones over during upgrade. While in this state, the resource groups are also prevented from automatically rebooting the zones when the nodes are rebooted after upgrade. The failover zone will be upgraded together with the node on which it is currently online.

4. Perform the standard or rolling upgrade on the nodes.

Perform all steps, including booting into the new boot environment, in [“How to Upgrade the Software \(Standard Upgrade\)” on page 25](#) or [“How to Upgrade the Software \(Rolling Upgrade\)” on page 46](#). Then return to this procedure.

5. Disable each sczbt resource that you recorded in [Step 2](#).

```
# clresource disable sczbt-resource
```

Disabling this resource ensures that the resource does not time out if the first boot of the upgraded zone takes more time than the resource's Start_timeout allows.

6. For each resource group that you recorded in [Step 1](#), bring the resource group online on the same node where it was online during the upgrade.

```
# clresourcegroup online -n node resource-group
```

The updated failover zones are currently only associated with the global zone boot environment, where they were upgraded, and can only successfully boot there.

7. Attach and boot the failover zone.

Do this for each zone that is managed by an sczbt resource that you identified in [Step 2](#).

```
# zoneadm -z zone-name attach
# zoneadm -z zone-name boot
```

8. Verify the zone console.

Do this for each zone that is managed by an sczbt resource that you identified in [Step 2](#).

```
# zlogin -C zone-name
```

9. Shut down and forcibly detach the failover zone.

Do this for each zone that is managed by an sczbt resource that you identified in [Step 2](#).

```
# zoneadm -z zone-name shutdown
# zoneadm -z zone-name detach -F
```

10. Enable each sczbt resource.

Do this for each zone that is managed by an sczbt resource that you identified in [Step 2](#).

```
# clresource enable sczbt-resource
```

11. Synchronize the universally unique IDs (UUID) of the updated boot environments.

Note - If you are using the Oracle Solaris 11.2 OS, this step is optional.

- If you omit this step, at failover of a `solaris` brand zone, the last zone boot environment that was booted is first cloned and then activated on the node.
 - If you perform this step, at failover of a `solaris` brand zone, the last zone boot environment that was booted is activated on the node without first creating a clone.
-

a. Determine the UUID of the updated boot-environment root dataset of one upgraded node.

```
phys-schost-1# /opt/SUNWsczone/sczbt/util/ha-solaris-zone-boot-env-id get [-
b bename]
UUID
```

`-b bename` Specifies the boot environment from which to retrieve the UUID.

`UUID` The UUID of the specified boot environment.

The `ha-solaris-zone-boot-env-id get` command output displays the UUID, similar to the following:

```
8fe53702-16c3-eb21-ed85-d19af92c6bbd
```

This UUID will be used in the next steps as the reference UUID for all other upgraded cluster nodes.

b. Set the reference UUID in [Step 11.a](#) on the updated boot environment for the other upgraded nodes.

```
phys-schost-N# /opt/SUNWsczone/sczbt/util/ha-solaris-zone-boot-env-id set [-
b bename] UUID
```

`-b bename` Specifies the boot environment in which to define the specified UUID.

`UUID` The reference UUID that you obtained in [Step 11.a](#).

12. Resume each resource group that you brought online in [Step 6](#).

```
# clresourcegroup resume resource-group
```

Next Steps Go to [Chapter 6, “Completing the Upgrade”](#).

◆◆◆ CHAPTER 3

Performing a Standard Upgrade

This chapter provides the following information to upgrade to Oracle Solaris Cluster 4.2 software or to a Software Repository Update (SRU) by using the standard nonrolling upgrade method:

- [“How to Upgrade Quorum Server Software” on page 22](#)
- [“How to Prepare the Cluster for Upgrade \(Standard Upgrade\)” on page 24](#)
- [“How to Upgrade the Software \(Standard Upgrade\)” on page 25](#)

Note - Upgrading to Oracle Solaris Cluster 4.2 software also upgrades the Oracle Solaris operating system to the latest compatible version.

Performing a Standard Upgrade of a Cluster

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

If you do not want to upgrade all of your software to the latest available version, you can prevent certain components from being upgraded. You can do this by freezing individual packages or incorporations, or by disabling the appropriate publisher. For instructions, see [Chapter 5, “Configuring Installed Images,” in “Adding and Updating Software in Oracle Solaris 11.2”](#) or the `pkg(1)` man page.

Any time you upgrade the Oracle Solaris Cluster software in the global cluster, by default you also upgrade the data services in the global cluster and in any zone clusters, and upgrade Geographic Edition software in the global cluster. However, you must always manually upgrade Geographic Edition software in a zone cluster. For information about upgrading Geographic Edition software, see [Chapter 3, “Upgrading or Updating Geographic Edition Software,” in “Oracle Solaris Cluster Geographic Edition Installation and Configuration Guide”](#).

Note - If you want to install any individual Oracle Solaris Cluster SRUs, instead see [“How to Update a Specific Package”](#) in [“Oracle Solaris Cluster System Administration Guide”](#).

The following table lists the tasks to upgrade to Oracle Solaris Cluster 4.2 software or to a 4.2 SRU. By default, all Oracle Solaris packages are automatically upgraded.

