Oracle® Grid Infrastructure

Grid Infrastructure Installation and Upgrade Guide

21c for HP-UX Itanium

F31406-01

September 2021

ORACLE®

Oracle Grid Infrastructure Grid Infrastructure Installation and Upgrade Guide, 21c for HP-UX Itanium

F31406-01

Contents

4   Configuring Networks for Oracle Grid Infrastructure and Oracle RAC

5   Configuring Users, Groups and Environments for Oracle Grid Infrastructure and Oracle Database

6   Supported Storage Options for Oracle Database and Oracle Grid Infrastructure

7   Configuring Storage for Oracle Grid Infrastructure

8 Installing Oracle Grid Infrastructure

9 Oracle Grid Infrastructure Postinstallation Tasks

10 Upgrading Oracle Grid Infrastructure

11 Removing Oracle Database Software

A Completing Preinstallation Tasks Manually

B Installing and Configuring Oracle Database Using Response Files

C   Optimal Flexible Architecture

List of Tables

Preface

This guide explains how to configure a server in preparation for installing and configuring an Oracle Grid Infrastructure installation (Oracle Clusterware and Oracle Automatic Storage Management).

It also explains how to configure a server and storage in preparation for an Oracle Real Application Clusters (Oracle RAC) installation.

- Audience
- Documentation Accessibility
- Diversity and Inclusion
- Set Up Java Access Bridge to Implement Java Accessibility
  Install Java Access Bridge so that assistive technologies on Microsoft Windows systems can use the Java Accessibility API.
- Command Syntax
  Refer to these command syntax conventions to understand command examples in this guide.
- Conventions

Audience

This guide provides configuration information for network and system administrators, and database installation information for database administrators (DBAs) who install and configure Oracle Clusterware and Oracle Automatic Storage Management in an Oracle Grid Infrastructure for a cluster installation.

For users with specialized system roles who intend to install Oracle RAC, this book is intended to be used by system administrators, network administrators, or storage administrators to configure a system in preparation for an Oracle Grid Infrastructure for a cluster installation, and complete all configuration tasks that require operating system root privileges. When Oracle Grid Infrastructure installation and configuration is completed successfully, a system administrator should only need to provide configuration information and to grant access to the database administrator to run scripts as root during an Oracle RAC installation.

This guide assumes that you are familiar with Oracle Database concepts.

Documentation Accessibility

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

Access to Oracle Support

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

Diversity and Inclusion

Oracle is fully committed to diversity and inclusion. Oracle respects and values having a diverse workforce that increases thought leadership and innovation. As part of our initiative to build a more inclusive culture that positively impacts our employees, customers, and partners, we are working to remove insensitive terms from our products and documentation. We are also mindful of the necessity to maintain compatibility with our customers' existing technologies and the need to ensure continuity of service as Oracle's offerings and industry standards evolve. Because of these technical constraints, our effort to remove insensitive terms is ongoing and will take time and external cooperation.

Set Up Java Access Bridge to Implement Java Accessibility

Install Java Access Bridge so that assistive technologies on Microsoft Windows systems can use the Java Accessibility API.

Java Access Bridge is a technology that enables Java applications and applets that implement the Java Accessibility API to be visible to assistive technologies on Microsoft Windows systems.

Refer to *Java Platform, Standard Edition Accessibility Guide* for information about the minimum supported versions of assistive technologies required to use Java Access Bridge. Also refer to this guide to obtain installation and testing instructions, and instructions for how to use Java Access Bridge.

Related Topics

• *Java Platform, Standard Edition Java Accessibility Guide*

Command Syntax

Refer to these command syntax conventions to understand command examples in this guide.

Convention	Description
`$`	Bourne or BASH shell prompt in a command example. Do not enter the prompt as part of the command.
`%`	C Shell prompt in a command example. Do not enter the prompt as part of the command.
`#`	Superuser (root) prompt in a command example. Do not enter the prompt as part of the command.
`monospace`	UNIX command syntax

Convention	Description

backslash \	A backslash is the UNIX and Linux command continuation character. It is used in command examples that are too long to fit on a single line. Enter the command as displayed (with a backslash) or enter it on a single line without a backslash:
	
	`dd if=/dev/rdsk/c0t1d0s6 of=/dev/rst0 bs=10b \`
	`count=10000`

braces { }	Braces indicate required items: `.DEFINE {macro1}`
brackets [ ]	Brackets indicate optional items: `cvtcrt termname [outfile]`
ellipses ...	Ellipses indicate an arbitrary number of similar items: `CHKVAL fieldname value1 value2 ... valueN`
italic	Italic type indicates a variable. Substitute a value for the variable: `library_name`
vertical line |	A vertical line indicates a choice within braces or brackets: `FILE filesize [K|M]`

Conventions

The following text conventions are used in this document:

Convention	Meaning
boldface	Boldface type indicates graphical user interface elements associated with an action, or terms defined in text or the glossary.
italic	Italic type indicates book titles, emphasis, or placeholder variables for which you supply particular values.
`monospace`	Monospace type indicates commands within a paragraph, URLs, code in examples, text that appears on the screen, or text that you enter.

1

Oracle Grid Infrastructure Installation Checklist

Use checklists to plan and carry out Oracle Grid Infrastructure (Oracle Clusterware and Oracle Automatic Storage Management) installation.

Oracle recommends that you use checklists as part of your installation planning process. Using this checklist can help you to confirm that your server hardware and configuration meet minimum requirements for this release, and to ensure you carry out a successful installation.

- Server Hardware Checklist for Oracle Grid Infrastructure
  Review server hardware requirements for Oracle Grid Infrastructure installation.

- Operating System Checklist for Oracle Database on HP-UX Itanium
  Use this checklist to check minimum operating system requirements for Oracle Database.

- Server Configuration Checklist for Oracle Grid Infrastructure
  Use this checklist to check minimum server configuration requirements for Oracle Grid Infrastructure installations.

- Network Checklist for Oracle Grid Infrastructure
  Review this network checklist for Oracle Grid Infrastructure installation to ensure that you have required hardware, names, and addresses for the cluster.

- User Environment Configuration Checklist for Oracle Grid Infrastructure
  Use this checklist to plan operating system users, groups, and environments for Oracle Grid Infrastructure installation.

- Storage Checklist for Oracle Grid Infrastructure
  Review the checklist for storage hardware and configuration requirements for Oracle Grid Infrastructure installation.

- Installer Planning Checklist for Oracle Grid Infrastructure
  Review the checklist for planning your Oracle Grid Infrastructure installation before starting Oracle Universal Installer.

Server Hardware Checklist for Oracle Grid Infrastructure

Review server hardware requirements for Oracle Grid Infrastructure installation.

Table 1-1   Server Hardware Checklist for Oracle Grid Infrastructure

Check	Task
Server make and architecture	Confirm that server makes, models, core architecture, and host bus adaptors (HBA) are supported to run with Oracle Grid Infrastructure and Oracle RAC.
Runlevel	3
Server Display Cards	At least 1024 x 768 display resolution for Oracle Universal Installer. Confirm display monitor.

Table 1-1   (Cont.) Server Hardware Checklist for Oracle Grid Infrastructure

Check	Task
Minimum Random Access Memory (RAM)	At least 8 GB RAM for Oracle Grid Infrastructure installations.
Intelligent Platform Management Interface (IPMI)	IPMI cards installed and configured, with IPMI administrator account information available to the person running the installation.
	Ensure baseboard management controller (BMC) interfaces are configured, and have an administration account username and password to provide when prompted during installation.

Operating System Checklist for Oracle Database on HP-UX Itanium

Use this checklist to check minimum operating system requirements for Oracle Database.

Table 1-2   Operating System General Checklist for Oracle Database on HP-UX Itanium

Item	Task
Operating system general requirements	The following HP-UX Itanium version is supported:
	HP-UX 11.31 May 2019 B.11.31.1905.441b or later
	Review the system requirements section for a list of patch requirements.

Server Configuration Checklist for Oracle Grid Infrastructure

Use this checklist to check minimum server configuration requirements for Oracle Grid Infrastructure installations.

Table 1-3   Server Configuration Checklist for Oracle Grid Infrastructure

Check	Task
Disk space allocated to the temporary file system	At least 1 GB of space in the temporary disk space (`/tmp`) directory

Table 1-3   (Cont.) Server Configuration Checklist for Oracle Grid Infrastructure

Check	Task
Swap space allocation relative to RAM	Between 4 GB and 16 GB: Equal to RAM
	More than 16 GB: 16 GB
	Note: If you enable HugePages for your Linux servers, then you should deduct the memory allocated to HugePages from the available RAM before calculating swap space.
Mount point paths for the software binaries	Oracle recommends that you create an Optimal Flexible Architecture configuration as described in the appendix "Optimal Flexible Architecture" in *Oracle Grid Infrastructure Installation and Upgrade Guide* for your platform.
Ensure that the Oracle home (the Oracle home path you select for Oracle Database) uses only ASCII characters	The ASCII character restriction includes installation owner user names, which are used as a default for some home paths, as well as other directory names you may select for paths.
Set locale (if needed)	Specify the language and the territory, or locale, in which you want to use Oracle components. A locale is a linguistic and cultural environment in which a system or program is running. NLS (National Language Support) parameters determine the locale-specific behavior on both servers and clients. The locale setting of a component determines the language of the user interface of the component, and the globalization behavior, such as date and number formatting.
Set Network Time Protocol for Cluster Time Synchronization	Oracle Clusterware requires the same time zone environment variable setting on all cluster nodes.
	Ensure that you set the time zone synchronization across all cluster nodes using either an operating system configured network time protocol (NTP) or Oracle Cluster Time Synchronization Service.

Network Checklist for Oracle Grid Infrastructure

Review this network checklist for Oracle Grid Infrastructure installation to ensure that you have required hardware, names, and addresses for the cluster.

About Oracle Grid Infrastructure Network Configuration

During installation, you designate interfaces for use as public, private, or Oracle ASM interfaces. You can also designate interfaces that are in use for other purposes, such as a network file system, and not available for Oracle Grid Infrastructure use.

If you use a third-party cluster software, then the public host name information is obtained from that software.

Table 1-4   Network Configuration Tasks for Oracle Grid Infrastructure and Oracle RAC

Check	Task
Public network hardware	• Public network switch (redundant switches recommended) connected to a public gateway and to the public interface ports for each cluster member node. • Ethernet interface card (redundant network cards recommended, bonded as one Ethernet port name). • The switches and network interfaces must be at least 1 GbE. • The network protocol is Transmission Control Protocol (TCP) and Internet Protocol (IP).
Private network hardware for the interconnect	• Private dedicated network switches (redundant switches recommended), connected to the private interface ports for each cluster member node. **Note:** If you have more than one private network interface card for each server, then Oracle Clusterware automatically associates these interfaces for the private network using Grid Interprocess Communication (GIPC) and Grid Infrastructure Redundant Interconnect, also known as Cluster High Availability IP (HAIP). • The switches and network interface adapters must be at least 1 GbE. • The interconnect must support the user datagram protocol (UDP). • Jumbo Frames (Ethernet frames greater than 1500 bits) are not an IEEE standard, but can reduce UDP overhead if properly configured. Oracle recommends the use of Jumbo Frames for interconnects. However, be aware that you must load-test your system, and ensure that they are enabled throughout the stack.
Oracle Flex ASM Network Hardware	Oracle Flex ASM can use either the same private networks as Oracle Clusterware, or use its own dedicated private networks. Each network can be classified PUBLIC or PRIVATE+ASM or PRIVATE or ASM. Oracle ASM networks use the TCP protocol.

Table 1-4   (Cont.) Network Configuration Tasks for Oracle Grid Infrastructure and Oracle RAC

Check	Task
Cluster Names and Addresses	Determine and configure the following names and addresses for the cluster: • **Cluster name**: Decide a name for the cluster, and be prepared to enter it during installation. The cluster name should have the following characteristics: Globally unique across all hosts, even across different DNS domains. At least one character long and less than or equal to 15 characters long. Consist of the same character set used for host names, in accordance with RFC 1123: Hyphens (-), and single-byte alphanumeric characters (a to z, A to Z, and 0 to 9). If you use third-party vendor clusterware, then Oracle recommends that you use the vendor cluster name. • **Grid Naming Service Virtual IP Address (GNS VIP)**: If you plan to use GNS, then configure a GNS name and fixed address in DNS for the GNS VIP, and configure a subdomain on your DNS delegated to the GNS VIP for resolution of cluster addresses. GNS domain delegation is mandatory with dynamic public networks (DHCP, autoconfiguration). • **Single Client Access Name (SCAN) and addresses** **Using Grid Naming Service Resolution**: Do not configure SCAN names and addresses in your DNS. SCAN names are managed by GNS. **Using Manual Configuration and DNS resolution**: Configure a SCAN name to resolve to three addresses on the domain name service (DNS).
Node Public, Private and Virtual IP names and Addresses	If you are not using GNS, then configure the following for each node: • **Public node name and address**, configured in the DNS and in `/etc/hosts` (for example, `node1.example.com`, address 192.0.2.10). The public node name should be the primary host name of each node, which is the name displayed by the `hostname` command. • **Private node address**, configured on the private interface for each node. The private subnet that the private interfaces use must connect all the nodes you intend to have as cluster members. Oracle recommends that the network you select for the private network uses an address range defined as private by RFC 1918. • **Public node virtual IP name and address** (for example, `node1-vip.example.com`, address 192.0.2.11). If you are not using dynamic networks with GNS and subdomain delegation, then determine a virtual host name for each node. A virtual host name is a public node name that is used to reroute client requests sent to the node if the node is down. Oracle Database uses VIPs for client-to-database connections, so the VIP address must be publicly accessible. Oracle recommends that you provide a name in the format *hostname*-vip. For example: `myclstr2-vip`.

User Environment Configuration Checklist for Oracle Grid Infrastructure

Use this checklist to plan operating system users, groups, and environments for Oracle Grid Infrastructure installation.

Table 1-5   User Environment Configuration for Oracle Grid Infrastructure

Check	Task
Review Oracle Inventory (oraInventory) and OINSTALL Group Requirements	The Oracle Inventory directory is the central inventory of Oracle software installed on your system. It should be the primary group for all Oracle software installation owners. Users who have the Oracle Inventory group as their primary group are granted the OINSTALL privilege to read and write to the central inventory.
	• If you have an existing installation, then OUI detects the existing oraInventory directory from the `/etc/oraInst.loc` file, and uses this location.
	• If you are installing Oracle software for the first time, then OUI creates an Oracle base and central inventory, and creates an Oracle inventory using information in the following priority:
	– In the path indicated in the ORACLE_BASE environment variable set for the installation owner user account.
	– In an Optimal Flexible Architecture (OFA) path (u[01–99]/app/*owner* where *owner* is the name of the user account running the installation), if that user account has permissions to write to that path.
	– In the user home directory, in the path /app/*owner*, where *owner* is the name of the user account running the installation.
	Ensure that the group designated as the OINSTALL group is available as the primary group for all planned Oracle software installation owners.
Create operating system groups and users for standard or role-allocated system privileges	Create operating system groups and users depending on your security requirements, as described in this installation guide.
	Set resource limits settings and other requirements for Oracle software installation owners.
	Group and user names must use only ASCII characters.

> **✎ Note:**
>
> Do not delete an existing daemon user. If a daemon user has been deleted, then you must add it back.

Table 1-5   (Cont.) User Environment Configuration for Oracle Grid Infrastructure

Check	Task
Unset Oracle Software Environment Variables	If you have an existing Oracle software installation, and you are using the same user to install this installation, then unset the following environment variables: `$ORACLE_HOME`; `$ORA_NLS10`; `$TNS_ADMIN`.
	If you have set `$ORA_CRS_HOME` as an environment variable, then unset it before starting an installation or upgrade. Do not use `$ORA_CRS_HOME` as a user environment variable, except as directed by Oracle Support.
Configure the Oracle Software Owner Environment	Configure the environment of the oracle or grid user by performing the following tasks: • Set the default file mode creation mask (umask) to 022 in the shell startup file. • Set the DISPLAY environment variable.
Determine root privilege delegation option for installation	During installation, you are asked to run configuration scripts as the root user. You can either run these scripts manually as root when prompted, or during installation you can provide configuration information and passwords using a root privilege delegation option.

	To run root scripts automatically, select **Automatically run configuration scripts**. during installation. To use the automatic configuration option, the root user credentials for all cluster member nodes must use the same password.
	
	• **Use root user credentials**
	  Provide the superuser password for cluster member node servers.
	• **Use `sudo`**
	  `sudo` is a UNIX and Linux utility that allows members of the `sudoers` list privileges to run individual commands as root. Provide the user name and password of an operating system user that is a member of `sudoers`, and is authorized to run `sudo` on each cluster member node.
	  To enable `sudo`, have a system administrator with the appropriate privileges configure a user that is a member of the `sudoers` list, and provide the user name and password when prompted during installation.

Storage Checklist for Oracle Grid Infrastructure

Review the checklist for storage hardware and configuration requirements for Oracle Grid Infrastructure installation.

Table 1-6   Oracle Grid Infrastructure Storage Configuration Checks

Check	Task
Minimum disk space (local or shared) for Oracle Grid Infrastructure Software	• At least 12 GB of space for the Oracle Grid Infrastructure for a cluster home (Grid home). Oracle recommends that you allocate 100 GB to allow additional space for patches. At least 9 GB for Oracle Database Enterprise Edition • Allocate additional storage space in accordance with cluster configuration, as described in *Oracle Clusterware Storage Space Requirements*.
Select Oracle ASM Storage Options	During installation, based on the cluster configuration, you are asked to provide Oracle ASM storage paths for the Oracle Clusterware files. These path locations must be writable by the Oracle Grid Infrastructure installation owner (Grid user). These locations must be shared across all nodes of the cluster on Oracle ASM because the files in the Oracle ASM disk group created during installation must be available to all cluster member nodes. • For Oracle Standalone Cluster deployment, shared storage, either Oracle ASM or shared file system, is locally mounted on each of the cluster nodes. **Voting files** are files that Oracle Clusterware uses to verify cluster node membership and status. **Oracle Cluster Registry** files (OCR) contain cluster and database configuration information for Oracle Clusterware.
Select Grid Infrastructure Management Repository (GIMR) Storage Option	Depending on the type of cluster you are installing, you can choose to either host the Grid Infrastructure Management Repository (GIMR) for a cluster on the same cluster or on a remote cluster. > ✎ **Note:** > Starting with Oracle Grid Infrastructure 19c, configuring GIMR is optional for Oracle Standalone Cluster deployments. For Oracle Standalone Cluster deployment with a local GIMR, you can specify the same or separate Oracle ASM disk group for the GIMR. For Oracle Standalone Cluster deployment with a remote GIMR, you can use the remote GIMR of the Oracle Domain Services Cluster. You must specify the GIMR details when you create the credentials file before installation. For Oracle Domain Services Cluster deployment, the GIMR must be configured on a separate Oracle ASM disk group.