TABLE 3-1 Task Map: Performing a Standard Upgrade for Oracle Solaris Cluster 4.2 Software

Task	Instructions
1. Read the upgrade requirements and restrictions. Determine the proper upgrade method for your configuration and needs.	“Upgrade Requirements and Software Support Guidelines” on page 9 “Choosing an Oracle Solaris Cluster Upgrade Method” on page 11
2. If failover zones of brand type solaris are configured in the cluster, upgrade the failover zones.	“Upgrading Failover Zones” on page 18
3. If a quorum server is used, upgrade the Quorum Server software.	“How to Upgrade Quorum Server Software” on page 22
4. Remove the cluster from production and back up shared data.	“How to Prepare the Cluster for Upgrade (Standard Upgrade)” on page 24
5. Upgrade to Oracle Solaris Cluster 4.2 framework, data-service, and Geographic Edition software (global cluster only). If necessary, upgrade applications that support alternate boot environments.	“How to Upgrade the Software (Standard Upgrade)” on page 25
6. Use the <code>scversions</code> command to commit the cluster to the upgrade.	“How to Commit the Upgraded Cluster” on page 49
7. Verify successful completion of upgrade to Oracle Solaris Cluster 4.2 software.	“How to Verify the Upgrade” on page 50
8. Enable resources and bring resource groups online. Migrate existing resources to new resource types. If necessary, boot into noncluster mode and upgrade applications that do not support alternate boot environments.	“How to Finish the Upgrade” on page 51

▼ 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 the root role on the cluster and on the quorum server.

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

See [“Adding a Quorum Device”](#) in [“Oracle Solaris Cluster System Administration Guide”](#).

If you add another quorum server as a temporary quorum device, the quorum server can run the same software version as the quorum server that you are upgrading, or it can run the 4.2 version of Quorum Server software.

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

```
phys-schost# clquorum remove quorum-server
```

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

```
quorumserver# clquorumserver show +
```

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

Note - If you have unconfigured the quorum server from a cluster but the `clquorumserver show` command still reports that the quorum server is serving that cluster, the command might be reporting stale configuration information. See [“Cleaning Up Stale Quorum Server Cluster Information”](#) in [“Oracle Solaris Cluster System Administration Guide”](#).

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

```
quorumserver# clquorumserver stop +
```

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

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

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

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

7. Install the Oracle Solaris Cluster 4.2 Quorum Server software, reconfigure the quorum server, and start the quorum server daemon.

Follow the steps in [“How to Install and Configure Oracle Solaris Cluster Quorum Server Software”](#) in [“Oracle Solaris Cluster Software Installation Guide”](#) for installing the Quorum Server software.

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

Follow the steps in [“How to Configure Quorum Devices”](#) in [“Oracle Solaris Cluster Software Installation Guide”](#).

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

```
phys-schost# clquorum remove temp-quorum
```

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

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

Before You Begin Perform the following tasks:

- Ensure that the configuration meets the requirements for upgrade. See [“Upgrade Requirements and Software Support Guidelines” on page 9](#).
- Have available the installation media, documentation, and software updates for all software products that you are upgrading, including the following software:
 - Oracle Solaris OS
 - Oracle Solaris Cluster
 - Applications that are managed by Oracle Solaris Cluster data services
 - Any other third-party applications to upgrade

For instructions on updating single or multiple packages, see [Chapter 11, “Updating Your Software,” in “Oracle Solaris Cluster System Administration Guide”](#).

- If you use role-based access control (RBAC) instead of the root role to access the cluster nodes, ensure that you can become an administrator with rights for all Oracle Solaris Cluster commands. This series of upgrade procedures requires the following Oracle Solaris Cluster RBAC authorizations if the user is not the root role:
 - `solaris.cluster.modify`
 - `solaris.cluster.admin`
 - `solaris.cluster.read`

See [“User Rights Management” in “Securing Users and Processes in Oracle Solaris 11.2”](#) for more information about using RBAC roles. See the Oracle Solaris Cluster man pages for the RBAC authorization that each Oracle Solaris Cluster subcommand requires.

1. **Ensure that the cluster is functioning normally.**
 - a. **View the current status of the cluster by running the following command from any node.**

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

See the [cluster\(1CL\)](#) man page for more information.

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

- c. **Check the volume-manager status.**
2. **Assume the root role on a node of the cluster.**
3. **Ensure that all shared data is backed up.**
4. **Ensure that each system disk is backed up.**

Next Steps Upgrade the cluster and OS software. Go to [“How to Upgrade the Software \(Standard Upgrade\)” on page 25.](#)

▼ How to Upgrade the Software (Standard Upgrade)

Perform this procedure on each node of the cluster. You can perform upgrade tasks in parallel on multiple nodes by using the `pconsole` utility. For more information about the `pconsole` utility, see [“How to Install pconsole Software on an Administrative Console” in “Oracle Solaris Cluster Software Installation Guide”](#).

- Before You Begin**
- If you intend to use the interactive `scinstall` utility, ensure that all cluster nodes are installed with a minimum of Oracle Solaris Cluster 4.0 SRU 2 software.
 - If you have failover zones of brand type `solaris` configured on the cluster, you must perform additional steps. Follow the instructions in [“How to Upgrade a Failover Zone” on page 18](#) before you begin this procedure.
 - If you have a `solaris10` brand zone in a zone cluster, follow the instructions in [“Upgrading a solaris10 Brand Zone in a Zone Cluster” on page 15.](#)

Note - If you have failover zones of brand type `solaris` configured on the cluster, you must perform additional steps. Follow the instructions in [“How to Upgrade a Failover Zone” on page 18](#) before you begin this procedure. If you have a `solaris10` brand zone in a zone cluster, follow the instructions in [“Upgrading a solaris10 Brand Zone in a Zone Cluster” on page 15.](#)

1. **Assume the root role or become an administrator with `solaris.cluster.admin` RBAC authorization.**
2. **Subscribe to the `ha-cluster` publisher that contains the software you want to upgrade to.**

```
# pkg set-publisher -G '*' -g URL_for_ha-cluster_publisher ha-cluster
```
3. **Ensure that the `solaris` publisher is valid.**

```
# pkg publisher
```

PUBLISHER	TYPE	STATUS	P	LOCATION
solaris	origin	online	F	solaris-repository

For information about setting the solaris publisher, see [“Adding, Modifying, or Removing Package Publishers”](#) in [“Adding and Updating Software in Oracle Solaris 11.2”](#).

4. Run the upgrade.

If you are upgrading a failover zone, follow the instructions in [“How to Upgrade a Failover Zone”](#) on page 18.

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

```
# scinstall -u update [-b bename] [-L accept,licenses]
```

You can choose to specify a name for the new boot environment with the **-b *bename*** option. If you specify the **-L *accept,licenses*** option, the command accepts and displays the licenses of the packages you upgrade. See the [scinstall\(1M\)](#) man page for more information about these options.

- **To use the interactive scinstall utility, perform the following:**

Note - To use the interactive scinstall utility, the cluster nodes must be installed with at least Oracle Solaris Cluster version 4.0 SRU 2.

- a. **Start the scinstall utility.**

```
phys-schost# scinstall
```

The scinstall Main Menu is displayed.

- b. **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) Upgrade this cluster node
* 3) Print release information for this cluster node
```

```
* ?) Help with menu options
* q) Quit
```

```
Option: 2
```

The Upgrade Menu is displayed.

- c. **Choose the menu item, Upgrade This Cluster Node.**

d. Follow the menu prompts to upgrade the cluster framework, data service agents, and underlying OS.

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

e. Quit the `scinstall` utility.

5. **As needed, upgrade other applications that support alternate boot environments.**
6. **After all cluster nodes are upgraded on their new boot environments, from one node shut down the cluster.**

Note - Do not perform this step until the new boot environment on each cluster nodes has been upgraded and is ready to put into production. You must boot the new boot environment on all cluster nodes at the same time.

```
phys-schost-1# scshutdown
```

7. **Boot all machines into the new boot environment at the same time.**
 - **If you need to upgrade any applications that do not support alternate boot environments, boot into the new boot environment in noncluster mode.**

```
# reboot -- -x
```
 - **Otherwise, follow the prompts to boot the machine into the new boot environment.**

Next Steps Go to [Chapter 6, “Completing the Upgrade”](#).

◆◆◆ CHAPTER 4

Performing a Dual-Partition Upgrade

This chapter provides the following information to upgrade a multiple-node cluster to Oracle Solaris Cluster 4.2 software by using the dual-partition upgrade method:

- [“How to Prepare the Cluster for Upgrade \(Dual-Partition\)” on page 31](#)
- [“How to Upgrade the Software \(Dual-Partition\)” on page 36](#)

Note - You cannot upgrade a failover zone when using the dual-partition upgrade method.

Performing a Dual-Partition Upgrade of a Cluster

The following table lists the tasks to upgrade to Oracle Solaris Cluster 4.2 software or to a 4.2 SRU. By default, all Oracle Solaris packages are automatically upgraded.

TABLE 4-1 Task Map: Performing a Dual-Partition Upgrade to Oracle Solaris Cluster 4.2 Software

Task	Instructions
1. Read the upgrade requirements and restrictions. Determine the proper upgrade method for your configuration and needs.	“Upgrade Requirements and Software Support Guidelines” on page 9 “Choosing an Oracle Solaris Cluster Upgrade Method” on page 11
2. If a quorum server is used, upgrade the Quorum Server software.	“How to Upgrade Quorum Server Software” on page 30
3. Partition the cluster into two groups of nodes.	“How to Prepare the Cluster for Upgrade (Dual-Partition)” on page 31
4. Upgrade to Oracle Solaris Cluster 4.2 framework, data-service, and Geographic Edition software (global cluster only). If necessary, upgrade applications that support alternate boot environments.	“How to Upgrade the Software (Dual-Partition)” on page 36
5. Use the <code>scversions</code> command to commit the cluster to the upgrade.	“How to Commit the Upgraded Cluster” on page 49
6. Verify successful completion of upgrade to Oracle Solaris Cluster 4.2 software.	“How to Verify the Upgrade” on page 50
7. Enable resources and bring resource groups online. Optionally, migrate existing resources to new resource types.	“How to Finish the Upgrade” on page 51

▼ 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 the root role on the cluster and on the quorum server.

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

See [“Adding a Quorum Device”](#) in [“Oracle Solaris Cluster System Administration Guide”](#).

If you add another quorum server as a temporary quorum device, the quorum server can run the same software version as the quorum server that you are upgrading, or it can run the 4.2 version of Quorum Server software.

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

```
phys-schost# clquorum remove quorum-server
```

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

```
quorumserver# clquorumserver show +
```

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

Note - If you have unconfigured the quorum server from a cluster but the `clquorumserver show` command still reports that the quorum server is serving that cluster, the command might be reporting stale configuration information. See [“Cleaning Up Stale Quorum Server Cluster Information”](#) in [“Oracle Solaris Cluster System Administration Guide”](#).

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

```
quorumserver# clquorumserver stop +
```

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

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

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

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

7. **Install the Oracle Solaris Cluster 4.2 Quorum Server software, reconfigure the quorum server, and start the quorum server daemon.**

Follow the steps in “[How to Install and Configure Oracle Solaris Cluster Quorum Server Software](#)” in “[Oracle Solaris Cluster Software Installation Guide](#)” for installing the Quorum Server software.

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

Follow the steps in “[How to Configure Quorum Devices](#)” in “[Oracle Solaris Cluster Software Installation Guide](#)”.

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