Installer Planning Checklist for Oracle Grid Infrastructure

Review the checklist for planning your Oracle Grid Infrastructure installation before starting Oracle Universal Installer.

Table 1-7   Oracle Universal Installer Checklist for Oracle Grid Infrastructure Installation

Check	Task
Read the Release Notes	Review release notes for your platform, which are available for your release at the following URL:
	http://www.oracle.com/technetwork/indexes/documentation/index.html
Review the Licensing Information	You are permitted to use only those components in the Oracle Database media pack for which you have purchased licenses. For more information, see:
	Oracle Database Licensing Information User Manual
Run OUI with CVU and use fixup scripts	Oracle Universal Installer is fully integrated with Cluster Verification Utility (CVU), automating many CVU prerequisite checks. Oracle Universal Installer runs all prerequisite checks and creates fixup scripts when you run the installer.
	You can also run CVU commands manually to check system readiness. For more information, see:
	Oracle Clusterware Administration and Deployment Guide
Download and run Oracle ORAchk for runtime and upgrade checks, or runtime health checks	The Oracle ORAchk utility provides system checks that can help to prevent issues after installation. These checks include kernel requirements, operating system resource allocations, and other system requirements.
	Use the Oracle ORAchk Upgrade Readiness Assessment to obtain an automated upgrade-specific system health check for upgrades. For example:
	`./orachk -u -o pre`
	The Oracle ORAchk Upgrade Readiness Assessment automates many of the manual pre- and post-upgrade checks described in Oracle upgrade documentation.
	Oracle ORAchk is supported on Windows platforms in a Cygwin environment only. For more information, see:
	https://support.oracle.com/epmos/faces/DocContentDisplay?id=2550798.1&parent=DOCUMENTATION&sourceId=USERGUIDE
Ensure `cron` jobs do not run during installation	If the installer is running when daily `cron` jobs start, then you may encounter unexplained installation problems if your cron job is performing cleanup, and temporary files are deleted before the installation is finished. Oracle recommends that you complete installation before daily `cron` jobs are run, or disable daily `cron` jobs that perform cleanup until after the installation is completed.
Obtain Your My Oracle Support account information	During installation, you require a My Oracle Support user name and password to configure security updates, download software updates, and other installation tasks. You can register for My Oracle Support at the following URL:
	https://support.oracle.com/

Table 1-7   (Cont.) Oracle Universal Installer Checklist for Oracle Grid Infrastructure Installation

Check	Task
Check running Oracle processes, and shut down processes if necessary	• On a node with a single-instance database not using Oracle ASM: You do not need to shut down the database while you install Oracle Grid Infrastructure. • On a node with a single-instance Oracle Database using Oracle ASM: Stop the existing Oracle ASM instances. The Oracle ASM instances are restarted during installation. • On an Oracle RAC Database node: This installation requires an upgrade of Oracle Clusterware, as Oracle Clusterware is required to run Oracle RAC. As part of the upgrade, you must shut down the database one node at a time as the rolling upgrade proceeds from node to node.

2

Checking and Configuring Server Hardware for Oracle Grid Infrastructure

Verify that servers where you install Oracle Grid Infrastructure meet the minimum requirements for installation.

This section provides minimum server requirements to complete installation of Oracle Grid Infrastructure. It does not provide system resource guidelines, or other tuning guidelines for particular workloads.

- Logging In to a Remote System Using X Window System
  Use this procedure to run Oracle Universal Installer (OUI) by logging on to a remote system where the runtime setting prohibits logging in directly to a graphical user interface (GUI).

- Checking Server Hardware and Memory Configuration
  Use this procedure to gather information about your server configuration.

Logging In to a Remote System Using X Window System

Use this procedure to run Oracle Universal Installer (OUI) by logging on to a remote system where the runtime setting prohibits logging in directly to a graphical user interface (GUI).

OUI is a graphical user interface (GUI) application. On servers where the runtime settings prevent GUI applications from running, you can redirect the GUI display to a client system connecting to the server.

> **Note:**
>
> If you log in as another user (for example, `oracle` or `grid`), then repeat this procedure for that user as well.

1. Start an X Window System session. If you are using an X Window System terminal emulator from a PC or similar system, then you may need to configure security settings to permit remote hosts to display X applications on your local system.

2. Enter a command using the following syntax to enable remote hosts to display X applications on the local X server:

```
# xhost + RemoteHost
```

RemoteHost is the fully qualified remote host name. For example:

```
# xhost + somehost.example.com
somehost.example.com being added to the access control list
```

3. If you are not installing the software on the local system, then use the `ssh` command to connect to the system where you want to install the software:

```
# ssh -Y RemoteHost
```

RemoteHost is the fully qualified remote host name. The `-Y` flag ("yes") enables remote X11 clients to have full access to the original X11 display. For example:

```
# ssh -Y somehost.example.com
```

4. If you are not logged in as the `root` user, and you are performing configuration steps that require `root` user privileges, then switch the user to `root`.

> **Note:**
>
> For more information about remote login using X Window System, refer to your X server documentation, or contact your X server vendor or system administrator. Depending on the X server software that you are using, you may have to complete the tasks in a different order.

Checking Server Hardware and Memory Configuration

Use this procedure to gather information about your server configuration.

1. To determine the physical RAM size on the server:

```
# /usr/contrib/bin/machinfo  | grep -i Memory
```

2. To determine the size of the configured swap space:

```
# /usr/sbin/swapinfo -a
```

If necessary, see your operating system documentation for information about how to configure additional swap space.

3. To determine the amount of space available in the `/tmp` directory:

```
# bdf /tmp
```

If the free space available in the /tmp directory is less than what is required, then complete one of the following steps:

- Delete unnecessary files from the /tmp directory to meet the disk space requirement.
- When you set the Oracle user's environment, also set the TMP and TMPDIR environment variables to the directory you want to use instead of `/tmp`.

4. To determine the amount of free disk space on the system:

```
# bdf
```

5. To determine if the system architecture can run the software:

```
# uname -m
```

Verify that the processor architecture matches the Oracle software release to install. For example, you should see the following for a x86-64 bit system:

```
ia64 (itanium)
```

If you do not see the expected output, then you cannot install the software on this system.

6. To view the shared memory kernel parameters (`shmmni`, `shmseg`, and `shmmax`) and determine if any changes are necessary for your system:

```
# kctune -v shmmni
# kctune -v shmseg
# kctune -v shmmax
```

Use the `kctune shmmni=1024` command to set the memory kernel parameter to the required value.

3

Configuring Operating Systems for Oracle Grid Infrastructure on HP-UX Itanium

Complete operating system configuration requirements and checks for HP-UX Itanium operating systems before you start installation.

- Guidelines for HP-UX Itanium Operating System Installation
  Be aware of these operating system guidelines before proceeding with an Oracle installation.

- Reviewing Operating System and Software Upgrade Best Practices
  These topics provide general planning guidelines and platform-specific information about upgrades and migration.

- Reviewing Operating System Security Common Practices
  Secure operating systems are an important basis for general system security.

- About Installation Fixup Scripts
  Oracle Universal Installer detects when the minimum requirements for an installation are not met, and creates shell scripts, called fixup scripts, to finish incomplete system configuration steps.

- About Operating System Requirements
  Depending on the products that you intend to install, verify that you have the required operating system kernel and packages installed.

- Operating System Requirements for HP-UX Itanium Systems
  The compilers and patches listed in this section are supported for this release on HP-UX Itanium systems for Oracle Database, Oracle Database Client, and Oracle Grid Infrastructure.

- Additional Drivers and Software Packages for HP-UX Itanium Systems
  You are not required to install additional drivers and packages, but you may choose to install or configure these drivers and packages.

- Checking the Software Requirements on HP-UX Itanium
  Check your software to see if they meet the minimum version, compiler, and patch requirements for installation.

- Enabling the Name Service Cache Daemon
  To allow Oracle Clusterware to better tolerate network failures with NAS devices or NFS mounts, enable the Name Service Cache Daemon (`nscd`).

- Using Automatic SSH Configuration During Installation
  To install Oracle software, configure secure shell (SSH) connectivity between all cluster member nodes.

- Setting Network Time Protocol for Cluster Time Synchronization
  Use either the Oracle Cluster Time Synchronization or network time protocol (NTP) option for time synchronization.

Guidelines for HP-UX Itanium Operating System Installation

Be aware of these operating system guidelines before proceeding with an Oracle installation.

Refer to your HP-UX documentation to obtain information about installing HP-UX Itanium on your servers.

Reviewing Operating System and Software Upgrade Best Practices

These topics provide general planning guidelines and platform-specific information about upgrades and migration.

- General Upgrade Best Practices
  Be aware of these guidelines as a best practice before you perform an upgrade.

- New Server Operating System Upgrade Option
  You can upgrade your operating system by installing a new operating system on a server, and then migrating your database either manually, or by using Export/Import method.

- Oracle ASM Upgrade Notifications
  Understand Oracle ASM upgrade options and restrictions.

General Upgrade Best Practices

Be aware of these guidelines as a best practice before you perform an upgrade.

If you have an existing Oracle Database installation, then do the following:

- Record the version numbers, patches, and other configuration information

- Review upgrade procedures for your existing installation

- Review Oracle Database upgrade documentation before proceeding with installation, to decide how you want to proceed

> ⚠️ **Caution:**
>
> Always create a backup of existing databases before starting any configuration change.

Refer to *Oracle Database Upgrade Guide* for more information about required software updates, pre-upgrade tasks, post-upgrade tasks, compatibility, and interoperability between different releases.

Related Topics

- *Oracle Database Upgrade Guide*

New Server Operating System Upgrade Option

You can upgrade your operating system by installing a new operating system on a server, and then migrating your database either manually, or by using Export/Import method.

> **Note:**
>
> Confirm that the server operating system is supported, and that kernel and package requirements for the operating system meet or exceed the minimum requirements for the Oracle Database release to which you want to migrate.

Manual, Command-Line Copy for Migrating Data and Upgrading Oracle Database

You can copy files to the new server and upgrade it manually. If you use this procedure, then you cannot use Oracle Database Upgrade Assistant. However, you can revert to your existing database if you encounter upgrade issues.

1. Copy the database files from the computer running the previous operating system to the one running the new operating system.

2. Re-create the control files on the computer running the new operating system.

3. Manually upgrade the database using command-line scripts and utilities.

> **See Also:**
>
> *Oracle Database Upgrade Guide* to review the procedure for upgrading the database manually, and to evaluate the risks and benefits of this option

Export/Import Method for Migrating Data and Upgrading Oracle Database

You can install the operating system on the new server, install the new Oracle Database release on the new server, and then use Oracle Data Pump Export and Import utilities to migrate a copy of data from your current database to a new database in the new release. Oracle Data Pump Export and Import are recommended for higher performance and to ensure support for new data types.

> **See Also:**
>
> *Oracle Database Upgrade Guide* to review the Export/Import method for migrating data and upgrading Oracle Database

Oracle ASM Upgrade Notifications

Understand Oracle ASM upgrade options and restrictions.

ORACLE®

- You can upgrade Oracle Automatic Storage Management (Oracle ASM) without shutting down an Oracle RAC database by performing a rolling upgrade either of individual nodes, or of a set of nodes in the cluster. However, if you have a single-instance database on a cluster that uses Oracle ASM, then you must shut down the single-instance database before upgrading.

- Oracle ASM is installed with Oracle Clusterware in the Oracle Grid Infrastructure home (Grid home).

- Two nodes of different releases cannot run in the cluster. When upgrading from Oracle Grid Infrastructure 12c release 2 (12.2), Oracle Grid Infrastructure 18c, or Oracle Grid Infrastructure 19c to a later release, if there is an outage during the rolling upgrade, then when you restart the upgrade, ensure that you start the earlier release of Oracle Grid Infrastructure and bring the Oracle ASM cluster back in the rolling migration mode.

Reviewing Operating System Security Common Practices

Secure operating systems are an important basis for general system security.

Ensure that your operating system deployment is in compliance with common security practices as described in your operating system vendor security guide.

About Installation Fixup Scripts

Oracle Universal Installer detects when the minimum requirements for an installation are not met, and creates shell scripts, called fixup scripts, to finish incomplete system configuration steps.

If Oracle Universal Installer detects an incomplete task, then it generates fixup scripts (`runfixup.sh`). You can run the fixup script and click **Fix and Check Again**. The fixup script modifies both persistent parameter settings and parameters in memory, so you do not have to restart the system.

The Fixup script does the following tasks:

- Sets kernel parameters, if necessary, to values required for successful installation, including:

  - Shared memory parameters.

  - Open file descriptor and UDP send/receive parameters.

- Creates and sets permissions on the Oracle Inventory (central inventory) directory.

- Creates or reconfigures primary and secondary group memberships for the installation owner, if necessary, for the Oracle Inventory directory and the operating system privileges groups.

- Sets shell limits, if necessary, to required values.

> ✎ **Note:**
>
> Using fixup scripts does not ensure that all the prerequisites for installing Oracle Database are met. You must still verify that all the preinstallation requirements are met to ensure a successful installation.

Oracle Universal Installer is fully integrated with Cluster Verification Utility (CVU) automating many prerequisite checks for your Oracle Grid Infrastructure or Oracle Real Application Clusters (Oracle RAC) installation. You can also manually perform various CVU verifications by running the `cluvfy` command.

Related Topics

- *Oracle Clusterware Administration and Deployment Guide*

About Operating System Requirements

Depending on the products that you intend to install, verify that you have the required operating system kernel and packages installed.

Requirements listed in this document are current as of the date listed on the title page.

Oracle Universal Installer performs checks on your system to verify that it meets the listed operating system package requirements. To ensure that these checks complete successfully, verify the requirements before you start OUI.

> **Note:**
>
> Oracle does not support running different operating system versions on cluster members, unless an operating system is being upgraded. You cannot run different operating system version binaries on members of the same cluster, even if each operating system is supported.

Operating System Requirements for HP-UX Itanium Systems

The compilers and patches listed in this section are supported for this release on HP-UX Itanium systems for Oracle Database, Oracle Database Client, and Oracle Grid Infrastructure.

The platform-specific hardware and software requirements included in this guide were current when this guide was published. However, because new platforms and operating system software versions might be certified after this guide is published, review the certification matrix on the My Oracle Support website for the most up-to-date list of certified hardware platforms and operating system versions:

https://support.oracle.com/

Identify the requirements for your HP-UX Itanium system, and ensure that you have supported compilers and required patches installed before starting installation.

- Supported HP-UX Release on Itanium
  Use the following information to check the minimum supported HP-UX Itanium distribution:

Supported HP-UX Release on Itanium

Use the following information to check the minimum supported HP-UX Itanium distribution:

Table 3-1   HP-UX Itanium Minimum Operating System Requirements

Item	Requirements
HP-UX 11.31 operating system	HP-UX 11.31 May 2019 B.11.31.1905.441b or later.
Patches for HP-UX 11.31	The following patches must be installed: PHCO_43503 - 11.31 diskowner(1M) cumulative patch PHKL_44248 - 11.31 SCSI cumulative I/O patch PHSS_37042 - 11.31 hppac PHSS_42686 - 11.31 assembler patch PHSS_42981 - 11.31 OS patch PHSS_43205 - 11.31 Math Library Cumulative Patch PHSS_43733 - 11.31 LIBCL patch PHSS_43740 - 11.31 Integrity Unwind Library
Oracle Clusterware	• HP Serviceguard A.11.20 • HP Serviceguard A.11.20 extension for RAC **Note:** HP Serviceguard is optional. It is required only if you want to use shared logical volumes for Oracle Clusterware or database files.

Additional Drivers and Software Packages for HP-UX Itanium Systems

You are not required to install additional drivers and packages, but you may choose to install or configure these drivers and packages.

- **Installing Oracle Messaging Gateway**
  Oracle Messaging Gateway is installed with Enterprise Edition of Oracle Database. However, you may require a CSD or Fix Packs.

- **Installation Requirements for ODBC and LDAP**
  Review these topics to install Open Database Connectivity (ODBC) and Lightweight Directory Access Protocol (LDAP).

- **Installation Requirements for Programming Environments for HP-UX Itanium Systems**
  Ensure that your system meets the requirements for the programming environment you want to configure.

- **Installation Requirements for Web Browsers**
  Web browsers are required only if you intend to use Oracle Enterprise Manager Database Express and Oracle Enterprise Manager Cloud Control. Web browsers must support JavaScript, and the HTML 4.0 and CSS 1.0 standards.

Installing Oracle Messaging Gateway

Oracle Messaging Gateway is installed with Enterprise Edition of Oracle Database. However, you may require a CSD or Fix Packs.

If you require a CSD or Fix Packs for IBM WebSphere MQ, then see the following website for more information:

https://www.ibm.com/support/

> **Note:**
>
> Oracle Messaging Gateway does not support the integration of Advanced Queuing with TIBCO Rendezvous on IBM: Linux on System z.

Related Topics

• *Oracle Database Advanced Queuing User's Guide*

Installation Requirements for ODBC and LDAP

Review these topics to install Open Database Connectivity (ODBC) and Lightweight Directory Access Protocol (LDAP).

• About ODBC Drivers and Oracle Database
  Open Database Connectivity (ODBC) is a set of database access APIs that connect to the database, prepare, and then run SQL statements on the database.

• Installing ODBC Drivers for HP-UX Itanium Systems
  If you intend to use ODBC, then download and install the most recent ODBC Driver Manager for HP-UX Itanium.

• About LDAP and Oracle Plug-ins
  Lightweight Directory Access Protocol (LDAP) is an application protocol for accessing and maintaining distributed directory information services over IP networks.

• Installing the LDAP Package
  LDAP is included in a default operating system installation.

About ODBC Drivers and Oracle Database

Open Database Connectivity (ODBC) is a set of database access APIs that connect to the database, prepare, and then run SQL statements on the database.

An application that uses an ODBC driver can access non-uniform data sources, such as spreadsheets and comma-delimited files.

Installing ODBC Drivers for HP-UX Itanium Systems

If you intend to use ODBC, then download and install the most recent ODBC Driver Manager for HP-UX Itanium.

http://www.unixodbc.org

Review the minimum supported ODBC driver releases, and install ODBC drivers of the following or later releases for HP-UX Itanium:

ODBC Driver Manager 2.3.4

To use ODBC, you must also install `gcc 4.2.3` or later.

About LDAP and Oracle Plug-ins

Lightweight Directory Access Protocol (LDAP) is an application protocol for accessing and maintaining distributed directory information services over IP networks.

You require the LDAP package if you want to use features requiring LDAP, including the Oracle Database scripts `odisrvreg` and `oidca` for Oracle Internet Directory, or `schemasync` for third-party LDAP directories.

Installing the LDAP Package

LDAP is included in a default operating system installation.

If you did not perform a default operating system installation, and you intend to use Oracle scripts requiring LDAP, then use a package management system for your distribution to install a supported LDAP package for your distribution, and install any other required packages for that LDAP package.