```
phys-schost# clquorum remove temp-quorum
```

▼ 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 refer to the two groups of nodes as the first partition and the second partition. The nodes that you assign to the second partition 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 \(Standard Upgrade\)](#)” on page 24.

Perform all steps from the global zone only.

Before You Begin Perform the following tasks:

- Ensure that the configuration meets the requirements for upgrade. See “[Upgrade Requirements and Software Support Guidelines](#)” on page 9.
- Have available the installation media, documentation, and software updates for all software products that you are upgrading, including the following software:
 - Oracle Solaris OS
 - Oracle Solaris Cluster
 - Applications that are managed by Oracle Solaris Cluster data services
 - Any other third-party applications to upgrade

For instructions on updating single or multiple packages, see [Chapter 11, “Updating Your Software,”](#) in “[Oracle Solaris Cluster System Administration Guide](#)”.

- If you use role-based access control (RBAC) instead of the root role to access the cluster nodes, ensure that you can become an administrator with rights for all Oracle Solaris Cluster commands. This series of upgrade procedures requires the following Oracle Solaris Cluster RBAC authorizations if the user is not the root role:

- `solaris.cluster.modify`
- `solaris.cluster.admin`
- `solaris.cluster.read`

See “[User Rights Management](#)” in “[Securing Users and Processes in Oracle Solaris 11.2](#)” for more information about using RBAC roles. See the Oracle Solaris Cluster man pages for the RBAC authorization that each Oracle Solaris Cluster subcommand requires.

1. Ensure that the cluster is functioning normally.

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

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

See the [cluster\(1CL\)](#) man page for more information.

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

- c. Check the volume-manager status.**

2. If Geographic Edition software is installed, ensure that all application resource groups have the `Auto_start_on_new_cluster` property set to `False`.

This setting prevents the unintended starting of application resource groups before cluster upgrade is completed.

```
# clresourcegroup show -p Auto_start_on_new_cluster resource-group
```

If necessary, change the property value to `False`.

```
# clresourcegroup set -p Auto_start_on_new_cluster=False resource-group
```

3. 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.

4. Assume the root role.

5. 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
```

6. **If you are upgrading a two-node cluster, skip to [Step 14](#).**

Otherwise, proceed to [Step 7](#) to determine the partitioning scheme to use. You determine which nodes each partition will contain, but interrupt the partitioning process. You 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.

7. **Assume the root role on a node of the cluster.**

8. **Start the `scinstall` utility in interactive mode.**

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

The `scinstall` Main Menu is displayed.

9. **Choose the menu item, `Manage a Dual-Partition Upgrade`.**

The `Manage a Dual-Partition Upgrade` Menu is displayed.

10. **Choose the menu item, `Display and Select Possible Partitioning Schemes`.**

11. **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 respond to this prompt in [Step 16](#) of this procedure.

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

13. On another node of the cluster, become superuser.

14. 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.

```
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.

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

15. Determine your next step.

- **If you are upgrading a two-node cluster, return to [Step 14](#) through [Step 11](#) 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 16](#).

- **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 16](#).

16. **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.
17. **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.
18. **After all nodes in the first partition are shut down, boot each node in that partition into noncluster mode.**
 - **SPARC:**

ok `boot -x`
 - **x86:**
 - a. **In the GRUB menu, use the arrow keys to select the appropriate Oracle Solaris entry and type `e` to edit its commands.**

For more information about GRUB based booting, see [“Booting a System” in “Booting and Shutting Down Oracle Solaris 11.2 Systems”](#).
 - b. **In the boot parameters screen, use the arrow keys to select the kernel entry and type `e` to edit the entry.**
 - c. **Add `-x` to the `multiboot` command to specify that the system boot into noncluster mode.**
 - d. **Press `Enter` to accept the change and return to the boot parameters screen.**

The screen displays the edited command.
 - e. **Type `b` to boot the node into noncluster mode.**

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

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

20. 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 RAC 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. Go to [“How to Upgrade the Software \(Dual-Partition\)”](#) on page 36.

▼ How to Upgrade the Software (Dual-Partition)

Perform this procedure to upgrade each node of the cluster to Oracle Solaris Cluster 4.2 software and, if necessary, upgrade the Oracle Solaris software.

Note - If you intend to upgrade failover zones, do not use the dual-partition upgrade method. Instead, use the standard upgrade or rolling upgrade method.

Perform all steps from the global zone only.

Before You Begin Perform the following tasks:

- Ensure that all steps in [“How to Prepare the Cluster for Upgrade \(Dual-Partition\)” on page 31](#) 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.

Tip - You can use the `pconsole` utility to perform this procedure on multiple nodes simultaneously. See the `pconsole(1)` man page that is installed with the `terminal/pconsole` package for more information.

1. **Assume the root role on a node that is a member of the partition that is in noncluster mode.**
2. **Subscribe to the `ha-cluster` publisher that contains the software you want to upgrade to.**

```
# pkg set-publisher -G '*' -g URL_for_ha-cluster_publisher ha-cluster
```

3. **Ensure that the `solaris` publisher is valid.**

```
# pkg publisher
PUBLISHER                                TYPE    STATUS  P  LOCATION
solaris                                  origin  online  F  solaris-repository
```

For information about setting the `solaris` publisher, see [“Adding, Modifying, or Removing Package Publishers” in “Adding and Updating Software in Oracle Solaris 11.2”](#).

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

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

The `scinstall` Main Menu is displayed.

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

The Upgrade Menu is displayed.

6. **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.

7. **Quit the `scinstall` utility.**

8. **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 [“Planning 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.

9. **As needed, manually upgrade any custom data services that are not supplied on the product media.**
10. **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.
11. **Install any Oracle Solaris Cluster 4.2 framework and data-service software updates.**
For instructions on updating your software, see [Chapter 11, “Updating Your Software,”](#) in [“Oracle Solaris Cluster System Administration Guide ”](#).
12. **Upgrade software applications that are installed on the cluster.**
Ensure that application levels are compatible with the current versions of Oracle Solaris Cluster and Oracle Solaris software. See your application documentation for installation instructions.

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

13. **Repeat all steps in this procedure up to this point on all remaining nodes that you need to upgrade in the partition.**
14. **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
```

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 in the second partition 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.**

15. **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 [Step 16](#).

a. **Boot each node in the second partition into noncluster mode.**

- **SPARC:**

```
ok boot -x
```

- **x86:**

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

For more information about GRUB based booting, see [“Booting a System”](#) in [“Booting and Shutting Down Oracle Solaris 11.2 Systems”](#).

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

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

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

The screen displays the edited command.

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

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

- b. **Upgrade the nodes in the second partition.**

Return to [Step 1](#).

- 16. **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 57.

Performing a Rolling Upgrade

This chapter provides procedures to perform a rolling upgrade of an Oracle Solaris Cluster 4 release to the Oracle Solaris Cluster 4.2 release, from the Oracle Solaris Cluster 4.2 release to a newer Oracle Solaris Cluster 4.2 SRU, or from a compatible Oracle Solaris 11 release to a newer Oracle Solaris 11 release or SRU.

In a rolling upgrade, you upgrade one cluster node at a time, while the other cluster nodes remain in production. After all nodes are upgraded and have rejoined the cluster, you must commit the cluster to the new software version before you can use any new features.

Note - Oracle Solaris Cluster software does not support rolling upgrade to Oracle Solaris 11 software from a previous marketing release such as Oracle Solaris 10 software, or to Oracle Solaris Cluster 4.2 software from a previous marketing release such as Oracle Solaris Cluster 3.3 software. You can only perform an Oracle Solaris Cluster rolling upgrade of Oracle Solaris Cluster or Oracle Solaris software to an update of the same marketing release, such as from the Oracle Solaris Cluster 4.1 release to the Oracle Solaris Cluster 4.2 release.

This chapter provides the following information to upgrade an Oracle Solaris Cluster 4 configuration to the Oracle Solaris Cluster 4.2 release, to an Oracle Solaris Cluster 4.2 SRU, or to an SRU of the Oracle Solaris OS by using the rolling upgrade method:

- [“Performing a Rolling Upgrade of a Cluster” on page 41](#)

Performing a Rolling Upgrade of a Cluster

The following table lists the tasks to upgrade to Oracle Solaris Cluster 4.2 software or to a 4.2 SRU. By default, all Oracle Solaris packages are automatically upgraded.

TABLE 5-1 Task Map: Performing a Rolling Upgrade to Oracle Solaris Cluster 4.2 Software

Task	Instructions
1. Read the upgrade requirements and restrictions.	“Upgrade Requirements and Software Support Guidelines” on page 9

Task	Instructions
2. If failover zones of brand type solaris are configured in the cluster, upgrade the failover zones.	“Upgrading Failover Zones” on page 18
3. If a quorum server is used, upgrade the Quorum Server software.	“How to Upgrade Quorum Server Software” on page 42
4. 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.	“How to Prepare a Cluster Node for Upgrade (Rolling Upgrade)” on page 43
5. Upgrade the cluster node Oracle Solaris Cluster 4.2 framework, data service, and Geographic Edition (global cluster only) software. If necessary, upgrade applications that support alternate boot environments.	“How to Upgrade the Software (Rolling Upgrade)” on page 46
6. Repeat Tasks 3 through 4 on each remaining node to upgrade.	
7. Use the <code>scversions</code> command to commit the cluster to the upgrade.	“How to Commit the Upgraded Cluster” on page 49
8. Verify successful completion of upgrade to Oracle Solaris Cluster 4.2 software.	“How to Verify the Upgrade” on page 50
9. Enable resources and bring resource groups online. Migrate existing resources to new resource types. If necessary, upgrade applications that do not support alternate boot environments.	“How to Finish the Upgrade” on page 51

▼ 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 the root role on the cluster and on the quorum server.

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

See [“Adding a Quorum Device”](#) in [“Oracle Solaris Cluster System Administration Guide”](#).

If you add another quorum server as a temporary quorum device, the quorum server can run the same software version as the quorum server that you are upgrading, or it can run the 4.2 version of Quorum Server software.

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

```
phys-schost# clquorum remove quorum-server
```

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

```
quorumserver# clquorumserver show +
```

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

Note - If you have unconfigured the quorum server from a cluster but the `clquorumserver` show command still reports that the quorum server is serving that cluster, the command might be reporting stale configuration information. See [“Cleaning Up Stale Quorum Server Cluster Information”](#) in [“Oracle Solaris Cluster System Administration Guide ”](#).

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

```
quorumserver# clquorumserver stop +
```

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

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

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

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

7. Install the Oracle Solaris Cluster 4.2 Quorum Server software, reconfigure the quorum server, and start the quorum server daemon.

Follow the steps in [“How to Install and Configure Oracle Solaris Cluster Quorum Server Software”](#) in [“Oracle Solaris Cluster Software Installation Guide ”](#) for installing the Quorum Server software.

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

Follow the steps in [“How to Configure Quorum Devices”](#) in [“Oracle Solaris Cluster Software Installation Guide ”](#).

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

```
phys-schost# clquorum remove temp-quorum
```

▼ How to Prepare a Cluster Node for Upgrade (Rolling Upgrade)

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

Before You Begin Perform the following tasks:

- Ensure that the configuration meets requirements for upgrade. See [“Upgrade Requirements and Software Support Guidelines”](#) on page 9.
- Have available the installation media, documentation, and upgrades for all the software products that you are upgrading, including the following software:
 - Oracle Solaris OS
 - Oracle Solaris Cluster
 - Applications that are managed by Oracle Solaris Cluster 4.2 data service agents
 - Any other third-party applications to upgrade

For instructions about updating single or multiple packages, see [Chapter 11, “Updating Your Software,”](#) in [“Oracle Solaris Cluster System Administration Guide”](#).

1. Ensure that the cluster is functioning normally.

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

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

See the [cluster\(1CL\)](#) man page for more information.

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

- c. **Check the volume-manager status.**

2. If Geographic Edition software is installed, ensure that all application resource groups have the `Auto_start_on_new_cluster` property set to `False`.

This setting prevents the unintended starting of application resource groups before cluster upgrade is completed.

```
# clresourcegroup show -p Auto_start_on_new_cluster resource-group
```

If necessary, change the property value to `False`.

```
# clresourcegroup set -p Auto_start_on_new_cluster=False resource-group
```

3. 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.