Installation Requirements for Programming Environments for HP-UX Itanium Systems

Ensure that your system meets the requirements for the programming environment you want to configure.

These requirements are not mandatory for the database installation.

Table 3-2   Requirements for Programming Environments for HP-UX Itanium Systems

Programming Environments	Support Requirements
Java Database Connectivity (JDBC) / Oracle Call Interface (OCI)	JDK 8 with the JNDI extension with Oracle Java Database Connectivity and Oracle Call Interface drivers.
Oracle C++ Oracle C++ Call Interface Pro*C/C++ Oracle XML Developer's Kit (XDK)	Compiler Version: HP C/aC++ B3910B A.06.29 [Oct 18 2016]
Pro*COBOL	Micro Focus Visual Cobol V2.3 Micro Focus Visual Cobol V6.0
Pro*FORTRAN	HP FORTRAN/90 - Sep 2008 - release
unzip	Install the latest version of Unzip. Unzip is required to extract the image files for Oracle Database and Oracle Grid Infrastructure installations.

Installation Requirements for Web Browsers

Web browsers are required only if you intend to use Oracle Enterprise Manager Database Express and Oracle Enterprise Manager Cloud Control. Web browsers must support JavaScript, and the HTML 4.0 and CSS 1.0 standards.

ORACLE®

https://support.oracle.com

Related Topics

• *Oracle Enterprise Manager Cloud Control Basic Installation Guide*

Checking the Software Requirements on HP-UX Itanium

Check your software to see if they meet the minimum version, compiler, and patch requirements for installation.

1. To determine the distribution and version of HP-UX installed:

    ```
    # uname -a
    ```

    ```
    HP-UX hostname B.11.31 U ia64 3504233448 unlimited-user license
    ```

    In this example, the version of HP-UX is 11.31 and the processor is Itanium.

2. To determine the compiler installed:

    ```
    # /usr/sbin/swlist -l product | grep -i compiler
    ```

3. To determine if a patch is installed:

    ```
    # /usr/sbin/swlist -l patch | grep PHSS_42686
    ```

    Alternatively, to list all installed patches:

    ```
    # /usr/sbin/swlist -l patch | more
    ```

    If a required patch is not installed, then download it from the HP Support Center website and install it:
    https://www.support.hp.com/

    If the website has a recent version of the patch, then download and install that version.

4. To verify if the system meets the minimum patch bundle requirements:

    ```
    # /usr/sbin/swlist -l bundle |grep QPK
    ```

    The QPK (Quality Pack) bundles have version numbers of the form `B.11.31.1905.441b` (for the May 2019 release).

    If a required bundle, product, or fileset is not installed, then install it. See your operating system or software documentation for information about installing products.

> **Note:**
>
> There may be more recent versions of the patches listed in the preceding paragraph that are installed on the system. If a listed patch is not installed, then determine if a more recent version is installed before installing the version listed.

5. If you require a CSD for WebSphere MQ, then refer to the following website for download and installation information:

   https://www.ibm.com/support/

Enabling the Name Service Cache Daemon

To allow Oracle Clusterware to better tolerate network failures with NAS devices or NFS mounts, enable the Name Service Cache Daemon (nscd).

To check to see if nscd is set to load when the system is restarted, enter the command chkconfig --list nscd. For example:

```
# chkconfig --list nscd
nscd                    0:off  1:off  2:off  3:on   4:off  5:off
6:off
```

nscd is turned on for run level 3, and turned off for run level 5. The nscd should be turned on for both run level 3 and run level 5.

To change the configuration to ensure that nscd is on for both run level 3 and run level 5, enter the following command as root:

```
# chkconfig --level 35 nscd on
```

To start up nscd in the current session, enter the command as root:

```
# service nscd start
```

To restart nscd with the new setting, enter the following command as root:

```
# service nscd restart

nscd

systemctl --all |grep nscd
nscd.service loaded active running Name Service Cache Daemon
```

Using Automatic SSH Configuration During Installation

To install Oracle software, configure secure shell (SSH) connectivity between all cluster member nodes.

Oracle Universal Installer (OUI) uses the `ssh` and `scp` commands during installation to run remote commands on and copy files to the other cluster nodes. You must configure SSH so that these commands do not prompt for a password.

> **Note:**
>
> Oracle configuration assistants use SSH for configuration operations from local to remote nodes. Oracle Enterprise Manager also uses SSH. RSH is no longer supported.

You can configure SSH from the OUI interface during installation for the user account running the installation. The automatic configuration creates passwordless SSH connectivity between all cluster member nodes. Oracle recommends that you use the automatic procedure if possible.

To enable the script to run, you must remove `stty` commands from the profiles of any existing Oracle software installation owners you want to use, and remove other security measures that are triggered during a login, and that generate messages to the terminal. These messages, mail checks, and other displays prevent Oracle software installation owners from using the SSH configuration script that is built into OUI. If they are not disabled, then SSH must be configured manually before an installation can be run.

In rare cases, Oracle Clusterware installation can fail during the "AttachHome" operation when the remote node closes the SSH connection. To avoid this problem, set the timeout wait to unlimited by setting the following parameter in the SSH daemon configuration file `/etc/ssh/sshd_config` on all cluster nodes:

```
LoginGraceTime 0
```

Setting Network Time Protocol for Cluster Time Synchronization

Use either the Oracle Cluster Time Synchronization or network time protocol (NTP) option for time synchronization.

Oracle Clusterware requires the same time zone environment variable setting on all cluster nodes. During installation, the installation process picks up the time zone (TZ) environment variable setting of the Grid installation owner on the node where Oracle Universal Installer (OUI) runs, and uses that time zone value on all nodes as the default TZ environment variable setting for all processes managed by Oracle Clusterware. The time zone default is used for databases, Oracle ASM, and any other managed processes. You have two options for time synchronization:

- An operating system configured network time protocol (NTP) such as `ntpd`
- Oracle Cluster Time Synchronization Service

Oracle Cluster Time Synchronization Service is designed for organizations whose cluster servers are unable to access NTP services. If you use NTP, then the Oracle Cluster Time Synchronization daemon (ctssd) starts up in observer mode. If you do not have NTP daemons, then `ctssd` starts up in active mode and synchronizes time among cluster members without contacting an external time server.

> **Note:**
>
> Before starting the installation of Oracle Grid Infrastructure, Oracle recommends that you ensure the clocks on all nodes are set to the same time.

If you have NTP daemons on your server but you cannot configure them to synchronize time with a time server, and you want to use Cluster Time Synchronization Service to provide synchronization service in the cluster, then deactivate and deinstall the NTP.

Deactivating the ntpd Service

To deactivate the NTP service, run the following command as `root`:

```
# /sbin/init.d/xntpd stop
# rm /etc/ntp.conf
```

or, `# mv /etc/ntp.conf` to `# /etc/ntp.conf.org`.

When the installer finds that the NTP protocol is not active, the Cluster Time Synchronization Service is installed in active mode and synchronizes the time across the nodes. If NTP is found configured, then the Cluster Time Synchronization Service is started in observer mode, and no active time synchronization is performed by Oracle Clusterware within the cluster.

Confirming Oracle Cluster Time Synchronization Service After Installation

To confirm that `ctssd` is active after installation, enter the following command as the Grid installation owner:

```
$ crsctl check ctss
```

4

Configuring Networks for Oracle Grid Infrastructure and Oracle RAC

Check that you have the networking hardware and internet protocol (IP) addresses required for an Oracle Grid Infrastructure for a cluster installation.

- About Oracle Grid Infrastructure Network Configuration Options
  Ensure that you have the networking hardware and internet protocol (IP) addresses required for an Oracle Grid Infrastructure for a cluster installation.

- Understanding Network Addresses
  During installation, you are asked to identify the planned use for each network interface that Oracle Universal Installer (OUI) detects on your cluster node.

- Network Interface Hardware Minimum Requirements
  Review these requirements to ensure that you have the minimum network hardware technology for Oracle Grid Infrastructure clusters.

- Private IP Interface Configuration Requirements
  Requirements for private interfaces depend on whether you are using single or multiple Interfaces.

- IPv4 and IPv6 Protocol Requirements
  Oracle Grid Infrastructure and Oracle RAC support the standard IPv6 address notations specified by RFC 2732 and global and site-local IPv6 addresses as defined by RFC 4193.

- Oracle Grid Infrastructure IP Name and Address Requirements
  Review this information for Oracle Grid Infrastructure IP Name and Address requirements.

- Broadcast Requirements for Networks Used by Oracle Grid Infrastructure
  Broadcast communications (ARP and UDP) must work properly across all the public and private interfaces configured for use by Oracle Grid Infrastructure.

- Multicast Requirements for Networks Used by Oracle Grid Infrastructure
  For each cluster member node, the Oracle mDNS daemon uses multicasting on all interfaces to communicate with other nodes in the cluster.

- Domain Delegation to Grid Naming Service
  If you are configuring Grid Naming Service (GNS) for a standard cluster, then before installing Oracle Grid Infrastructure you must configure DNS to send to GNS any name resolution requests for the subdomain served by GNS.

- Configuration Requirements for Oracle Flex Clusters
  Understand Oracle Flex Clusters and their configuration requirements.

- Grid Naming Service Cluster Configuration Example
  Review this example to understand Grid Naming Service configuration.

- Manual IP Address Configuration Example
  If you choose not to use GNS, then before installation you must configure public, virtual, and private IP addresses.

- **Network Interface Configuration Options**
  During installation, you are asked to identify the planned use for each network adapter (or network interface) that Oracle Universal Installer (OUI) detects on your cluster node.

About Oracle Grid Infrastructure Network Configuration Options

Ensure that you have the networking hardware and internet protocol (IP) addresses required for an Oracle Grid Infrastructure for a cluster installation.

Oracle Clusterware Networks

An Oracle Clusterware configuration requires at least two interfaces:

- A public network interface, on which users and application servers connect to access data on the database server.

- A private network interface for internode communication.

You can configure a network interface to use either the IPv4 protocol, or the IPv6 protocol on a given network. If you use redundant network interfaces (bonded or teamed interfaces), then be aware that Oracle does not support configuring one interface to support IPv4 addresses and the other to support IPv6 addresses. You must configure network interfaces of a redundant interface pair with the same IP protocol.

All the nodes in the cluster must use the same IP protocol configuration. Either all the nodes use only IPv4, or all the nodes use only IPv6. You cannot have some nodes in the cluster configured to support only IPv6 addresses, and other nodes in the cluster configured to support only IPv4 addresses.

The VIP agent supports the generation of IPv6 addresses using the Stateless Address Autoconfiguration Protocol (RFC 2462), and advertises these addresses with GNS. Run the `srvctl config` network command to determine if Dynamic Host Configuration Protocol (DHCP) or stateless address autoconfiguration is being used.

See the Certify page on My Oracle Support for the most up-to-date information about supported network protocols and hardware for Oracle RAC:

https://support.oracle.com

Understanding Network Addresses

During installation, you are asked to identify the planned use for each network interface that Oracle Universal Installer (OUI) detects on your cluster node.

Identify each interface as a public or private interface, or as an interface that you do not want Oracle Grid Infrastructure or Oracle ASM to use. Public and virtual internet protocol (VIP) addresses are configured on public interfaces. Private addresses are configured on private interfaces.

- **About the Public IP Address**
  The public IP address is assigned dynamically using Dynamic Host Configuration Protocol (DHCP), or defined statically in a Domain Name System (DNS) or in a hosts file.

- About the Private IP Address
  Oracle Clusterware uses interfaces marked as private for internode communication.

- About the Virtual IP Address
  The virtual IP (VIP) address is registered in the grid naming service (GNS), the DNS, or in a hosts file.

- About the Grid Naming Service (GNS) Virtual IP Address
  The GNS virtual IP address is a static IP address configured in the Domain Name System (DNS).

- About the SCAN
  Oracle Database clients connect to the database using a Single Client Access Name (SCAN).

- About Shared SCAN
  Shared SCAN provides the capability to share SCAN VIPs across multiple clusters, thus reducing the number of IP addresses that must be assigned when deploying Oracle Clusters.

About the Public IP Address

The public IP address is assigned dynamically using Dynamic Host Configuration Protocol (DHCP), or defined statically in a Domain Name System (DNS) or in a hosts file.

The public IP address uses the public interface (the interface with access available to clients). The public IP address is the primary address for a cluster member node, and should be the address that resolves to the name returned when you enter the command hostname.

If you configure IP addresses manually, then avoid changing host names after you complete the Oracle Grid Infrastructure installation, including adding or deleting domain qualifications. A node with a new host name is considered a new host, and must be added to the cluster. A node under the old name appears to be down until it is removed from the cluster.

About the Private IP Address

Oracle Clusterware uses interfaces marked as private for internode communication.

Each cluster node must have an interface that you identify during installation as a private interface. Private interfaces must have addresses configured for the interface itself, but no additional configuration is required. Oracle Clusterware uses the interfaces you identify as private for the cluster interconnect. If you identify multiple interfaces during information for the private network, then Oracle Clusterware configures them with Redundant Interconnect Usage. Any interface that you identify as private must be on a subnet that connects to every node of the cluster. Oracle Clusterware uses all the interfaces you identify for use as private interfaces.

For the private interconnects, because of Cache Fusion and other traffic between nodes, Oracle strongly recommends using a physically separate, private network. If you configure addresses using a DNS, then you should ensure that the private IP addresses are reachable only by the cluster nodes.

You can choose multiple interconnects either during installation or postinstallation using the `oifcfg setif` command.

After installation, if you modify the interconnect for Oracle Real Application Clusters (Oracle RAC) with the `CLUSTER_INTERCONNECTS` initialization parameter, then you must change the interconnect to a private IP address, on a subnet that is not used with a public IP address,

nor marked as a public subnet by `oifcfg`. Oracle does not support changing the interconnect to an interface using a subnet that you have designated as a public subnet.

You should not use a firewall on the network with the private network IP addresses, because this can block interconnect traffic.

About the Virtual IP Address

The virtual IP (VIP) address is registered in the grid naming service (GNS), the DNS, or in a hosts file.

> **Note:**
>
> Starting with Oracle Grid Infrastructure 18c, using VIP is optional for Oracle Clusterware deployments. You can specify VIPs for all or none of the cluster nodes. However, specifying VIPs for selected cluster nodes is not supported.

Select an address for your VIP that meets the following requirements:

- The IP address and host name are currently unused (it can be registered in a DNS, but should not be accessible by a ping command)
- The VIP is on the same subnet as your public interface

If you are not using Grid Naming Service (GNS), then determine a virtual host name for each node. A virtual host name is a public node name that reroutes client requests sent to the node if the node is down. Oracle Database uses VIPs for client-to-database connections, so the VIP address must be publicly accessible. Oracle recommends that you provide a name in the format *hostname*-vip. For example: `myclstr2-vip`.

About the Grid Naming Service (GNS) Virtual IP Address

The GNS virtual IP address is a static IP address configured in the Domain Name System (DNS).

The DNS delegates queries to the GNS virtual IP address, and the GNS daemon responds to incoming name resolution requests at that address. Within the subdomain, the GNS enables the cluster to map host names and IP addresses dynamically as nodes are added and removed from the cluster, without requiring additional host configuration in the DNS.

To enable GNS, you must have your network administrator provide a set of IP addresses for a subdomain assigned to the cluster (for example, `grid.example.com`), and delegate DNS requests for that subdomain to the GNS virtual IP address for the cluster, which GNS serves. DHCP provides the set of IP addresses to the cluster. DHCP must be available on the public network for the cluster.

Related Topics

- *Oracle Clusterware Administration and Deployment Guide*

About the SCAN

Oracle Database clients connect to the database using a Single Client Access Name (SCAN).

The SCAN and its associated IP addresses provide a stable name for clients to use for connections, independent of the nodes that make up the cluster. SCAN addresses, virtual IP addresses, and public IP addresses must all be on the same subnet.

The SCAN is a virtual IP name, similar to the names used for virtual IP addresses, such as `node1-vip`. However, unlike a virtual IP, the SCAN is associated with the entire cluster, rather than an individual node, and associated with multiple IP addresses, not just one address.

The SCAN resolves to multiple IP addresses reflecting multiple listeners in the cluster that handle public client connections. When a client submits a request, the SCAN listener listening on a SCAN IP address and the SCAN port is made available to a client. Because all services on the cluster are registered with the SCAN listener, the SCAN listener replies with the address of the local listener on the least-loaded node where the service is currently being offered. Finally, the client establishes connection to the service through the listener on the node where service is offered. All of these actions take place transparently to the client without any explicit configuration required in the client.

During installation, listeners are created. These SCAN listeners listen on the SCAN IP addresses. The SCAN listeners are started on nodes determined by Oracle Clusterware. Oracle Net Services routes application requests to the least-loaded instance providing the service. Because the SCAN addresses resolve to the cluster, rather than to a node address in the cluster, nodes can be added to or removed from the cluster without affecting the SCAN address configuration. The SCAN listener also supports HTTP protocol for communication with Oracle XML Database (XDB).

The SCAN should be configured so that it is resolvable either by using Grid Naming Service (GNS) within the cluster, or by using Domain Name Service (DNS) resolution. For high availability and scalability, Oracle recommends that you configure the SCAN name so that it resolves to three IP addresses. Ensure that the SCAN resolves to at least one IP address. However, configuring less than the recommended three IP addresses may impact the availability to connect to the cluster.

If you specify a GNS domain, then the SCAN name defaults to *clustername-scan.cluster_name.GNS_domain*. Otherwise, it defaults to *clustername-scan.current_domain*. For example, if you start Oracle Grid Infrastructure installation from the server `node1`, the cluster name is `mycluster`, and the GNS domain is `grid.example.com`, then the SCAN Name is `mycluster-scan.mycluster.grid.example.com`.

The database registers with the SCAN listener through the remote listener parameter in the `init.ora` file. The `REMOTE_LISTENER` parameter must be set to `SCAN:PORT`. Do not set it to a `TNSNAMES` alias with a single address for the SCAN, for example, using `HOST=` *SCAN_name*.

The SCAN is optional for most deployments. However, clients using policy-managed databases using server pools must access the database using the SCAN. This is required because policy-managed databases can run on different servers at different times, so connecting to a particular node by using the virtual IP address for a policy-managed database is not possible.

Provide SCAN addresses for client access to the cluster. These addresses must be configured as round robin addresses on the domain name service (DNS), if DNS is used. Oracle recommends that you supply three SCAN addresses.

Identify public and private interfaces. Oracle Universal Installer configures public interfaces for use by public and virtual IP addresses, and configures private IP addresses on private interfaces. The private subnet that the private interfaces use must connect all the nodes you intend to have as cluster members. The SCAN must be in the same subnet as the public interface.

Related Topics

• *Oracle Real Application Clusters Administration and Deployment Guide*

About Shared SCAN

Shared SCAN provides the capability to share SCAN VIPs across multiple clusters, thus reducing the number of IP addresses that must be assigned when deploying Oracle Clusters.