4. Assume the root role on a node of the cluster.

5. Evacuate resource groups and device groups from the node, except those used by failover zones on the node.

- **If the node is configured with failover zones, suspend the resource groups used by the failover zones and evacuate all other resource groups and device groups on the node.**

Failover zones refers to solaris brand zones that are configured with the HA for Oracle Solaris Zones data service.

```
# clresourcegroup suspend -n node-to-evacuate failover-zone-resource-group ...
# clresourcegroup evacuate -n node-to-evacuate resource-group-unused-by-failover-zone ...
# cldevicegroup switch -n node-to-evacuate device-group ...
```

Alternatively, you can use the `clnode evacuate` command to evacuate all resource groups and device groups from the node, then use the `clresourcegroup switch` command to move back to the node you will upgrade only those resource groups that are used by a failover zone. See the [clresourcegroup\(1CL\)](#) and [clnode\(1CL\)](#) man pages for more information.

For more information about failover zones, see [“Upgrading Failover Zones” on page 18](#).

- **If there are no failover zones configured on the node, evacuate all resource groups and device groups on the node.**

Note - Do not use this command if failover zones are configured on the node. If a resource group that is used by a failover zone is evacuated from the node, even though the resource group is suspended, the evacuate operation will take the resource group offline.

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

See the [clnode\(1CL\)](#) man page for more information.

6. Evacuate any resource groups that are running in a zone-cluster node that is on the node to upgrade.

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

7. Verify that the move was completed successfully.

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

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

Next Steps Go to [“How to Upgrade the Software \(Rolling Upgrade\)” on page 46](#).

▼ How to Upgrade the Software (Rolling Upgrade)

Perform this procedure to upgrade to Oracle Solaris Cluster 4.2 software or to an Oracle Solaris Cluster 4.2 SRU 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.

Before You Begin If you have failover zones of brand type `solaris` configured on the cluster, you must perform additional steps. Follow the instructions in [“How to Upgrade a Failover Zone” on page 18](#) before you begin this procedure. If you have a `solaris10` brand zone in a zone cluster, follow the instructions in [“Upgrading a solaris10 Brand Zone in a Zone Cluster” on page 15](#).

1. **Assume the root role or become an administrator with `solaris.cluster.admin` RBAC authorization on the node of the cluster you want to upgrade.**
2. **Subscribe to the `ha-cluster` publisher that contains the software you want to upgrade to.**

```
# pkg set-publisher -G '*' -g URL_for_ha-cluster_publisher ha-cluster
```

3. **Ensure that the `solaris` publisher is valid.**

```
# pkg publisher
PUBLISHER                                TYPE    STATUS  P  LOCATION
solaris                                   origin  online  F  solaris-repository
```

For information about setting the `solaris` publisher, see [“Adding, Modifying, or Removing Package Publishers” in “Adding and Updating Software in Oracle Solaris 11.2”](#).

4. **Run the upgrade on the node you want to upgrade.**

If you are upgrading a failover zone, follow the instructions in [“How to Upgrade a Failover Zone” on page 18](#).

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

```
# scinstall -u update [-b bename] [-L accept,licenses]
```

You can choose to specify a name for the new boot environment with the `-b bename` option. If you specify the `-L accept,licenses` option, the command accepts and displays the licenses of the packages you upgrade. See the [`scinstall\(1M\)`](#) man page for more information about these options.

- **To use the interactive `scinstall` utility, perform the following:**

Note - To use the interactive `scinstall` utility, the cluster nodes must be installed with at least 4.0 SRU 2.

a. Start the `scinstall` utility.

```
phys-schost# scinstall
```

The `scinstall` Main Menu is displayed.

b. Choose the Upgrade This Cluster menu item.

The Upgrade Menu is displayed.

c. If you are upgrading a global-cluster node, follow the menu prompts to upgrade the cluster framework, data service agents, or underlying OS.

If you are upgrading a zone cluster node, follow the menu prompts to upgrade the cluster framework, data service agents, or Oracle RAC for this node.

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

d. Quit the `scinstall` utility.

5. **As needed, upgrade other applications that support alternate boot environments.**
6. **Boot the node into the new boot environment.**
 - **If you need to upgrade any applications that do not support alternate boot environments, boot into the new boot environment in noncluster mode.**

```
# reboot -- -x
```
 - **Otherwise, follow the prompts to boot the node into the new boot environment.**
7. **Repeat these steps on each remaining node, one at a time, that you want to upgrade.**

Next Steps When all nodes in the cluster are upgraded, go to [Chapter 6, “Completing the Upgrade”](#).

Completing the Upgrade

This chapter provides the following information to complete all Oracle Solaris Cluster 4.2 software upgrade methods:

- [“How to Commit the Upgraded Cluster” on page 49](#)
- [“How to Verify the Upgrade” on page 50](#)
- [“How to Finish the Upgrade” on page 51](#)

Completing a Cluster Upgrade

▼ How to Commit the Upgraded Cluster

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

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

```
phys-schost# scversions
```

2. **From the following table, perform the action that is listed for the output message from [Step 1](#).**

Output Message	Action
Upgrade commit is needed.	Proceed to Step 3 .
Upgrade commit is NOT needed. All versions match.	Go to “How to Verify the Upgrade” on page 50 .
Upgrade commit cannot be performed until all cluster nodes are upgraded. Please run <code>scinstall(1m)</code> on cluster nodes to identify older versions.	Return to the Oracle Solaris Cluster upgrade procedures that you used and upgrade the remaining cluster nodes.
Check upgrade cannot be performed until all cluster nodes are upgraded. Please run <code>scinstall(1m)</code> on cluster nodes to identify older versions.	Return to the Oracle Solaris Cluster upgrade procedures that you used and upgrade the remaining cluster nodes.

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

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

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

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

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

Next Steps Go to [“How to Verify the Upgrade” on page 50](#).

▼ How to Verify the Upgrade

Perform this procedure to verify that the cluster is successfully upgraded to Oracle Solaris Cluster 4.2 software. Perform all steps from the global zone only.