In earlier versions of the Oracle Clusterware, SCAN VIPs were configured on a per cluster basis. With shared SCAN, the same SCAN is used among multiple clusters, so that only one of these clusters runs SCAN listeners. The databases of all clusters use the SCAN VIPs of this cluster, for all their database connections. Each cluster has its own set of ports, instead of SCAN VIPs. Clusters using shared SCAN can name their database services as desired, without naming conflicts if one or more of these clusters are configured with services of the same name.

Network Interface Hardware Minimum Requirements

Review these requirements to ensure that you have the minimum network hardware technology for Oracle Grid Infrastructure clusters.

Public Network for Each Node

Public networks provide access to clients for database services. Public networks must meet these minimum requirements:

• Adapters: Each node must have at least one public network adapter or network interface cards (NIC).

  Oracle supports the use of link aggregations, bonded, trunked or teamed networks for improved bandwidth and high availability.

• Protocol: Each public interface must support TCP/IP.

Private Network for Each Node

Private networks (also called interconnects) are networks that only cluster member nodes can access. They use switches for connections. Private networks must meet these minimum requirements:

• Adapters: Each node must have at least one private network adapter or network interface cards (NIC).

  Oracle recommends that you configure interconnects using Redundant Interconnect Usage, in which multiple network adapters are configured with addresses in the link-local range to provide highly available IP (HAIP) addresses for the interconnect. You can configure Redundant Interconnect Usage either during installation, or after installation by using Oracle Interface Configuration Tool (OIFCFG), to provide improved bandwidth and high availability.

Oracle also supports the use of link aggregations, bonded, trunked or teamed networks for improved bandwidth and high availability.

• Protocol: User datagram protocol (UDP) using high-speed network adapters and switches that support TCP/IP, or Reliable Datagram Sockets (RDS) with Infiniband.

Switches: You must use switches for interconnects that support TCP/IP. Oracle recommends that you use dedicated switches. The minimum switch speed is 1 Gigabit Ethernet.

Local Area Network Technology

Oracle does not support token-rings or crossover cables for the interconnect. Oracle supports Jumbo Frames and Infiniband. When you use Infiniband on the interconnect, Oracle supports using the RDS protocol.

If you have a shared Ethernet VLAN deployment, with shared physical adapter, ensure that you apply standard Ethernet design, deployment, and monitoring best practices to protect against cluster outages and performance degradation due to common shared Ethernet switch network events.

Storage Networks

Oracle Automatic Storage Management and Oracle Real Application Clusters require network-attached storage.

Oracle Automatic Storage Management (Oracle ASM): The network interfaces used for Oracle Clusterware files are also used for Oracle ASM.

Third-party storage: Oracle recommends that you configure additional interfaces for storage.

Private IP Interface Configuration Requirements

Requirements for private interfaces depend on whether you are using single or multiple Interfaces.

Network Requirements for Single Interface Private Network Clusters

• Each node's private interface for interconnects must be on the same subnet.

• The subnet must connect to every node of the cluster.

For example, if the private interfaces have a subnet mask of 255.255.255.0, then your private network is in the range 192.168.0.0--192.168.0.255, and your private addresses must be in the range of 192.168.0.[0-255]. If the private interfaces have a subnet mask of 255.255.0.0, then your private addresses can be in the range of 192.168.[0-255].[0-255]

• Both IPv4 and IPv6 addresses are supported.

Network Requirements for Redundant Interconnect Usage Clusters

With Redundant Interconnect Usage, you can identify multiple interfaces to use for the cluster private network, without the need of using bonding or other technologies.

When you define multiple interfaces, Oracle Clusterware creates from one to four highly available IP (HAIP) addresses. Oracle RAC and Oracle Automatic Storage Management (Oracle ASM) instances use these interface addresses to ensure highly available, load-balanced interface communication between nodes. The installer enables Redundant Interconnect Usage to provide a high availability private network. By default, Oracle Grid Infrastructure software uses all of the HAIP addresses for private network communication,

providing load-balancing across the set of interfaces you identify for the private network. If a private interconnect interface fails or become non-communicative, then Oracle Clusterware transparently moves the corresponding HAIP address to one of the remaining functional interfaces.

- Each private interface should be on a different subnet.

- Each cluster member node must have an interface on each private interconnect subnet, and these subnets must connect to every node of the cluster.

  For example, you can have private networks on subnets 192.168.0 and 10.0.0, but each cluster member node must have an interface connected to the 192.168.0 and 10.0.0 subnets.

- Endpoints of all designated interconnect interfaces must be completely reachable on the network. There should be no node that is not connected to every private network interface.

  You can test if an interconnect interface is reachable using `ping`.

- You can use IPv4 and IPv6 addresses for the interfaces with Oracle Clusterware Redundant interconnects.

> **Note:**
>
> During installation, you can define up to four interfaces for the private network. The number of HAIP addresses created during installation is based on both physical and logical interfaces configured for the network adapter. After installation, you can define additional interfaces. If you define more than four interfaces as private network interfaces, then be aware that Oracle Clusterware activates only four of the interfaces at a time. However, if one of the four active interfaces fails, then Oracle Clusterware transitions the HAIP addresses configured to the failed interface to one of the reserve interfaces in the defined set of private interfaces.

Related Topics

- *Oracle Clusterware Administration and Deployment Guide*

IPv4 and IPv6 Protocol Requirements

Oracle Grid Infrastructure and Oracle RAC support the standard IPv6 address notations specified by RFC 2732 and global and site-local IPv6 addresses as defined by RFC 4193.

Configuring Public VIPs

Cluster member node interfaces can be configured to use IPv4, IPv6, or both types of Internet protocol addresses. During installation, you can configure VIPs for a given public network as IPv4 or IPv6 types of addresses. You can configure an IPv6 cluster by selecting VIP and SCAN names that resolve to addresses in an IPv6 subnet for the cluster, and selecting that subnet as public during installation. After installation, you can also configure cluster member nodes with a mixture of IPv4 and IPv6 addresses.

If you install using static virtual IP (VIP) addresses in an IPv4 cluster, then the VIP names you supply during installation should resolve only to IPv4 addresses. If you

install using static IPv6 addresses, then the VIP names you supply during installation should resolve only to IPv6 addresses.

During installation, you cannot configure the cluster with VIP and SCAN names that resolve to both IPv4 and IPv6 addresses. You cannot configure VIPs and SCANS on some cluster member nodes to resolve to IPv4 addresses, and VIPs and SCANs on other cluster member nodes to resolve to IPv6 addresses. Oracle does not support this configuration.

Configuring Private IP Interfaces (Interconnects)

You can configure the private network either as an IPv4 network or IPv6 network.

Redundant Network Interfaces

If you configure redundant network interfaces for a public or VIP node name, then configure both interfaces of a redundant pair to the same address protocol. Also ensure that private IP interfaces use the same IP protocol. Oracle does not support names using redundant interface configurations with mixed IP protocols. You must configure both network interfaces of a redundant pair with the same IP protocol.

GNS or Multi-Cluster Addresses

Oracle Grid Infrastructure supports IPv4 DHCP addresses, and IPv6 addresses configured with the Stateless Address Autoconfiguration protocol, as described in RFC 2462.

> ✎ **Note:**
>
> Link-local and site-local IPv6 addresses as defined in RFC 1884 are not supported.

Oracle Grid Infrastructure IP Name and Address Requirements

Review this information for Oracle Grid Infrastructure IP Name and Address requirements.

For small clusters, you can use a static configuration of IP addresses. For large clusters, manually maintaining the large number of required IP addresses becomes too cumbersome. Use Oracle Grid Naming Service with large clusters to ease network administration costs.

- About Oracle Grid Infrastructure Name Resolution Options
  Before starting the installation, you must have at least two interfaces configured on each node: One for the private IP address, and one for the public IP address.

- Cluster Name and SCAN Requirements
  Review this information before you select the cluster name and SCAN.

- IP Name and Address Requirements For Grid Naming Service (GNS)
  Review this information for IP name and address requirements for Grid Naming Service (GNS).

- IP Name and Address Requirements For Multi-Cluster GNS
  Multi-cluster GNS differs from standard GNS in that Multi-cluster GNS provides a single networking service across a set of clusters, rather than a networking service for a single cluster.

- IP Name and Address Requirements for Manual Configuration of Cluster
  For Oracle Flex Clusters, configure static cluster node names and addresses if you do not enable GNS.

- Confirming the DNS Configuration for SCAN
  Use the `nslookup` command to confirm that the DNS is correctly associating the SCAN with the addresses.

About Oracle Grid Infrastructure Name Resolution Options

Before starting the installation, you must have at least two interfaces configured on each node: One for the private IP address, and one for the public IP address.

During installation, you are asked to identify the planned use for each network interface that Oracle Universal Installer (OUI) detects on your cluster node. Identify each interface as a public or private interface, or as an interface that you do not want Oracle Grid Infrastructure or Oracle ASM to use. Public and virtual internet protocol (VIP) addresses are configured on public interfaces. Private addresses are configured on private interfaces.

> ✎ **Note:**
>
> All name servers in the DNS configuration for a cluster must resolve to all host names used in the cluster, such as cluster node host name, VIP host name, and SCAN host name.

Configure IP addresses with one of the following options:

Dynamic IP address assignment using Multi-cluster or standard Oracle Grid Naming Service (GNS)

If you select this option, then network administrators delegate a subdomain that GNS resolves (standard or multicluster). GNS requirements are different depending on how you configure zone delegation. If you configure GNS with zone delegation, then GNS is delegated a domain for which it resolves service requests. If you configure GNS without zone delegation, then GNS has a virtual IP address that is resolved by a DNS.

For GNS with zone delegation:

- For IPv4, a DHCP service running on the public network the cluster uses
- For IPv6, an autoconfiguration service running on the public network the cluster uses
- Enough addresses on the DHCP server to provide one IP address for each node, and three IP addresses for the cluster used by the Single Client Access Name (SCAN) for the cluster

Use an existing GNS configuration

Multiple clusters can use a single GNS instance. To use GNS for multiple clusters, the DNS administrator must have delegated a zone for use by GNS. Also, there must be an instance of GNS started somewhere on the network, and the GNS instance must be accessible (not blocked by a firewall). All of the node names registered with the GNS instance must be unique.

Static IP address assignment using DNS or host file resolution

If you select this option, then network administrators assign a fixed IP address for each physical host name in the cluster, and for IPs for VIPs managed by Oracle

Clusterware. In addition, either domain name server (DNS) based static name resolution is used for each node, or host files for both the clusters and clients have to be updated, resulting in limited SCAN functionality. Selecting this option requires that you request network administration updates when you modify the cluster.

For GNS without zone delegation, configure a GNS virtual IP address (VIP) for the cluster. To enable Oracle Flex Cluster, you must at least configure a GNS virtual IP address.

Cluster Name and SCAN Requirements

Review this information before you select the cluster name and SCAN.

Cluster Name and SCAN Requirements

Cluster Name must meet the following requirements:

- The cluster name is case-insensitive, must be unique across your enterprise, must be at least one character long and no more than 15 characters in length, must be alphanumeric and may contain hyphens (-). Underscore characters (_) are not allowed.

- The SCAN and cluster name are entered in separate fields during installation, so cluster name requirements do not apply to the name used for the SCAN, and the SCAN can be longer than 15 characters. If you enter a domain with the SCAN name, and you want to use GNS with zone delegation, then the domain must be the GNS domain.

> **Note:**
>
> Select your cluster name carefully. After installation, you can only change the cluster name by reinstalling Oracle Grid Infrastructure.

IP Name and Address Requirements For Grid Naming Service (GNS)

Review this information for IP name and address requirements for Grid Naming Service (GNS).

IP Name and Address Requirements For Grid Naming Service (GNS)

If you enable Grid Naming Service (GNS), then name resolution requests to the cluster are delegated to the GNS, which is listening on the GNS virtual IP address. The domain name server (DNS) must be configured to delegate resolution requests for cluster names (any names in the subdomain delegated to the cluster) to the GNS. When a request comes to the domain, GNS processes the requests and responds with the appropriate addresses for the name requested. To use GNS, you must specify a static IP address for the GNS VIP address.

IP Name and Address Requirements For Multi-Cluster GNS

Multi-cluster GNS differs from standard GNS in that Multi-cluster GNS provides a single networking service across a set of clusters, rather than a networking service for a single cluster.

- About Multi-Cluster GNS Networks
  The general requirements for multi-cluster GNS are similar to those for standard GNS. Multi-cluster GNS differs from standard GNS in that multi-cluster GNS provides a single

networking service across a set of clusters, rather than a networking service for a single cluster.

- Requirements for Configuring GNS Server Clusters
  Review these requirements to configure GNS server clusters.

- Requirements for Configuring GNS Client Clusters
  Review these requirements to configure GNS client clusters.

- Creating and Using a GNS Client Data File
  Generate a GNS client data file and copy the file to the GNS client cluster member node on which you are running the Oracle Grid Infrastructure installation.

About Multi-Cluster GNS Networks

The general requirements for multi-cluster GNS are similar to those for standard GNS. Multi-cluster GNS differs from standard GNS in that multi-cluster GNS provides a single networking service across a set of clusters, rather than a networking service for a single cluster.

Requirements for Multi-Cluster GNS Networks

To provide networking service, multi-cluster Grid Naming Service (GNS) is configured using DHCP addresses, and name advertisement and resolution is carried out with the following components:

- The GNS server cluster performs address resolution for GNS client clusters. A GNS server cluster is the cluster where multi-cluster GNS runs, and where name resolution takes place for the subdomain delegated to the set of clusters.

- GNS client clusters receive address resolution from the GNS server cluster. A GNS client cluster is a cluster that advertises its cluster member node names using the GNS server cluster.

- If you choose to use GNS, then the GNS configured at the time of installation is the primary. A secondary GNS for high availability can be configured at a later time.

Requirements for Configuring GNS Server Clusters

Review these requirements to configure GNS server clusters.

To configure a GNS server cluster, check to ensure all of the following requirements are completed:

- Your network administrators must have delegated a subdomain to GNS for resolution.

- Before installation, create a static IP address for the GNS VIP address, and provide a subdomain that your DNS servers delegate to that static GNS IP address for resolution.

Requirements for Configuring GNS Client Clusters

Review these requirements to configure GNS client clusters.

To configure a GNS client cluster, check to ensure all of the following requirements are completed:

- A GNS server instance must be running on your network, and it must be accessible (for example, not blocked by a firewall).

- All of the node names in the GNS domain must be unique; address ranges and cluster names must be unique for both GNS server and GNS client clusters.

- You must have a GNS client data file that you generated on the GNS server cluster, so that the GNS client cluster has the information needed to delegate its name resolution to the GNS server cluster, and you must have copied that file to the GNS client cluster member node on which you are running the Oracle Grid Infrastructure installation.

Creating and Using a GNS Client Data File

Generate a GNS client data file and copy the file to the GNS client cluster member node on which you are running the Oracle Grid Infrastructure installation.

On a GNS server cluster member, run the following command, where `path_to_file` is the name and path location of the GNS client data file you create:

```
srvctl export gns -clientdata path_to_file -role client
```

For example:

```
$ srvctl export gns -clientdata /home/grid/gns_client_data -role client
```

Copy the GNS Client data file to a secure path on the GNS Client node where you run the GNS Client cluster installation. The Oracle installation user must have permissions to access that file. Oracle recommends that no other user is granted permissions to access the GNS Client data file. During installation, you are prompted to provide a path to that file.

```
srvctl add gns -clientdata path_to_file
```

For example:

```
$ srvctl add gns -clientdata /home/grid/gns_client_data
```

Related Topics

- *Oracle Clusterware Administration and Deployment Guide*

IP Name and Address Requirements for Manual Configuration of Cluster

For Oracle Flex Clusters, configure static cluster node names and addresses if you do not enable GNS.

IP Address Requirements for Static Clusters

Public and virtual IP names must conform with the RFC 952 standard, which allows alphanumeric characters and hyphens ("-"), but does not allow underscores ("_").

Oracle Clusterware manages private IP addresses in the private subnet on interfaces you identify as private during the installation interview.

Public IP Address Requirements

The cluster must have a public IP address for each node, with the following characteristics:

- Static IP address

- Configured before installation for each node, and resolvable to that node before installation

- On the same subnet as all other public IP addresses, VIP addresses, and SCAN addresses in the cluster

Virtual IP Address Requirements

The cluster must have a virtual IP address for each node, with the following characteristics:

- Static IP address

- Configured before installation for each node, but not currently in use

- On the same subnet as all other public IP addresses, VIP addresses, and SCAN addresses in the cluster

Single Client Access Name Requirements

The cluster must have a Single Client Access Name (SCAN) for the cluster, with the following characteristics:

- Three static IP addresses configured on the domain name server (DNS) before installation so that the three IP addresses are associated with the name provided as the SCAN, and all three addresses are returned in random order by the DNS to the requestor

- Configured before installation in the DNS to resolve to addresses that are not currently in use

- Given addresses on the same subnet as all other public IP addresses, VIP addresses, and SCAN addresses in the cluster

- Given a name conforms with the RFC 952 standard, which allows alphanumeric characters and hyphens ("-"), but does not allow underscores ("_")

Private IP Address Requirements

The cluster must have a private IP address for each node, with the following characteristics:

- Static IP address

- Configured before installation, but on a separate, private network, with its own subnet, that is not resolvable except by other cluster member nodes

The SCAN is a name used to provide service access for clients to the cluster. Because the SCAN is associated with the cluster as a whole, rather than to a particular node, the SCAN makes it possible to add or remove nodes from the cluster without needing to reconfigure clients. It also adds location independence for the databases, so that client configuration does not have to depend on which nodes are running a particular database. Clients can continue to access the cluster in the same way as with previous releases, but Oracle recommends that clients accessing the cluster use the SCAN.

> **Note:**
>
> The SCAN and cluster name are entered in separate fields during installation, so cluster name requirements do not apply to the SCAN name.
>
> Oracle strongly recommends that you do not configure SCAN VIP addresses in the hosts file. Use DNS resolution for SCAN VIPs. If you use the hosts file to resolve SCANs, then the SCAN can resolve to one IP address only.
>
> Configuring SCANs in a DNS or a hosts file is the only supported configuration. Configuring SCANs in a Network Information Service (NIS) is not supported.

Confirming the DNS Configuration for SCAN

Use the `nslookup` command to confirm that the DNS is correctly associating the SCAN with the addresses.

> **Note:**
>
> All name servers in the DNS configuration for a cluster must resolve to all host names used in the cluster, such as cluster node host name, VIP host name, and SCAN host name.

The following example shows how to use the `nslookup` command to confirm that the DNS is correctly associating the SCAN with the addresses:

```
root@node1]$ nslookup mycluster-scan
Server:         dns.example.com
Address:        192.0.2.001

Name:   mycluster-scan.example.com
Address: 192.0.2.201
Name:   mycluster-scan.example.com
Address: 192.0.2.202
Name:   mycluster-scan.example.com
Address: 192.0.2.203
```

After installation, when a client sends a request to the cluster, the Oracle Clusterware SCAN listeners redirect client requests to servers in the cluster.

Oracle strongly recommends that you do not configure SCAN VIP addresses in the hosts file. Use DNS resolution for SCAN VIPs. If you use the hosts file to resolve SCANs, then the SCAN can resolve to one IP address only.

Configuring SCANs in a DNS or a hosts file is the only supported configuration. Configuring SCANs in a Network Information Service (NIS) is not supported.

Broadcast Requirements for Networks Used by Oracle Grid Infrastructure

Broadcast communications (ARP and UDP) must work properly across all the public and private interfaces configured for use by Oracle Grid Infrastructure.

The broadcast must work across any configured VLANs as used by the public or private interfaces.

When configuring public and private network interfaces for Oracle RAC, you must enable Address Resolution Protocol (ARP). Highly Available IP (HAIP) addresses do not require ARP on the public network, but for VIP failover, you need to enable ARP. Do not configure `NOARP`.

Multicast Requirements for Networks Used by Oracle Grid Infrastructure

For each cluster member node, the Oracle mDNS daemon uses multicasting on all interfaces to communicate with other nodes in the cluster.

Multicast Requirements for Networks Used by Oracle Grid Infrastructure

Multicasting is required on the private interconnect. For this reason, at a minimum, you must enable multicasting for the cluster:

- Across the broadcast domain as defined for the private interconnect
- On the IP address subnet ranges 224.0.0.0/24 and optionally 230.0.1.0/24

You do not need to enable multicast communications across routers.

Domain Delegation to Grid Naming Service

If you are configuring Grid Naming Service (GNS) for a standard cluster, then before installing Oracle Grid Infrastructure you must configure DNS to send to GNS any name resolution requests for the subdomain served by GNS.

The subdomain that GNS serves represents the cluster member nodes.

- Choosing a Subdomain Name for Use with Grid Naming Service
  To implement GNS, your network administrator must configure the DNS to set up a domain for the cluster, and delegate resolution of that domain to the GNS VIP.

- Configuring DNS for Cluster Domain Delegation to Grid Naming Service
  If you plan to use Grid Naming Service (GNS) with a delegated domain, then before Oracle Grid Infrastructure installation, configure your domain name server (DNS) to send to GNS name resolution requests for the subdomain GNS serves, which are the cluster member nodes.

Choosing a Subdomain Name for Use with Grid Naming Service

To implement GNS, your network administrator must configure the DNS to set up a domain for the cluster, and delegate resolution of that domain to the GNS VIP.

Requirements for Choosing a Subdomain Name for Use with GNS

You can use a separate domain, or you can create a subdomain of an existing domain for the cluster. The subdomain name can be any supported DNS name such as `sales-cluster.rac.com`.

Oracle recommends that the subdomain name is distinct from your corporate domain. For example, if your corporate domain is `mycorp.example.com`, the subdomain for GNS might be `rac-gns.mycorp.example.com`.

If the subdomain is not distinct, then it should be for the exclusive use of GNS. For example, if you delegate the subdomain `mydomain.example.com` to GNS, then there should be no other domains that share it such as `lab1.mydomain.example.com`.

Configuring DNS for Cluster Domain Delegation to Grid Naming Service

If you plan to use Grid Naming Service (GNS) with a delegated domain, then before Oracle Grid Infrastructure installation, configure your domain name server (DNS) to send to GNS name resolution requests for the subdomain GNS serves, which are the cluster member nodes.

GNS domain delegation is mandatory with dynamic public networks (DHCP, autoconfiguration). GNS domain delegation is not required with static public networks (static addresses, manual configuration).

The following is an overview of the steps to be performed for domain delegation. Your actual procedure may be different from this example.

Configure the DNS to send GNS name resolution requests using delegation:

1. In the DNS, create an entry for the GNS virtual IP address, where the address uses the form `gns-server.clustername.domainname`. For example, where the cluster name is `mycluster`, and the domain name is example.com, and the IP address is `192.0.2.1`, create an entry similar to the following:

   ```
   mycluster-gns-vip.example.com  A  192.0.2.1
   ```

   The address you provide must be routable.

2. Set up forwarding of the GNS subdomain to the GNS virtual IP address, so that GNS resolves addresses to the GNS subdomain. To do this, create a `BIND` configuration entry similar to the following for the delegated domain, where `cluster01.example.com` is the subdomain you want to delegate:

   ```
   cluster01.example.com  NS  mycluster-gns-vip.example.com
   ```

3. When using GNS, you must configure `resolve.conf` on the nodes in the cluster (or the file on your system that provides resolution information) to contain name server entries that are resolvable to corporate DNS servers. The total timeout period configured—a combination of options attempts (retries) and options timeout (exponential backoff)—

should be less than 30 seconds. For example, where `xxx.xxx.xxx.42` and `xxx.xxx.xxx.15` are valid name server addresses in your network, provide an entry similar to the following in `/etc/resolv.conf`:

```
options attempts: 2
options timeout: 1

search cluster01.example.com example.com
nameserver xxx.xxx.xxx.42
nameserver xxx.xxx.xxx.15
```

`/etc/nsswitch.conf` controls name service lookup order. In some system configurations, the Network Information System (NIS) can cause problems with SCAN address resolution. Oracle recommends that you place the NIS entry at the end of the search list. For example:

```
/etc/nsswitch.conf
     hosts:    files   dns   nis
```

Be aware that use of NIS is a frequent source of problems when doing cable pull tests, as host name and user name resolution can fail.

Configuration Requirements for Oracle Flex Clusters

Understand Oracle Flex Clusters and their configuration requirements.

- Understanding Oracle Flex Clusters
  Starting with Oracle Grid Infrastructure 12c Release 2 (12.2), Oracle Grid Infrastructure cluster configurations are Oracle Flex Clusters deployments.

- About Oracle Flex ASM Clusters Networks
  Oracle ASM is configured within Oracle Grid Infrastructure to provide storage services as part of an Oracle Flex Cluster installation.

- General Requirements for Oracle Flex Cluster Configuration
  Review this information about network requirements for Oracle Flex Cluster configuration.

- Oracle Flex Cluster DHCP-Assigned Virtual IP (VIP) Addresses
  Configure cluster node VIP names for all the cluster nodes.

- Oracle Flex Cluster Manually-Assigned Addresses
  Review this information to manually assign cluster node VIP names for the cluster nodes.

Understanding Oracle Flex Clusters

Starting with Oracle Grid Infrastructure 12c Release 2 (12.2), Oracle Grid Infrastructure cluster configurations are Oracle Flex Clusters deployments.

Oracle Grid Infrastructure installed in an Oracle Flex Cluster configuration is a scalable, dynamic, robust network of nodes. Oracle Flex Clusters provide a platform for Oracle Real Application Clusters databases with large numbers of nodes, to support massive parallel query operations. Oracle Flex Clusters also provide a platform for other service deployments that require coordination and automation for high availability.

All nodes in an Oracle Flex Cluster belong to a single Oracle Grid Infrastructure cluster. This architecture centralizes policy decisions for deployment of resources based on application needs, to account for various service levels, loads, failure responses, and recovery. Nodes in Oracle Flex Clusters are tightly connected, and have direct access to shared storage

About Oracle Flex ASM Clusters Networks

Oracle ASM is configured within Oracle Grid Infrastructure to provide storage services as part of an Oracle Flex Cluster installation.

Oracle Flex ASM enables an Oracle ASM instance to run on a separate physical server from the database servers. Many Oracle ASM instances can be clustered to support numerous database clients. Each Oracle Flex ASM cluster has its own name that is globally unique within the enterprise.

You can consolidate all the storage requirements into a single set of disk groups. All these disk groups are managed by a small set of Oracle ASM instances running in a single Oracle Flex Cluster.

Every Oracle Flex ASM cluster has one or more cluster nodes on which Oracle ASM instances are running.

Oracle Flex ASM can use either the same private networks as Oracle Clusterware, or use its own dedicated private networks. Each network can be classified `PUBLIC`, `ASM & PRIVATE`, `PRIVATE`, or `ASM`.

The Oracle ASM network can be configured during installation, or configured or modified after installation.

About Oracle Flex ASM Cluster Configuration on Cluster Nodes

Oracle Flex ASM cluster nodes can be configured with the following characteristics:

- Are similar to prior release Oracle Grid Infrastructure cluster member nodes.

- Have direct connections to the Oracle ASM disks.

- Run a Direct ASM client process.

- Run an Oracle ASM Filter Driver, part of whose function is to provide cluster fencing security for the Oracle Flex ASM cluster.

- Respond to service requests delegated to them through the global Oracle ASM listener configured for the Oracle Flex ASM cluster, which designates three of the Oracle Flex ASM cluster member node listeners as remote listeners for the Oracle Flex ASM cluster.

- Can provide database clients that are running on cluster nodes of the Oracle ASM cluster remote access to Oracle ASM for metadata, and allow database clients to perform block I/O operations directly to Oracle ASM disks. The hosts running the Oracle ASM server and the remote database client must both be cluster nodes.

General Requirements for Oracle Flex Cluster Configuration

Review this information about network requirements for Oracle Flex Cluster configuration.

Network Requirements for Oracle Flex Cluster Configuration

- You must use Grid Naming Service (GNS) with an Oracle Flex Cluster deployment.

- You must configure the GNS VIP as a static IP address.

- On Multi-cluster configurations, you must identify the GNS client data file location. The GNS client data files are copied over from the GNS server before you start configuring a GNS client cluster.

- All public network addresses, whether assigned manually or automatically, must be in the same subnet range.

- All Oracle Flex Cluster addresses must be either static IP addresses, DHCP addresses assigned through DHCP (IPv4) or autoconfiguration addresses assigned through an autoconfiguration service (IPv6), registered in the cluster through GNS.

Oracle Flex Cluster DHCP-Assigned Virtual IP (VIP) Addresses

Configure cluster node VIP names for all the cluster nodes.

Requirements for DHCP-Assigned VIP Addresses

If you want to configure DHCP-assigned VIPs, then during installation, configure cluster node VIP names as follows:

- Automatically Assigned Names: Select the **Configure nodes Virtual IPs assigned by the Dynamic Networks** option to allow the installer to assign names to VIP addresses generated through DHCP automatically. Addresses are assigned through DHCP, and resolved by GNS. Oracle Clusterware sends DHCP requests with `client ID nodename-vip` and without a MAC address. You can verify the availability of DHCP addresses using the `cluvfy comp dhcp` command.

Oracle Flex Cluster Manually-Assigned Addresses

Review this information to manually assign cluster node VIP names for the cluster nodes.

Requirements for Manually-Assigned Addresses

If you choose to configure manually-assigned VIPs, then during installation, you must configure cluster node VIP names for all cluster nodes using one of the following options:

- Manual Names: Enter the host name and virtual IP name for each node manually. The names you provide must resolve to addresses configured on the DNS. Names must conform with the RFC 952 standard, which allows alphanumeric characters and hyphens ("-"), but does not allow underscores ("_").

- Automatically Assigned Names: Enter string variables for values corresponding to host names that you have configured on the DNS. String variables allow you to assign a large number of names rapidly during installation. Configure addresses on the DNS with the following characteristics:

  - Hostname prefix: a prefix string used in each address configured on the DNS for use by cluster member nodes. For example: `mycloud`.

  - Range: A range of numbers to be assigned to the cluster member nodes, consisting of a starting node number and an ending node number, designating the end of the range. For example: `001` and `999`.

  - Node name suffix: A suffix added after the end of a range number to a public node name. For example: `nd`.

–   VIP name suffix: A suffix added after the end of a virtual IP node name. For example: `-vip`.

Syntax

You can create manual addresses using alphanumeric strings.

Example 4-1   Examples of Manually-Assigned Addresses

`mycloud001nd; mycloud046nd; mycloud046-vip; mycloud348nd; mycloud784-vip`

Grid Naming Service Cluster Configuration Example

Review this example to understand Grid Naming Service configuration.

To use GNS, you must specify a static IP address for the GNS VIP address, and you must have a subdomain configured on your DNS to delegate resolution for that subdomain to the static GNS IP address.

As nodes are added to the cluster, your organization's DHCP server can provide addresses for these nodes dynamically. These addresses are then registered automatically in GNS, and GNS provides resolution within the subdomain to cluster node addresses registered with GNS.

Because allocation and configuration of addresses is performed automatically with GNS, no further configuration is required. Oracle Clusterware provides dynamic network configuration as nodes are added to or removed from the cluster. The following example is provided only for information.

With IPv6 networks, the IPv6 auto configuration feature assigns IP addresses and no DHCP server is required.

With a two node cluster where you have defined the GNS VIP, after installation you might have a configuration similar to the following for a two-node cluster, where the cluster name is `mycluster`, the GNS parent domain is `gns.example.com`, the subdomain is `cluster01.example.com`, the `192.0.2` portion of the IP addresses represents the cluster public IP address subdomain, and `192.168` represents the private IP address subdomain:

Table 4-1   Grid Naming Service Cluster Configuration Example

Identity	Home Node	Host Node	Given Name	Type	Address	Address Assigned By	Resolved By
GNS VIP	None	Selected by Oracle Clusterware	mycluster-gns-vip.example.com	virtual	192.0.2.1	Fixed by net administrator	DNS
Node 1 Public	Node 1	node1	node1	public	192.0.2.101	Fixed	GNS
Node 1 VIP	Node 1	Selected by Oracle Clusterware	node1-vip	virtual	192.0.2.104	DHCP	GNS
Node 1 Private	Node 1	node1	node1-priv	private	192.168.0.1	Fixed or DHCP	GNS

Table 4-1   (Cont.) Grid Naming Service Cluster Configuration Example

Identity	Home Node	Host Node	Given Name	Type	Address	Address Assigned By	Resolved By
Node 2 Public	Node 2	node2	node2	public	192.0.2.102	Fixed	GNS
Node 2 VIP	Node 2	Selected by Oracle Clusterware	node2-vip	virtual	192.0.2.105	DHCP	GNS
Node 2 Private	Node 2	node2	node2-priv	private	192.168.0.2	Fixed or DHCP	GNS
SCAN VIP 1	none	Selected by Oracle Clusterware	mycluster-scan.mycluster.cluster01.example.com	virtual	192.0.2.201	DHCP	GNS
SCAN VIP 2	none	Selected by Oracle Clusterware	mycluster-scan.mycluster.cluster01.example.com	virtual	192.0.2.202	DHCP	GNS
SCAN VIP 3	none	Selected by Oracle Clusterware	mycluster-scan.mycluster.cluster01.example.com	virtual	192.0.2.203	DHCP	GNS

Manual IP Address Configuration Example

If you choose not to use GNS, then before installation you must configure public, virtual, and private IP addresses.

Check that the default gateway can be accessed by a ping command. To find the default gateway, use the `route` command, as described in your operating system's help utility.

For example, with a two-node cluster where each node has one public and one private interface, and you have defined a SCAN domain address to resolve on your DNS to one of three IP addresses, you might have the configuration shown in the following table for your network interfaces:

Table 4-2   Manual Network Configuration Example

Identity	Home Node	Host Node	Given Name	Type	Address	Address Assigned By	Resolved By
Node 1 Public	Node 1	node1	node1	public	192.0.2.101	Fixed	DNS

Table 4-2   (Cont.) Manual Network Configuration Example

Identity	Home Node	Host Node	Given Name	Type	Address	Address Assigned By	Resolved By
Node 1 VIP	Node 1	Selected by Oracle Clusterware	node1-vip	virtual	192.0.2.104	Fixed	DNS and hosts file
Node 1 Private	Node 1	node1	node1-priv	private	192.168.0.1	Fixed	DNS and hosts file, or none
Node 2 Public	Node 2	node2	node2	public	192.0.2.102	Fixed	DNS
Node 2 VIP	Node 2	Selected by Oracle Clusterware	node2-vip	virtual	192.0.2.105	Fixed	DNS and hosts file
Node 2 Private	Node 2	node2	node2-priv	private	192.168.0.2	Fixed	DNS and hosts file, or none
SCAN VIP 1	none	Selected by Oracle Clusterware	mycluster-scan	virtual	192.0.2.201	Fixed	DNS
SCAN VIP 2	none	Selected by Oracle Clusterware	mycluster-scan	virtual	192.0.2.202	Fixed	DNS
SCAN VIP 3	none	Selected by Oracle Clusterware	mycluster-scan	virtual	192.0.2.203	Fixed	DNS

You do not need to provide a private name for the interconnect. If you want name resolution for the interconnect, then you can configure private IP names in the hosts file or the DNS. However, Oracle Clusterware assigns interconnect addresses on the interface defined during installation as the private interface (eth1, for example), and to the subnet used for the private subnet.

The addresses to which the SCAN resolves are assigned by Oracle Clusterware, so they are not fixed to a particular node. To enable VIP failover, the configuration shown in the preceding table defines the SCAN addresses and the public and VIP addresses of both nodes on the same subnet, 192.0.2.0/24.

> ✏ **Note:**
>
> All host names must conform to the RFC–952 standard, which permits alphanumeric characters, but does not allow underscores ("_").

Network Interface Configuration Options

During installation, you are asked to identify the planned use for each network adapter (or network interface) that Oracle Universal Installer (OUI) detects on your cluster node.

Each NIC can be configured to perform only one of the following roles:

- Public

- Private

- ASM

- ASM & Private

- Do Not Use

Network Interface Configuration Options

You must use the same private adapters for both Oracle Clusterware and Oracle RAC. The precise configuration you choose for your network depends on the size and use of the cluster you want to configure, and the level of availability you require. Network interfaces must be at least 1 GbE, with 10 GbE recommended. Alternatively, use InfiniBand for the interconnect.

If certified Network-attached Storage (NAS) is used for Oracle RAC and this storage is connected through Ethernet-based networks, then you must have a third network interface for NAS I/O. Failing to provide three separate interfaces in this case can cause performance and stability problems under load.

Redundant interconnect usage cannot protect network adapters used for public communication. If you require high availability or load balancing for public adapters, then use a third party solution. Typically, bonding, trunking or similar technologies can be used for this purpose.

You can enable redundant interconnect usage for the private network by selecting multiple network adapters to use as private adapters. Redundant interconnect usage creates a redundant interconnect when you identify more than one network adapter as private.

5

Configuring Users, Groups and Environments for Oracle Grid Infrastructure and Oracle Database

Before installation, create operating system groups and users, and configure user environments.

- Creating Groups, Users and Paths for Oracle Grid Infrastructure
  Log in as `root`, and use the following instructions to locate or create the Oracle Inventory group, and create a Oracle Grid Infrastructure software owner, and Oracle home directories.

- Oracle Installations with Standard and Job Role Separation Groups and Users
  A job role separation configuration of Oracle Database and Oracle ASM is a configuration with groups and users to provide separate groups for operating system authentication.

- Creating Operating System Privileges Groups
  The following sections describe how to create operating system groups for Oracle Grid Infrastructure and Oracle Database:

- Creating Operating System Oracle Installation User Accounts
  Before starting installation, create Oracle software owner user accounts, and configure their environments.