- Before You Begin**
- Ensure that all upgrade procedures are completed for all cluster nodes that you are upgrading.
 - Ensure that all steps in [“How to Commit the Upgraded Cluster” on page 49](#) are completed successfully.

1. **On each node, assume the root role.**
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.

4. **From any node, view the boot environment (BE) created by the upgrade.**

```
# beadm list
```

Record the name of the upgraded BE and any other BEs that you might want to boot back into if needed.

Example 6-1 Verifying Upgrade to Oracle Solaris Cluster 4.2 Software

The following example shows the commands used to verify upgrade of a two-node cluster to Oracle Solaris Cluster 4.2 software. The cluster node names are `phys-schost-1` and `phys-schost-2`.

```
phys-schost# clnode show-rev -v
4.2
...
phys-schost# clnode status
=== Cluster Nodes ===

--- Node Status ---

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

Next Steps Go to [“How to Finish the Upgrade” on page 51](#).

▼ How to Finish the Upgrade

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

Before You Begin Ensure that all steps in [“How to Verify the Upgrade” on page 50](#) are completed.

- 1. 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.

- 2. If necessary, reset the `resource_security` property.**

After upgrade, the `resource_security` property for the cluster is reset to `COMPATIBLE`. To use a different security policy for RGM resources, run the following command from one node of the cluster:

```
phys-schost# cluster set -p resource_security=policy clustername
```

You can alternatively use the `clsetup` utility from the Other Cluster Tasks menu option. For more information about the `resource_security` property, see the [`cluster\(1CL\)`](#) man page.

- 3. Migrate resources to new resource type versions.**

You must migrate all resources to the Oracle Solaris Cluster 4.2 resource-type version to use the new features and bug fixes that are provided in this release.

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

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

Note - The Oracle Solaris Cluster 4.2 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.

4. In the global zone, re-enable all disabled resources and bring online all resource groups.

- **To use the `clsetup` utility, perform the following steps:**

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

```
phys-schost# clsetup
```

The `clsetup` Main Menu is displayed.

b. Choose the menu item, Resource Groups.

The Resource Group Menu is displayed.

c. Choose the menu item, Enable/Disable a Resource.

d. Choose a resource to enable and follow the prompts.

e. Repeat [Step 4.1.d](#) for each disabled resource.

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

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

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

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

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

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

a. **Enable each disabled resource.**

```
# clresource enable resource
```

b. **Verify that each resource is enabled.**

```
# clresource status
```

c. **Bring online each resource group.**

```
# clresourcegroup online -eM resourcegroup
```

d. **Verify that each resource group is online.**

```
# clresourcegroup status
```

5. **If zone clusters are configured in the cluster, in each zone cluster re-enable all disabled resources and bring online all resource groups.**

```
# clresourcegroup online -Z zonecluster resource-group
# clresource enable -Z zonecluster resource
# clresourcegroup online -eM -Z zonecluster resource-group
```

6. **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.**

7. **Enable the automatic reboot feature.**

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

-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.

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

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

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

9. Revalidate the upgraded cluster configuration.

See [“How to Validate the Cluster”](#) in [“Oracle Solaris Cluster Software Installation Guide”](#).

10. (Optional) Capture the ZFS root pool property information for future reference.

```
phys-schost# zpool get all rootpool > filename
```

Store the file in a location outside the cluster. If you make any root pool configuration changes, run this command again to capture the changed configuration. If necessary, you can use this information to restore the root pool partition configuration. For more information, see the [zpool\(1M\)](#) man page.

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

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

For more information, see [“How to Back Up the Cluster Configuration”](#) in [“Oracle Solaris Cluster System Administration Guide”](#).

Troubleshooting

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

```
phys-schost - Resource depends on a SUNW.HAStoragePlus type resource that is not
online anywhere.
```

```
(C189917) VALIDATE on resource nfsrs, resource group rg, exited with non-zero
exit status.
```

(C720144) Validation of resource nfsrs in resource group rg on node phys-schost failed.

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

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

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

```
No suitable Java runtime found. Java 1.7 or higher is required.
```

```
Jan 3 17:10:26 ppups3 cacao: No suitable Java runtime found. Java 1.7 or higher is required.
```

```
Cannot locate all the dependencies
```

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

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

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

Next Steps The cluster upgrade is complete.

Recovering From an Incomplete Upgrade

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

- [“How to Recover from a Failed Dual-Partition Upgrade” on page 57](#)
- [“How to Recover From a Partially Completed Dual-Partition Upgrade” on page 59](#)
- [“Recovering From Storage Configuration Changes During Upgrade” on page 61](#)

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 dual-partition 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. **Assume the root role on each node of the cluster.**
2. **Boot each node into noncluster mode.**
 - **SPARC:**

```
ok boot -x
```

- **x86:**
 - a. **In the GRUB menu, use the arrow keys to select the appropriate Oracle Solaris entry and type `e` to edit its commands.**

For more information about GRUB based booting, see [“Booting a System” in “Booting and Shutting Down Oracle Solaris 11.2 Systems”](#).
 - b. **In the boot parameters screen, use the arrow keys to select the kernel entry and type `e` to edit the entry.**
 - c. **Add `-x` to the `multiboot` command to specify that the system boot into noncluster mode.**
 - d. **Press Enter to accept the change and return to the boot parameters screen.**

The screen displays the edited command.
 - e. **Type `b` to boot the node into noncluster mode.**

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

3. On each node, run the upgrade recovery script.

```
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 3-1](#). 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.

▼ 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 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 57](#).

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 the root role.

1. Boot each node in the second partition into noncluster mode by completing the following steps.

- **SPARC:**

```
ok boot -x
```

- **x86:**

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

For more information about GRUB based booting, see [“Booting a System” in “Booting and Shutting Down Oracle Solaris 11.2 Systems”](#).