- Configuring Grid Infrastructure Software Owner User Environments
  Understand the software owner user environments to configure before installing Oracle Grid Infrastructure.

- Enabling Intelligent Platform Management Interface (IPMI)
  Intelligent Platform Management Interface (IPMI) provides a set of common interfaces to computer hardware and firmware that system administrators can use to monitor system health and manage the system.

- Granting MLOCK Privilege to OSDBA
  Grant the MLOCK privilege to `dba` to avoid Oracle Database installation errors.

Creating Groups, Users and Paths for Oracle Grid Infrastructure

Log in as `root`, and use the following instructions to locate or create the Oracle Inventory group, and create a Oracle Grid Infrastructure software owner, and Oracle home directories.

Oracle software installations require an installation owner, an Oracle Inventory group, which is the primary group of all Oracle installation owners, and at least one group designated as a system privileges group. Review group and user options with your system administrator.

- Determining If an Oracle Inventory and Oracle Inventory Group Exist
  Determine if you have an existing Oracle central inventory, and ensure that you use the same Oracle Inventory for all Oracle software installations. Also, ensure that all Oracle software users you intend to use for installation have permissions to write to this directory.

- Creating the Oracle Inventory Group If an Oracle Inventory Does Not Exist
  If the `oraInst.loc` file does not exist, then create the Oracle Inventory group.

- About Oracle Installation Owner Accounts
  Select or create an Oracle installation owner for your installation, depending on the group and user management plan you want to use for your installations.

- Restrictions for Oracle Software Installation Owners
  Review the following restrictions for users created to own Oracle software.

- Identifying an Oracle Software Owner User Account
  You must create at least one software owner user account the first time you install Oracle software on the system. Either use an existing Oracle software user account, or create an Oracle software owner user account for your installation.

- About the Oracle Base Directory for the grid User
  Review this information about creating the Oracle base directory on each cluster node.

- About the Oracle Home Directory for Oracle Grid Infrastructure Software
  Review this information about creating the Oracle home directory location on each cluster node.

- About the Oracle Home Directory for Grid Infrastructure Management Repository
  Review this information about creating the Oracle home directory location for Grid Infrastructure Management Repository (GIMR) on each cluster node.

- About Creating the Oracle Home and Oracle Base Directory
  Create Grid home and Oracle base home directories on each cluster node.

Determining If an Oracle Inventory and Oracle Inventory Group Exist

Determine if you have an existing Oracle central inventory, and ensure that you use the same Oracle Inventory for all Oracle software installations. Also, ensure that all Oracle software users you intend to use for installation have permissions to write to this directory.

When you install Oracle software on the system for the first time, OUI creates the `oraInst.loc` file. This file identifies the name of the Oracle Inventory group (by default, `oinstall`), and the path of the Oracle central inventory directory. If you have an existing Oracle central inventory, then ensure that you use the same Oracle Inventory for all Oracle software installations, and ensure that all Oracle software users you intend to use for installation have permissions to write to this directory.

```
oraInst.loccentral_inventory_locationgroup

inventory_loc=central_inventory_location
inst_group=group
```

Use the `more` command to determine if you have an Oracle central inventory on your system. For example:

```
# more /var/opt/oracle/oraInst.loc

inventory_loc=/u01/app/oraInventory
inst_group=oinstall
```

Use the command `grep groupname /etc/group` to confirm that the group specified as the Oracle Inventory group still exists on the system. For example:

```
$ grep oinstall /etc/group
oinstall:x:54321:grid,oracle
```

> **Note:**
>
> Do not put the `oraInventory` directory under the Oracle base directory for a new installation, because that can result in user permission errors for other installations.

Creating the Oracle Inventory Group If an Oracle Inventory Does Not Exist

If the `oraInst.loc` file does not exist, then create the Oracle Inventory group.

Members of the OINSTALL group are granted privileges to write to the Oracle central inventory (`oraInventory`), and other system privileges for Oracle installation owner users.

An Oracle installation owner should always have the group you want to have designated as the OINSTALL group (`oinstall`) as its primary group. Ensure that this group is available as the primary group for all planned Oracle software installation owners. By default, if an `oraInst.loc` file does not exist and an Oracle central inventory (`oraInventory`) is not identified, then the installer designates the primary group of the installation owner running the installation as the OINSTALL group.

> **Note:**
>
> For installations on Oracle Clusterware, group and user IDs must be identical on all nodes in the cluster. Ensure that the group and user IDs you want to use are available on each cluster member node, and confirm that the primary group for each Oracle Grid Infrastructure for a cluster installation owner has the same name and group ID.

The following example creates the oraInventory group `oinstall`, with the group ID number 54321.

```
# /usr/sbin/groupadd -g 54321 oinstall
```

About Oracle Installation Owner Accounts

Select or create an Oracle installation owner for your installation, depending on the group and user management plan you want to use for your installations.

You must create a software owner for your installation in the following circumstances:

* If an Oracle software owner user does not exist; for example, if this is the first installation of Oracle software on the system.

- If an Oracle software owner user exists, but you want to use a different operating system user, with different group membership, to separate Oracle Grid Infrastructure administrative privileges from Oracle Database administrative privileges.

In Oracle documentation, a user created to own only Oracle Grid Infrastructure software installations is called the Grid user (`grid`). This user owns both the Oracle Clusterware and Oracle Automatic Storage Management binaries. A user created to own either all Oracle installations, or one or more Oracle database installations, is called the Oracle user (`oracle`). You can have only one Oracle Grid Infrastructure installation owner, but you can have different Oracle users to own different installations.

Oracle software owners must have the Oracle Inventory group as their primary group, so that each Oracle software installation owner can write to the central inventory (oraInventory), and so that OCR and Oracle Clusterware resource permissions are set correctly. The database software owner must also have the OSDBA group and (if you create them) the OSOPER, OSBACKUPDBA, OSDGDBA, OSRACDBA, and OSKMDBA groups as secondary groups.

Restrictions for Oracle Software Installation Owners

Review the following restrictions for users created to own Oracle software.

- If you intend to use multiple Oracle software owners for different Oracle Database homes, then Oracle recommends that you create a separate software owner for Oracle Grid Infrastructure software (Oracle Clusterware and Oracle ASM), and use that owner to run the Oracle Grid Infrastructure installation.

- During installation, SSH must be set up between cluster member nodes. SSH can be set up automatically by Oracle Universal Installer (the installer). To enable SSH to be set up automatically, create Oracle installation owners without any `stty` commands in their profiles, and remove other security measures that are triggered during a login that generate messages to the terminal. These messages, mail checks, and other displays prevent Oracle software installation owner accounts from using the SSH configuration script that is built into the installer. If they are not disabled, then SSH must be configured manually before an installation can be run.

- If you plan to install Oracle Database or Oracle RAC, then Oracle recommends that you create separate users for the Oracle Grid Infrastructure and the Oracle Database installations. If you use one installation owner, then when you want to perform administration tasks, you must change the value for `$ORACLE_HOME` to the instance you want to administer (Oracle ASM, in the Oracle Grid Infrastructure home, or the database in the Oracle home), using command syntax such as the following example, where `/u01/app/21.0.0/grid` is the Oracle Grid Infrastructure home:

  ```
  $ export ORACLE_HOME=/u01/app/21.0.0/grid
  ```

- If you try to administer an Oracle home or Grid home instance using `sqlplus`, `lsnrctl`, or `asmcmd` commands while the environment variable `$ORACLE_HOME` is set to a different Oracle home or Grid home path, then you encounter errors. For example, when you start SRVCTL from a database home, `$ORACLE_HOME` should be set to that database home, or SRVCTL fails. The exception is when you are using SRVCTL in the Oracle Grid Infrastructure home. In that case, `$ORACLE_HOME` is ignored, and the Oracle home environment variable does not affect SRVCTL

commands. In all other cases, you must change `$ORACLE_HOME` to the instance that you want to administer.

- To create separate Oracle software owners and separate operating system privileges groups for different Oracle software installations, note that each of these users must have the Oracle central inventory group (`oraInventory` group) as their primary group. Members of this group are granted the OINSTALL system privileges to write to the Oracle central inventory (`oraInventory`) directory, and are also granted permissions for various Oracle Clusterware resources, OCR keys, directories in the Oracle Clusterware home to which DBAs need write access, and other necessary privileges. Members of this group are also granted execute permissions to start and stop Clusterware infrastructure resources and databases. In Oracle documentation, this group is represented as `oinstall` in code examples.

- Each Oracle software owner must be a member of the same central inventory `oraInventory` group, and they must have this group as their primary group, so that all Oracle software installation owners share the same OINSTALL system privileges. Oracle recommends that you do not have more than one central inventory for Oracle installations. If an Oracle software owner has a different central inventory group, then you may corrupt the central inventory.

Identifying an Oracle Software Owner User Account

You must create at least one software owner user account the first time you install Oracle software on the system. Either use an existing Oracle software user account, or create an Oracle software owner user account for your installation.

To use an existing user account, obtain the name of an existing Oracle installation owner from your system administrator. Confirm that the existing owner is a member of the Oracle Inventory group.

For example, if you know that the name of the Oracle Inventory group is `oinstall`, then an Oracle software owner should be listed as a member of `oinstall`:

```
$ grep "oinstall" /etc/group
oinstall:x:54321:grid,oracle
```

You can then use the ID command to verify that the Oracle installation owners you intend to use have the Oracle Inventory group as their primary group. For example:

```
$ id oracle
uid=54321(oracle) gid=54321(oinstall) groups=54321(oinstall),54322(dba),
54323(oper),54324(backupdba),54325(dgdba),54326(kmdba),54327(asmdba),54330(ra
cdba)
```

```
$ id grid
uid=54331(grid) gid=54321(oinstall) groups=54321(oinstall),54322(dba),
54327(asmdba),54328(asmoper),54329(asmadmin),54330(racdba)
```

For Oracle Restart installations, to successfully install Oracle Database, ensure that the `grid` user is a member of the `racdba` group.

After you create operating system groups, create or modify Oracle user accounts in accordance with your operating system authentication planning.

About the Oracle Base Directory for the grid User

Review this information about creating the Oracle base directory on each cluster node.

The Oracle base directory for the Oracle Grid Infrastructure installation is the location where diagnostic and administrative logs, and other logs associated with Oracle ASM and Oracle Clusterware are stored. For Oracle installations other than Oracle Grid Infrastructure for a cluster, it is also the location under which an Oracle home is placed.

However, in the case of an Oracle Grid Infrastructure installation, you must create a *different* path, so that the path for Oracle bases remains available for other Oracle installations.

For OUI to recognize the Oracle base path, it must be in the form `u[`*`00-99`*`][00-99]/app`, and it must be writable by any member of the oraInventory (`oinstall`) group. The OFA path for the Oracle base is `u[00-99][00-99]/app/`*`user`*, where *user* is the name of the software installation owner. For example:

```
/u01/app/grid
```

> **Note:**
>
> Oracle home or Oracle base cannot be `symlinks`, nor can any of their parent directories, all the way to up to the `root` directory.

About the Oracle Home Directory for Oracle Grid Infrastructure Software

Review this information about creating the Oracle home directory location on each cluster node.

The Oracle home for Oracle Grid Infrastructure software (Grid home) should be located in a path that is different from the Oracle home directory paths for any other Oracle software. The Optimal Flexible Architecture guideline for a Grid home is to create a path in the form `/pm/v/u`, where `p` is a string constant, `m` is a unique fixed-length key (typically a two-digit number), `v` is the version of the software, and `u` is the installation owner of the Oracle Grid Infrastructure software (`grid` user). During Oracle Grid Infrastructure for a cluster installation, the path of the Grid home is changed to the root user, so any other users are unable to read, write, or execute commands in that path. For example, to create a Grid home in the standard mount point path format `u[00-99][00-99]/app/`*`release`*`/grid`, where release is the release number of the Oracle Grid Infrastructure software, create the following path:

```
/u01/app/21.0.0/grid
```

> **✏️ Note:**
>
> Oracle home or Oracle base cannot be `symlinks`, nor can any of their parent directories, all the way to up to the `root` directory.

During installation, ownership of the entire path to the Grid home is changed to `root` (`/u01`, `/u01/app`, `/u01/app/21.0.0`, `/u01/app/21.0.0/grid`). If you do not create a unique path to the Grid home, then after the Grid install, you can encounter permission errors for other installations, including any existing installations under the same path. To avoid placing the application directory in the mount point under root ownership, you can create and select paths such as the following for the Grid home:

```
/u01/21.0.0/grid
```

> **⚠️ Caution:**
>
> For Oracle Grid Infrastructure for a cluster installations, note the following restrictions for the Oracle Grid Infrastructure binary home (Grid home directory for Oracle Grid Infrastructure):
>
> • It must not be placed under one of the Oracle base directories, including the Oracle base directory of the Oracle Grid Infrastructure installation owner.
>
> • It must not be placed in the home directory of an installation owner. These requirements are specific to Oracle Grid Infrastructure for a cluster installations.
>
> Oracle Grid Infrastructure for a standalone server (Oracle Restart) can be installed under the Oracle base for the Oracle Database installation.

About the Oracle Home Directory for Grid Infrastructure Management Repository

Review this information about creating the Oracle home directory location for Grid Infrastructure Management Repository (GIMR) on each cluster node.

The Oracle home for Grid Infrastructure Management Repository (GIMR home) should be located in a path that is different from the Oracle home directory paths for any other Oracle software. The Optimal Flexible Architecture guideline for a GIMR home is to create a path in the form `/pm/v/u`, where `p` is a string constant, `m` is a unique fixed-length key (typically a two-digit number), `v` is the version of the software, and `u` is the directory name (`gimr`). For example, to create a GIMR home in the standard mount point path format `u[00-99][00-99]/app/release/gimr`, where release is the release number of the Oracle Grid Infrastructure software, create the following path:

```
/u01/app/21.0.0/gimr
```

> **✎ Note:**
>
> The GIMR home directory must be owned by the `grid` user and you must not create user databases in this directory.

> **⚠ Caution:**
>
> For Grid Infrastructure Management Repository installations, note the following restrictions for the Grid Infrastructure Management Repository binary home (GIMR home directory):
>
> - It must not be placed under one of the Oracle base directories, including the Oracle base directory of the Oracle Grid Infrastructure installation owner.
>
> - It must not be placed in the home directory of an installation owner. These requirements are specific to Grid Infrastructure Management Repository installations.

About Creating the Oracle Home and Oracle Base Directory

Create Grid home and Oracle base home directories on each cluster node.

Oracle recommends that you create Oracle Grid Infrastructure Grid home and Oracle base homes manually, particularly if you have separate Oracle Grid Infrastructure for a cluster and Oracle Database software owners, so that you can separate log files for the Oracle Grid Infrastructure installation owner in a separate Oracle base, and prevent accidental placement of the Grid home under an Oracle base path. For example:

```
# mkdir -p /u01/app/19.0.0/grid
# mkdir -p /u01/app/grid
# mkdir -p /u01/app/oracle
# chown -R grid:oinstall /u01
# chown oracle:oinstall /u01/app/oracle
# chmod -R 775 /u01/
```

> **✎ Note:**
>
> Placing Oracle Grid Infrastructure for a cluster binaries on a cluster file system is not supported.
>
> If you plan to install an Oracle RAC home on a shared OCFS2 location, then you must upgrade OCFS2 to at least version 1.4.1, which supports shared writable `maps`.
>
> Oracle recommends that you install Oracle Grid Infrastructure locally, on each cluster member node. Using a shared Grid home prevents rolling upgrades, and creates a single point of failure for the cluster.

Oracle Installations with Standard and Job Role Separation Groups and Users

A job role separation configuration of Oracle Database and Oracle ASM is a configuration with groups and users to provide separate groups for operating system authentication.

Review the following sections to understand more about a Job Role Separation deployment:

- About Oracle Installations with Job Role Separation
  Job role separation requires that you create different operating system groups for each set of system privileges that you grant through operating system authorization.

- Standard Oracle Database Groups for Database Administrators
  Oracle Database has two standard administration groups: OSDBA, which is required, and OSOPER, which is optional.

- Extended Oracle Database Groups for Job Role Separation
  Oracle Database provides an extended set of database groups to grant task-specific system privileges for database administration.

- Creating an ASMSNMP User
  The `ASMSNMP` user is an Oracle ASM user with privileges to monitor Oracle ASM instances. You are prompted to provide a password for this user during installation.

- Oracle Automatic Storage Management Groups for Job Role Separation
  Oracle Grid Infrastructure operating system groups provide their members task-specific system privileges to access and to administer Oracle Automatic Storage Management.

About Oracle Installations with Job Role Separation

Job role separation requires that you create different operating system groups for each set of system privileges that you grant through operating system authorization.

With Oracle Grid Infrastructure job role separation, Oracle ASM has separate operating system groups that provide operating system authorization for Oracle ASM system privileges for storage tier administration. This operating system authorization is separated from Oracle Database operating system authorization. In addition, the Oracle Grid Infrastructure installation owner provides operating system user authorization for modifications to Oracle Grid Infrastructure binaries.

With Oracle Database job role separation, each Oracle Database installation has separate operating system groups to provide authorization for system privileges on that Oracle Database. Multiple databases can, therefore, be installed on the cluster without sharing operating system authorization for system privileges. In addition, each Oracle software installation is owned by a separate installation owner, to provide operating system user authorization for modifications to Oracle Database binaries. Note that any Oracle software owner can start and stop all databases and shared Oracle Grid Infrastructure resources such as Oracle ASM or Virtual IP (VIP). Job role separation configuration enables database security, and does not restrict user roles in starting and stopping various Oracle Clusterware resources.

You can choose to create one administrative user and one group for operating system authentication for all system privileges on the storage and database tiers. For example, you can designate the `oracle` user to be the installation owner for all Oracle software, and designate `oinstall` to be the group whose members are granted all system privileges for

Oracle Clusterware; all system privileges for Oracle ASM; all system privileges for all Oracle Databases on the servers; and all `OINSTALL` system privileges for installation owners. This group must also be the Oracle Inventory group.

If you do not want to use role allocation groups, then Oracle strongly recommends that you use at least two groups:

- A system privileges group whose members are granted administrative system privileges, including OSDBA, OSASM, and other system privileges groups.

- An installation owner group (the `oraInventory` group) whose members are granted Oracle installation owner system privileges (the `OINSTALL` system privilege).

> **✎ Note:**
>
> To configure users for installation that are on a network directory service such as Network Information Services (NIS), refer to your directory service documentation.

Related Topics

- *Oracle Database Administrator's Guide*
- *Oracle Automatic Storage Management Administrator's Guide*

Standard Oracle Database Groups for Database Administrators

Oracle Database has two standard administration groups: OSDBA, which is required, and OSOPER, which is optional.

- The OSDBA group (typically, `dba`)

  You must create this group the first time you install Oracle Database software on the system. This group identifies operating system user accounts that have database administrative privileges (the `SYSDBA` privilege).

  If you do not create separate OSDBA, OSOPER, and OSASM groups for the Oracle ASM instance, then operating system user accounts that have the `SYSOPER` and `SYSASM` privileges must be members of this group. The name used for this group in Oracle code examples is `dba`. If you do not designate a separate group as the OSASM group, then the OSDBA group you define is also by default the OSASM group.

- The OSOPER group for Oracle Database (typically, `oper`)

  OSOPER grants the OPERATOR privilege to start up and shut down the database (the `SYSOPER` privilege). By default, members of the OSDBA group have all privileges granted by the `SYSOPER` privilege.

Extended Oracle Database Groups for Job Role Separation

Oracle Database provides an extended set of database groups to grant task-specific system privileges for database administration.

The extended set of Oracle Database system privileges groups are task-specific and less privileged than the OSDBA/SYSDBA system privileges. They are designed to

provide privileges to carry out everyday database operations. Users granted these system privileges are also authorized through operating system group membership.

You do not have to create these specific group names, but during interactive and silent installation, you must assign operating system groups whose members are granted access to these system privileges. You can assign the same group to provide authorization for these privileges, but Oracle recommends that you provide a unique group to designate each privilege.

The subset of OSDBA job role separation privileges and groups consist of the following:

- OSBACKUPDBA group for Oracle Database (typically, `backupdba`)

  Create this group if you want a separate group of operating system users to have a limited set of database backup and recovery related administrative privileges (the SYSBACKUP privilege).

- OSDGDBA group for Oracle Data Guard (typically, `dgdba`)

  Create this group if you want a separate group of operating system users to have a limited set of privileges to administer and monitor Oracle Data Guard (the SYSDG privilege). To use this privilege, add the Oracle Database installation owners as members of this group.

- The OSKMDBA group for encryption key management (typically, `kmdba`)

  Create this group if you want a separate group of operating system users to have a limited set of privileges for encryption key management such as Oracle Wallet Manager management (the SYSKM privilege). To use this privilege, add the Oracle Database installation owners as members of this group.

- The OSRACDBA group for Oracle Real Application Clusters Administration (typically, `racdba`)

  Create this group if you want a separate group of operating system users to have a limited set of Oracle Real Application Clusters (RAC) administrative privileges (the SYSRAC privilege). To use this privilege:

  – Add the Oracle Database installation owners as members of this group.

  – For Oracle Restart configurations, if you have a separate Oracle Grid Infrastructure installation owner user (`grid`), then you must also add the `grid` user as a member of the OSRACDBA group of the database to enable Oracle Grid Infrastructure components to connect to the database.

Related Topics

- *Oracle Database Administrator's Guide*
- *Oracle Database Security Guide*

Creating an ASMSNMP User

The `ASMSNMP` user is an Oracle ASM user with privileges to monitor Oracle ASM instances. You are prompted to provide a password for this user during installation.

In addition to the OSASM group, whose members are granted the SYSASM system privilege to administer Oracle ASM, Oracle recommends that you create a less privileged user, `ASMSNMP`, and grant that user SYSDBA privileges to monitor the Oracle ASM instance. Oracle Enterprise Manager uses the `ASMSNMP` user to monitor Oracle ASM status.

During installation, you are prompted to provide a password for the `ASMSNMP` user. You can create an operating system authenticated user, or you can create an Oracle Database user called `asmsnmp`. In either case, grant the user SYSDBA privileges.

Oracle Automatic Storage Management Groups for Job Role Separation

Oracle Grid Infrastructure operating system groups provide their members task-specific system privileges to access and to administer Oracle Automatic Storage Management.

- The OSASM group for Oracle ASM Administration (typically, `asmadmin`)

  Create this group as a separate group to separate administration privileges groups for Oracle ASM and Oracle Database administrators. Members of this group are granted the SYSASM system privileges to administer Oracle ASM. In Oracle documentation, the operating system group whose members are granted privileges is called the OSASM group, and in code examples, where there is a group specifically created to grant this privilege, it is referred to as `asmadmin`.

  Oracle ASM can support multiple databases. If you have multiple databases on your system, and use multiple OSDBA groups so that you can provide separate SYSDBA privileges for each database, then you should create a group whose members are granted the OSASM/SYSASM administrative privileges, and create a grid infrastructure user (grid) that does not own a database installation, so that you separate Oracle Grid Infrastructure SYSASM administrative privileges from a database administrative privileges group.

  Members of the OSASM group can use SQL to connect to an Oracle ASM instance as SYSASM using operating system authentication. The SYSASM privileges permit mounting and dismounting disk groups, and other storage administration tasks. SYSASM privileges provide no access privileges on an RDBMS instance.

  If you do not designate a separate group as the OSASM group, but you do define an OSDBA group for database administration, then by default the OSDBA group you define is also defined as the OSASM group.

- The OSOPER group for Oracle ASM (typically, `asmoper`)

  This is an optional group. Create this group if you want a separate group of operating system users to have a limited set of Oracle instance administrative privileges (the SYSOPER for ASM privilege), including starting up and stopping the Oracle ASM instance. By default, members of the OSASM group also have all privileges granted by the SYSOPER for ASM privilege.

Creating Operating System Privileges Groups

The following sections describe how to create operating system groups for Oracle Grid Infrastructure and Oracle Database:

- Creating the OSASM Group
  If the OSASM group does not exist, or if you require a new OSASM group, then create it.

- **Creating the OSDBA for ASM Group**
  You must designate a group as the OSDBA for ASM (`asmdba`) group during installation.
  Members of this group are granted access privileges to Oracle Automatic Storage
  Management.

- **Creating the OSOPER for ASM Group**
  You can choose to designate a group as the OSOPER for ASM group (`asmoper`) during
  installation. Members of this group are granted startup and shutdown privileges to Oracle
  Automatic Storage Management.

- **Creating the OSDBA Group for Database Installations**
  Each Oracle Database requires an operating system group to be designated as the
  OSDBA group. Members of this group are granted the SYSDBA system privileges to
  administer the database.

- **Creating an OSOPER Group for Database Installations**
  Create an OSOPER group only if you want to identify a group of operating system users
  with a limited set of database administrative privileges (SYSOPER operator privileges).

- **Creating the OSBACKUPDBA Group for Database Installations**
  You must designate a group as the OSBACKUPDBA group during installation. Members
  of this group are granted the SYSBACKUP privileges to perform backup and recovery
  operations using RMAN or SQL*Plus.

- **Creating the OSDGDBA Group for Database Installations**
  You must designate a group as the OSDGDBA group during installation. Members of this
  group are granted the SYSDG privileges to perform Data Guard operations.

- **Creating the OSKMDBA Group for Database Installations**
  You must designate a group as the OSKMDBA group during installation. Members of this
  group are granted the SYSKM privileges to perform Transparent Data Encryption
  keystore operations.

- **Creating the OSRACDBA Group for Database Installations**
  You must designate a group as the OSRACDBA group during database installation.
  Members of this group are granted the SYSRAC privileges to perform day–to–day
  administration of Oracle databases on an Oracle RAC cluster.

Creating the OSASM Group

If the OSASM group does not exist, or if you require a new OSASM group, then create it.

Use the group name `asmadmin` unless a group with that name already exists. For example:

```
# groupadd -g 54329 asmadmin
```

Creating the OSDBA for ASM Group

You must designate a group as the OSDBA for ASM (`asmdba`) group during installation.
Members of this group are granted access privileges to Oracle Automatic Storage
Management.

Create an OSDBA for ASM group using the group name `asmdba` unless a group with that
name already exists:

```
# /usr/sbin/groupadd -g 54327 asmdba
```

Creating the OSOPER for ASM Group

You can choose to designate a group as the OSOPER for ASM group (`asmoper`) during installation. Members of this group are granted startup and shutdown privileges to Oracle Automatic Storage Management.

If you want to create an OSOPER for ASM group, use the group name `asmoper` unless a group with that name already exists:

```
# /usr/sbin/groupadd -g 54328 asmoper
```

Creating the OSDBA Group for Database Installations

Each Oracle Database requires an operating system group to be designated as the OSDBA group. Members of this group are granted the SYSDBA system privileges to administer the database.

You must create an OSDBA group in the following circumstances:

- An OSDBA group does not exist, for example, if this is the first installation of Oracle Database software on the system

- An OSDBA group exists, but you want to give a different group of operating system users database administrative privileges for a new Oracle Database installation

Create the OSDBA group using the group name `dba`, unless a group with that name already exists:

```
# /usr/sbin/groupadd -g 54322 dba
```

Creating an OSOPER Group for Database Installations

Create an OSOPER group only if you want to identify a group of operating system users with a limited set of database administrative privileges (SYSOPER operator privileges).

For most installations, it is sufficient to create only the OSDBA group. However, to use an OSOPER group, create it in the following circumstances:

- If an OSOPER group does not exist; for example, if this is the first installation of Oracle Database software on the system

- If an OSOPER group exists, but you want to give a different group of operating system users database operator privileges in a new Oracle installation

If the OSOPER group does not exist, or if you require a new OSOPER group, then create it. Use the group name `oper` unless a group with that name already exists. For example:

```
# groupadd -g 54323 oper
```

ORACLE®

Creating the OSBACKUPDBA Group for Database Installations

You must designate a group as the OSBACKUPDBA group during installation. Members of this group are granted the SYSBACKUP privileges to perform backup and recovery operations using RMAN or SQL*Plus.

Create the OSBACKUPDBA group using the group name `backupdba`, unless a group with that name already exists:

```
# /usr/sbin/groupadd -g 54324 backupdba
```

Creating the OSDGDBA Group for Database Installations

You must designate a group as the OSDGDBA group during installation. Members of this group are granted the SYSDG privileges to perform Data Guard operations.

Create the OSDGDBA group using the group name `dgdba`, unless a group with that name already exists:

```
# /usr/sbin/groupadd -g 54325 dgdba
```

Creating the OSKMDBA Group for Database Installations

You must designate a group as the OSKMDBA group during installation. Members of this group are granted the SYSKM privileges to perform Transparent Data Encryption keystore operations.

If you want a separate group for Transparent Data Encryption, then create the OSKMDBA group using the group name `kmdba` unless a group with that name already exists:

```
# /usr/sbin/groupadd -g 54326 kmdba
```

Creating the OSRACDBA Group for Database Installations

You must designate a group as the OSRACDBA group during database installation. Members of this group are granted the SYSRAC privileges to perform day–to–day administration of Oracle databases on an Oracle RAC cluster.

Create the OSRACDBA group using the groups name `racdba` unless a group with that name already exists:

```
# /usr/sbin/groupadd -g 54330 racdba
```

Creating Operating System Oracle Installation User Accounts

Before starting installation, create Oracle software owner user accounts, and configure their environments.

Oracle software owner user accounts require resource settings and other environment configuration. To protect against accidents, Oracle recommends that you create one software installation owner account for each Oracle software program you install.

- Creating an Oracle Software Owner User
  If the Oracle software owner user (`oracle` or `grid`) does not exist, or if you require a new Oracle software owner user, then create it as described in this section.

- Modifying Oracle Owner User Groups
  If you have created an Oracle software installation owner account, but it is not a member of the groups you want to designate as the OSDBA, OSOPER, OSDBA for ASM, ASMADMIN, or other system privileges group, then modify the group settings for that user before installation.

- Identifying Existing User and Group IDs
  To create identical users and groups, you must identify the user ID and group IDs assigned them on the node where you created them, and then create the user and groups with the same name and ID on the other cluster nodes.

- Creating Identical Database Users and Groups on Other Cluster Nodes
  Oracle software owner users and the Oracle Inventory, OSDBA, and OSOPER groups must exist and be identical on all cluster nodes.

- Example of Creating Minimal Groups, Users, and Paths
  You can create a minimal operating system authentication configuration as described in this example.

- Example of Creating Role-allocated Groups, Users, and Paths
  Understand this example of how to create role-allocated groups and users that is compliant with an Optimal Flexible Architecture (OFA) deployment.

Creating an Oracle Software Owner User

If the Oracle software owner user (`oracle` or `grid`) does not exist, or if you require a new Oracle software owner user, then create it as described in this section.

The following example shows how to create the user `oracle` with the user ID 54321; with the primary group `oinstall`; and with secondary groups `dba`, `asmdba`, `backupdba`, `dgdba`, `kmdba`, and `racdba`:

```
# /usr/sbin/useradd -u 54321 -g oinstall -G
dba,asmdba,backupdba,dgdba,kmdba,racdba oracle
```

You must note the user ID number for installation users, because you need it during preinstallation.

For Oracle Grid Infrastructure installations, user IDs and group IDs must be identical on all candidate nodes.

Modifying Oracle Owner User Groups

If you have created an Oracle software installation owner account, but it is not a member of the groups you want to designate as the OSDBA, OSOPER, OSDBA for

ASM, ASMADMIN, or other system privileges group, then modify the group settings for that user before installation.

> **⬦ Warning:**
>
> Each Oracle software owner must be a member of the same central inventory group. Do not modify the primary group of an existing Oracle software owner account, or designate different groups as the OINSTALL group. If Oracle software owner accounts have different groups as their primary group, then you can corrupt the central inventory.

During installation, the user that is installing the software should have the OINSTALL group as its primary group, and it must be a member of the operating system groups appropriate for your installation. For example:

```
# /usr/sbin/usermod -g oinstall -G
dba,asmdba,backupdba,dgdba,kmdba,racdba[,oper] oracle
```

Identifying Existing User and Group IDs

To create identical users and groups, you must identify the user ID and group IDs assigned them on the node where you created them, and then create the user and groups with the same name and ID on the other cluster nodes.

1. Enter a command similar to the following (in this case, to determine a user ID for the `oracle` user):

   ```
   # id oracle
   ```

   The output from this command is similar to the following:

   ```
   uid=54321(oracle) gid=54421(oinstall)
   groups=54322(dba),54323(oper),54327(asmdba)
   ```

2. From the output, identify the user ID (`uid`) for the user and the group identities (`gids`) for the groups to which it belongs.

   Ensure that these ID numbers are identical on each node of the cluster. The user's primary group is listed after `gid`. Secondary groups are listed after groups.

Creating Identical Database Users and Groups on Other Cluster Nodes

Oracle software owner users and the Oracle Inventory, OSDBA, and OSOPER groups must exist and be identical on all cluster nodes.

To create users and groups on the other cluster nodes, repeat the following procedure on each node:

You must complete the following procedures only if you are using local users and groups. If you are using users and groups defined in a directory service such as NIS, then they are already identical on each cluster node.

1. Log in to the node as `root`.

2. Enter commands similar to the following to create the `asmadmin`, `asmdba`, `backupdba`, `dgdba`, `kmdba`, `asmoper`, `racdba`, and `oper` groups, and if not configured by the Oracle Database Preinstallation RPM or prior installations, then the `oinstall` and `dba` groups.

   Use the `-g` option to specify the correct group ID for each group.

   ```
   # groupadd -g 54421 oinstall
   # groupadd -g 54322 dba
   # groupadd -g 54323 oper
   # groupadd -g 54324 backupdba
   # groupadd -g 54325 dgdba
   # groupadd -g 54326 kmdba
   # groupadd -g 54327 asmdba
   # groupadd -g 54328 asmoper
   # groupadd -g 54329 asmadmin
   # groupadd -g 54330 racdba
   ```

   > **Note:**
   >
   > You are not required to use the UIDs and GIDs in this example. If a group already exists, then use the `groupmod` command to modify it if necessary. If you cannot use the same group ID for a particular group on a node, then view the `/etc/group` file on all nodes to identify a group ID that is available on every node. You must then change the group ID on all nodes to the same group ID.

3. To create the Oracle Grid Infrastructure (`grid`) user, enter a command similar to the following:

   ```
   # useradd -u 54322 -g oinstall -G asmadmin,asmdba,racdba grid
   ```

   - The `-u` option specifies the user ID, which must be the user ID that you identified earlier.

   - The `-g` option specifies the primary group for the Grid user, which must be the Oracle Inventory group (OINSTALL), which grants the OINSTALL system privileges. In this example, the OINSTALL group is `oinstall`.

   - The `-G` option specifies the secondary groups. The Grid user must be a member of the OSASM group (`asmadmin`) and the OSDBA for ASM group (`asmdba`).

   > **Note:**
   >
   > If the user already exists, then use the `usermod` command to modify it if necessary. If you cannot use the same user ID for the user on every node, then view the `/etc/passwd` file on all nodes to identify a user ID that is available on every node. You must then specify that ID for the user on all of the nodes.

ORACLE

4. Set the password of the user.

   For example:

   ```
   # passwd grid
   ```

5. Complete user environment configuration tasks for each user.

Example of Creating Minimal Groups, Users, and Paths

You can create a minimal operating system authentication configuration as described in this example.

This configuration example shows the following:

- Creation of the Oracle Inventory group (`oinstall`)

- Creation of a single group (`dba`) as the only system privileges group to assign for all Oracle Grid Infrastructure, Oracle ASM, and Oracle Database system privileges

- Creation of the Oracle Grid Infrastructure software owner (`grid`), and one Oracle Database owner (`oracle`) with correct group memberships

- Creation and configuration of an Oracle base path compliant with OFA structure with correct permissions

Enter the following commands to create a minimal operating system authentication configuration:

```
# groupadd -g 54421 oinstall
# groupadd -g 54422 dba
# useradd -u 54321 -g oinstall -G dba oracle
# useradd -u 54322 -g oinstall -G dba grid
# mkdir -p  /u01/app/21.0.0/grid
# mkdir -p /u01/app/grid
# mkdir -p /u01/app/oracle
# chown -R grid:oinstall /u01
# chown oracle:oinstall /u01/app/oracle
# chmod -R 775 /u01/
```

After running these commands, you have the following groups and users:

- An Oracle central inventory group, or oraInventory group (`oinstall`). Members who have the central inventory group as their primary group, are granted the OINSTALL permission to write to the oraInventory directory.

- One system privileges group, `dba`, for Oracle Grid Infrastructure, Oracle ASM and Oracle Database system privileges. Members who have the `dba` group as their primary or secondary group are granted operating system authentication for OSASM/SYSASM, OSDBA/SYSDBA, OSOPER/SYSOPER, OSBACKUPDBA/SYSBACKUP, OSDGDBA/SYSDG, OSKMDBA/SYSKM, OSDBA for ASM/SYSDBA for ASM, and OSOPER for ASM/SYSOPER for Oracle ASM to administer Oracle Clusterware, Oracle ASM, and Oracle Database, and are granted SYSASM and OSOPER for Oracle ASM access to the Oracle ASM storage.

- An Oracle Grid Infrastructure for a cluster owner, or Grid user (`grid`), with the oraInventory group (`oinstall`) as its primary group, and with the OSASM group (`dba`) as the secondary group, with its Oracle base directory `/u01/app/grid`.

- An Oracle Database owner (`oracle`) with the oraInventory group (`oinstall`) as its primary group, and the OSDBA group (`dba`) as its secondary group, with its Oracle base directory `/u01/app/oracle`.