- b. In the boot parameters screen, use the arrow keys to select the kernel entry and type `e` to edit the entry.**
- c. Add `-x` to the `multiboot` command to specify that the system boot into noncluster mode.**
- d. Press Enter to accept the change and return to the boot parameters screen.**

The screen displays the edited command.
- e. Type `b` to boot the node into noncluster mode.**

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

- 2. On each node in the second partition, run the `scinstall -u recover` command.**

```
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.

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 61](#)
- [“How to Resolve Mistaken Storage Changes During an Upgrade” on page 62](#)

▼ 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 Oracle Solaris 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, assume the root role.**
2. **Manually update the unverified device for the node list.**

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

See the [cldevice\(1CL\)](#) man page for more information.

3. **Update the DID driver for the node list.**
- ```
phys-schost# cldevice refresh -n node[,...]
```
4. **Repeat [Step 2](#) through [Step 3](#) on all other nodes that are attached to the unverified device.**

**Next Steps** Return to the remaining upgrade tasks for your upgrade method.

## ▼ 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 61.

---

**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, assume the root role.**
2. **Update the DID driver on each node of the cluster.**

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

See the [cldevice\(1CL\)](#) man page for more information.

3. **If the `cldevice` 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 for your upgrade method.

# Index

---

## Numbers and Symbols

- /etc/system file
  - LOFS setting, 37
- /usr/cluster/bin/cldevice command
  - migrating device IDs after upgrade, 61
- /usr/cluster/bin/clnode command
  - moving resource groups and device groups, 45
  - verifying cluster mode, 50
  - verifying device-ID migration, 50
- /usr/cluster/bin/clresourcegroup command
  - displaying status, 18
  - resuming, 20
  - suspending, 18
- /usr/cluster/bin/scinstall command
  - verifying Oracle Solaris Cluster software, 50
- /usr/cluster/bin/scversions command, 49

## C

- cldevice command
  - migrating device IDs after upgrade, 61
- clnode command
  - moving resource groups and device groups, 45
  - verifying cluster mode, 50
  - verifying device-ID migration, 50
- clresourcegroup command
  - displaying status, 18
  - resuming, 20
  - suspending, 18
- cluster mode
  - verifying, 50
- clzonecluster show- rev command
  - running, 15
- committing the upgrade, 49

## D

- data services
  - upgrading
    - rolling upgrade, 46
    - standard upgrade, 25
- device groups
  - moving, 45
  - verifying evacuation, 45
- device-ID names
  - migrating after upgrade, 61
- DID driver
  - updating, 62
- disabling
  - LOFS after upgrade, 37
- dual-partition upgrade, 29
  - Oracle Solaris Cluster software, 36
  - preparing the cluster, 31
  - recovering from unrecoverable error, 57
  - requirements, 12
  - troubleshooting, 40

## E

- evacuating *See* moving
- examples
  - live upgrade verification, 51

## F

- failover zones
  - upgrading, 18

## H

- HA for Oracle Solaris Zones

- upgrading, 18
- L**
- LOFS
  - disabling after upgrade, 37
- loopback file system (LOFS)
  - disabling after upgrade, 37
- M**
- moving resource groups and device groups, 45
- N**
- nonrolling upgrade *See* standard upgrade
- O**
- Oracle Solaris software
  - publisher, 25, 37, 46
  - verifying device-ID migration, 50
- P**
- publisher
  - Oracle Solaris software, 25, 37, 46
- Q**
- quorum server
  - upgrading, 22
- R**
- recovering, 57, 57
  - See also* repairing
  - See also* troubleshooting
  - dual-partition upgrade unrecoverable error, 57
  - storage configuration changes during upgrade, 61
- repairing, 57, 57
  - See also* recovering
  - See also* troubleshooting
- requirements
  - all upgrade methods, 9
  - dual-partition upgrade, 12
  - rolling upgrade, 13
  - standard upgrade, 11
- resource groups
  - displaying status, 18
  - moving, 45
  - resuming, 20
  - suspending, 18
  - verifying evacuation, 45
- resource types
  - registering after upgrade, 51
- rolling upgrade, 41
  - Oracle Solaris Cluster software, 46
  - preparing the cluster, 43
  - requirements, 13
- storage reconfiguration during upgrade, 61
- S**
- scinstall command
  - upgrading Oracle Solaris Cluster software, 26, 37, 47
  - verifying Oracle Solaris Cluster software, 50
- scversions command, 49
- solaris brand failover zone
  - upgrading, 18
- solaris10 brand zone
  - upgrading, 15
- standard upgrade, 21
  - data services, 25
  - Oracle Solaris Cluster software, 25
  - Oracle Solaris software, 25, 25
  - preparing the cluster, 24
  - requirements, 11
- T**
- troubleshooting, 57, 57
  - See also* recovering
  - See also* repairing
  - dual-partition upgrade, 40
  - recovering from unrecoverable error during dual-partition upgrade, 57



storage configuration changes during upgrade, 61

device-ID migration, 50

Oracle Solaris Cluster software version, 50

resource group configurations, 45

## U

### upgrading

choosing an upgrade method, 11

committing the cluster, 49

dual-partition upgrade, 29

Oracle Solaris Cluster software, 36

preparing the cluster, 31

troubleshooting, 40

HA for Oracle Solaris Zones, 18

nonrolling upgrade *See* standard upgrade

quorum server software, 22

### recovering

storage configuration changes, 61

unrecoverable error during dual-partition

upgrade, 57

requirements, 9

resource types, 51

### rolling

preparing the cluster, 43

rolling upgrade, 41

data services, 46

Oracle Solaris Cluster software, 46

solaris brand failover zone, 18

solaris10 brand zone, 15

standard upgrade, 21

data services, 25

Oracle Solaris Cluster software, 25

Oracle Solaris software, 25, 25

preparing the cluster, 24

### troubleshooting

dual-partition upgrade, 40

recovering from unrecoverable error during dual-

partition upgrade, 57

storage configuration changes, 61

### verifying

device-ID conversion, 50

Oracle Solaris Cluster, 50

## V

### verifying

device group configurations, 45