- `/u01/app` owned by `grid:oinstall` with 775 permissions before installation, and by root after the `root.sh` script is run during installation. This ownership and permissions enables OUI to create the Oracle Inventory directory, in the path `/u01/app/oraInventory`.

- `/u01` owned by `grid:oinstall` before installation, and by root after the `root.sh` script is run during installation.

- `/u01/app/21.0.0/grid` owned by `grid:oinstall` with 775 permissions. These permissions are required for installation, and are changed during the installation process.

- `/u01/app/grid` owned by `grid:oinstall` with 775 permissions. These permissions are required for installation, and are changed during the installation process.

- `/u01/app/oracle` owned by `oracle:oinstall` with 775 permissions.

> **Note:**
>
> You can use one installation owner for both Oracle Grid Infrastructure and any other Oracle installations. However, Oracle recommends that you use separate installation owner accounts for each Oracle software installation.

Example of Creating Role-allocated Groups, Users, and Paths

Understand this example of how to create role-allocated groups and users that is compliant with an Optimal Flexible Architecture (OFA) deployment.

Purpose

This example illustrates the following scenario:

- An Oracle Grid Infrastructure installation

- Two separate Oracle Database installations planned for the cluster, `DB1` and `DB2`

- Separate installation owners for Oracle Grid Infrastructure, and for each Oracle Database

- Full role allocation of system privileges for Oracle ASM, and for each Oracle Database

- Oracle Database owner `oracle1` granted the right to start up and shut down the Oracle ASM instance

Create groups and users for a role-allocated configuration for this scenario using the following commands:

```
# groupadd -g 54321 oinstall
# groupadd -g 54322 dba1
# groupadd -g 54332 dba2
# groupadd -g 54323 oper1
# groupadd -g 54333 oper2
```

```
# groupadd -g 54324 backupdba1
# groupadd -g 54334 backupdba2
# groupadd -g 54325 dgdba1
# groupadd -g 54335 dgdba2
# groupadd -g 54326 kmdba1
# groupadd -g 54336 kmdba2
# groupadd -g 54327 asmdba
# groupadd -g 54328 asmoper
# groupadd -g 54329 asmadmin
# groupadd -g 54330 racdba1
# groupadd -g 54340 racdba2
# useradd -u 54322 -g oinstall -G asmadmin,asmdba,racdba1,racdba2 grid
# useradd -u 54321 -g oinstall -G
dba1,backupdba1,dgdba1,kmdba1,asmdba,racdba1,asmoper oracle1
# useradd -u 54323 -g oinstall -G
dba2,backupdba2,dgdba2,kmdba2,asmdba,racdba2 oracle2
# mkdir -p /u01/app/21.0.0/grid
# mkdir -p /u01/app/grid
# mkdir -p /u01/app/oracle1
# mkdir -p /u01/app/oracle2
# chown -R grid:oinstall /u01
# chmod -R 775 /u01/
# chown oracle1:oinstall /u01/app/oracle1
# chown oracle2:oinstall /u01/app/oracle2
```

After running these commands, you have a set of administrative privileges groups and users for Oracle Grid Infrastructure, and for two separate Oracle databases (DB1 and DB2):

Example 5-1   Oracle Grid Infrastructure Groups and Users Example

The command creates the following Oracle Grid Infrastructure groups and users:

- An Oracle central inventory group, or oraInventory group (oinstall), whose members that have this group as their primary group. Members of this group are granted the OINSTALL system privileges, which grants permissions to write to the oraInventory directory, and other associated install binary privileges.

- An OSASM group (asmadmin), associated with Oracle Grid Infrastructure during installation, whose members are granted the SYSASM privileges to administer Oracle ASM.

- An OSDBA for ASM group (asmdba), associated with Oracle Grid Infrastructure storage during installation. Its members include grid and any database installation owners, such as oracle1 and oracle2, who are granted access to Oracle ASM. Any additional installation owners that use Oracle ASM for storage must also be made members of this group.

- An OSOPER for ASM group for Oracle ASM (asmoper), associated with Oracle Grid Infrastructure during installation. Members of asmoper group are granted limited Oracle ASM administrator privileges, including the permissions to start and stop the Oracle ASM instance.

- An Oracle Grid Infrastructure installation owner (grid), with the oraInventory group (oinstall) as its primary group, and with the OSASM (asmadmin) group and the OSDBA for ASM (asmdba) group as secondary groups.

- `/u01/app/oraInventory`. The central inventory of Oracle installations on the cluster. This path remains owned by `grid:oinstall`, to enable other Oracle software owners to write to the central inventory.

- An OFA-compliant mount point `/u01` owned by `grid:oinstall` before installation, so that Oracle Universal Installer can write to that path.

- An Oracle base for the grid installation owner `/u01/app/grid` owned by `grid:oinstall` with 775 permissions, and changed during the installation process to 755 permissions.

- A Grid home `/u01/app/21.0.0/grid` owned by `grid:oinstall` with 775 (drwxdrwxr-x) permissions. These permissions are required for installation, and are changed during the installation process to `root:oinstall` with 755 permissions (`drwxr-xr-x`).

Example 5-2   Oracle Database DB1 Groups and Users Example

The command creates the following Oracle Database (`DB1`) groups and users:

- An Oracle Database software owner (`oracle1`), which owns the Oracle Database binaries for `DB1`. The `oracle1` user has the oraInventory group as its primary group, and the OSDBA group for its database (`dba1`) and the OSDBA for ASM group for Oracle Grid Infrastructure (`asmdba`) as secondary groups. In addition, the `oracle1` user is a member of `asmoper`, granting that user privileges to start up and shut down Oracle ASM.

- An OSDBA group (`dba1`). During installation, you identify the group `dba1` as the OSDBA group for the database installed by the user `oracle1`. Members of `dba1` are granted the SYSDBA privileges for the Oracle Database `DB1`. Users who connect as SYSDBA are identified as user SYS on `DB1`.

- An OSBACKUPDBA group (`backupdba1`). During installation, you identify the group `backupdba1` as the OSDBA group for the database installed by the user `oracle1`. Members of `backupdba1` are granted the SYSBACKUP privileges for the database installed by the user `oracle1` to back up the database.

- An OSDGDBA group (`dgdba1`). During installation, you identify the group `dgdba1` as the OSDGDBA group for the database installed by the user `oracle1`. Members of `dgdba1` are granted the SYSDG privileges to administer Oracle Data Guard for the database installed by the user `oracle1`.

- An OSKMDBA group (`kmdba1`). During installation, you identify the group `kmdba1` as the OSKMDBA group for the database installed by the user `oracle1`. Members of `kmdba1` are granted the SYSKM privileges to administer encryption keys for the database installed by the user `oracle1`.

- An OSOPER group (`oper1`). During installation, you identify the group `oper1` as the OSOPER group for the database installed by the user `oracle1`. Members of `oper1` are granted the SYSOPER privileges (a limited set of the SYSDBA privileges), including the right to start up and shut down the `DB1` database. Users who connect as OSOPER privileges are identified as user PUBLIC on `DB1`.

- An Oracle base `/u01/app/oracle1` owned by `oracle1:oinstall` with 775 permissions. The user `oracle1` has permissions to install software in this directory, but in no other directory in the `/u01/app` path.

Example 5-3   Oracle Database DB2 Groups and Users Example

The command creates the following Oracle Database (DB2) groups and users:

- An Oracle Database software owner (`oracle2`), which owns the Oracle Database binaries for DB2. The `oracle2` user has the oraInventory group as its primary group, and the OSDBA group for its database (`dba2`) and the OSDBA for ASM group for Oracle Grid Infrastructure (`asmdba`) as secondary groups. However, the `oracle2` user is not a member of the `asmoper` group, so `oracle2` cannot shut down or start up Oracle ASM.

- An OSDBA group (`dba2`). During installation, you identify the group `dba2` as the OSDBA group for the database installed by the user `oracle2`. Members of `dba2` are granted the SYSDBA privileges for the Oracle Database DB2. Users who connect as SYSDBA are identified as user SYS on DB2.

- An OSBACKUPDBA group (`backupdba2`). During installation, you identify the group `backupdba2` as the OSDBA group for the database installed by the user `oracle2`. Members of `backupdba2` are granted the SYSBACKUP privileges for the database installed by the user `oracle2` to back up the database.

- An OSDGDBA group (`dgdba2`). During installation, you identify the group `dgdba2` as the OSDGDBA group for the database installed by the user `oracle2`. Members of `dgdba2` are granted the SYSDG privileges to administer Oracle Data Guard for the database installed by the user `oracle2`.

- An OSKMDBA group (`kmdba2`). During installation, you identify the group `kmdba2` as the OSKMDBA group for the database installed by the user `oracle2`. Members of `kmdba2` are granted the SYSKM privileges to administer encryption keys for the database installed by the user `oracle2`.

- An OSOPER group (`oper2`). During installation, you identify the group `oper2` as the OSOPER group for the database installed by the user `oracle2`. Members of `oper2` are granted the SYSOPER privileges (a limited set of the SYSDBA privileges), including the right to start up and shut down the DB2 database. Users who connect as OSOPER privileges are identified as user PUBLIC on DB2.

- An Oracle base `/u01/app/oracle2` owned by `oracle2:oinstall` with 775 permissions. The user `oracle2` has permissions to install software in this directory, but in no other directory in the `/u01/app` path.

Configuring Grid Infrastructure Software Owner User Environments

Understand the software owner user environments to configure before installing Oracle Grid Infrastructure.

You run the installer software with the Oracle Grid Infrastructure installation owner user account (`oracle` or `grid`). However, before you start the installer, you must configure the environment of the installation owner user account. If needed, you must also create other required Oracle software owners.

- Environment Requirements for Oracle Software Owners
  Review these requirements for Oracle software owner environments.

- Procedure for Configuring Oracle Software Owner Environments
  Configure each Oracle installation owner user account environment:

- Checking Resource Limits for Oracle Software Installation Users
  For each installation software owner user account, check the resource limits for installation.

- Setting Remote Display and X11 Forwarding Configuration
  If you are on a remote terminal, and the local system has only one visual (which is typical), then use the following syntax to set your user account `DISPLAY` environment variable:

- Preventing Installation Errors Caused by Terminal Output Commands
  During an Oracle Grid Infrastructure installation, OUI uses SSH to run commands and copy files to the other nodes. During the installation, hidden files on the system (for example, `.bashrc` or `.cshrc`) can cause `makefile` and other installation errors if they contain terminal output commands.

Environment Requirements for Oracle Software Owners

Review these requirements for Oracle software owner environments.

You must make the following changes to configure Oracle software owner environments:

- Set the installation software owner user (`grid`, `oracle`) default file mode creation mask (umask) to 022 in the shell startup file. Setting the mask to 022 ensures that the user performing the software installation creates files with 644 permissions.

- Set ulimit settings for file descriptors and processes for the installation software owner (`grid`, `oracle`).

- Set the DISPLAY environment variable in preparation for running an Oracle Universal Installer (OUI) installation.

> ⚠️ **Caution:**
>
> If you have existing Oracle installations that you installed with the user ID that is your Oracle Grid Infrastructure software owner, then unset all Oracle environment variable settings for that user.

Procedure for Configuring Oracle Software Owner Environments

Configure each Oracle installation owner user account environment:

1. Start an X terminal session (`xterm`) on the server where you are running the installation.

2. Enter the following command to ensure that X Window applications can display on this system, where hostname is the fully qualified name of the local host from which you are accessing the server:

   `$ xhost + hostname`

3. If you are not logged in as the software owner user, then switch to the software owner user you are configuring. For example, with the user `grid`:

   `$ su - grid`

On systems where you cannot run `su` commands, use `sudo` instead:

```
$ sudo -u grid -s
```

4. To determine the default shell for the user, enter the following command:

```
$ echo $SHELL
```

5. Open the user's shell startup file in any text editor:
   - Bash shell (`bash`):

     ```
     $ vi .bash_profile
     ```

   - Bourne shell (`sh`) or Korn shell (`ksh`):

     ```
     $ vi .profile
     ```

   - C shell (`csh` or `tcsh`):

     ```
     % vi .login
     ```

6. Enter or edit the following line, specifying a value of `022` for the default file mode creation mask:

```
umask 022
```

7. If the `ORACLE_SID`, `ORACLE_HOME`, or `ORACLE_BASE` environment variables are set in the file, then remove these lines from the file.

8. Save the file, and exit from the text editor.

9. To run the shell startup script, enter one of the following commands:
   - Bash shell:

     ```
     $ . ./.bash_profile
     ```

   - Bourne, Bash, or Korn shell:

     ```
     $ . ./.profile
     ```

   - C shell:

     ```
     % source ./.login
     ```

10. Use the following command to check the PATH environment variable:

```
$ echo $PATH
```

Remove any Oracle environment variables.

11. Unset any Oracle environment variables.

If you have an existing Oracle software installation, and you are using the same user to install this installation, then unset the $ORACLE_HOME, $ORA_NLS10, and $TNS_ADMIN environment variables.

If you have set $ORA_CRS_HOME as an environment variable, then unset it before starting an installation or upgrade. Do not use $ORA_CRS_HOME as a user environment variable, except as directed by Oracle Support.

12. If you are not installing the software on the local system, then enter a command similar to the following to direct X applications to display on the local system:

- Bourne, Bash, or Korn shell:

```
$ export DISPLAY=local_host:0.0
```

- C shell:

```
% setenv DISPLAY local_host:0.0
```

In this example, `local_host` is the host name or IP address of the system (your workstation, or another client) on which you want to display the installer.

13. If the `/tmp` directory has less than 1 GB of free space, then identify a file system with at least 1 GB of free space and set the `TMP` and `TMPDIR` environment variables to specify a temporary directory on this file system:

> **Note:**
>
> You cannot use a shared file system as the location of the temporary file directory (typically `/tmp`) for Oracle RAC installations. If you place `/tmp` on a shared file system, then the installation fails.

a. Use the `df -h` command to identify a suitable file system with sufficient free space.

b. If necessary, enter commands similar to the following to create a temporary directory on the file system that you identified, and set the appropriate permissions on the directory:

```
$ sudo -s
# mkdir /mount_point/tmp
# chmod 775 /mount_point/tmp
# exit
```

c. Enter commands similar to the following to set the `TMP` and `TMPDIR` environment variables:

Bourne, Bash, or Korn shell:

```
$ TMP=/mount_point/tmp
$ TMPDIR=/mount_point/tmp
$ export TMP TMPDIR
```

C shell:

```
% setenv TMP /mount_point/tmp
% setenv TMPDIR /mount_point/tmp
```

14. To verify that the environment has been set correctly, enter the following commands:

```
$ umask
$ env | more
```

Verify that the `umask` command displays a value of `22`, `022`, or `0022` and that the environment variables you set in this section have the correct values.

Checking Resource Limits for Oracle Software Installation Users

For each installation software owner user account, check the resource limits for installation.

On Oracle Linux systems, Oracle recommends that you install Oracle Database Preinstallation RPMs to meet preinstallation requirements like configuring your operating system to set the resource limits in the `limits.conf` file. Oracle Database Preinstallation RPM only configures the `limits.conf` file for the `oracle` user. If you are implementing Oracle Grid Infrastructure job role separation, then copy the values from the `oracle` user to the `grid` user in the `limits.conf` file.

Use the following ranges as guidelines for resource allocation to Oracle installation owners:

Table 5-1   Installation Owner Resource Limit Recommended Ranges

Resource Shell Limit	Resource	Soft Limit	Hard Limit
Open file descriptors	nofile	at least 1024	at least 65536
Number of processes available to a single user	nproc	at least 2047	at least 16384
Size of the stack segment of the process	stack	at least 10240 KB	at least 10240 KB, and at most 32768 KB
Maximum locked memory limit	memlock	at least 90 percent of the current RAM when HugePages memory is enabled and at least 3145728 KB (3 GB) when HugePages memory is disabled	at least 90 percent of the current RAM when HugePages memory is enabled and at least 3145728 KB (3 GB) when HugePages memory is disabled

To check resource limits:

1. Log in as an installation owner.

2. Check the soft and hard limits for the file descriptor setting. Ensure that the result is in the recommended range. For example:

```
$ ulimit -Sn
1024
$ ulimit -Hn
65536
```

3. Check the soft and hard limits for the number of processes available to a user. Ensure that the result is in the recommended range. For example:

```
$ ulimit -Su
2047
$ ulimit -Hu
16384
```

4. Check the soft limit for the stack setting. Ensure that the result is in the recommended range. For example:

```
$ ulimit -Ss
10240
$ ulimit -Hs
32768
```

5. Repeat this procedure for each Oracle software installation owner.

If necessary, update the resource limits in the `/etc/security/limits.conf` configuration file for the installation owner. However, the configuration file may be distribution specific. Contact your system administrator for distribution specific configuration file information.

> **✎ Note:**
>
> If you make changes to an Oracle installation user account and that user account is logged in, then changes to the `limits.conf` file do not take effect until you log these users out and log them back in. You must do this before you use these accounts for installation.

Setting Remote Display and X11 Forwarding Configuration

If you are on a remote terminal, and the local system has only one visual (which is typical), then use the following syntax to set your user account `DISPLAY` environment variable:

Remote Display

```
Bourne, Korn, and Bash shells
```

```
$ export DISPLAY=hostname:0
```

```
C shell
```

```
% setenv DISPLAY hostname:0
```

For example, if you are using the Bash shell and if your host name is `local_host`, then enter the following command:

```
$ export DISPLAY=node1:0
```

X11 Forwarding

To ensure that X11 forwarding does not cause the installation to fail, use the following procedure to create a user-level SSH client configuration file for Oracle installation owner user accounts:

1. Using any text editor, edit or create the software installation owner's `~/.ssh/config` file.

2. Ensure that the `ForwardX11` attribute in the `~/.ssh/config` file is set to `no`. For example:

```
Host *
    ForwardX11 no
```

3. Ensure that the permissions on `~/.ssh` are secured to the Oracle installation owner user account. For example:

```
$ ls -al .ssh
total 28
drwx------  2 grid oinstall 4096 Jun 21 2020
drwx------ 19 grid oinstall 4096 Jun 21 2020
-rw-r--r--  1 grid oinstall 1202 Jun 21 2020 authorized_keys
-rwx------  1 grid oinstall  668 Jun 21 2020 id_dsa
-rwx------  1 grid oinstall  601 Jun 21 2020 id_dsa.pub
-rwx------  1 grid oinstall 1610 Jun 21 2020 known_hosts
```

Preventing Installation Errors Caused by Terminal Output Commands

During an Oracle Grid Infrastructure installation, OUI uses SSH to run commands and copy files to the other nodes. During the installation, hidden files on the system (for example, `.bashrc` or `.cshrc`) can cause `makefile` and other installation errors if they contain terminal output commands.

To avoid this problem, you must modify hidden files in each Oracle installation owner user home directory to suppress all output on STDOUT or STDERR (for example, `stty`, `xtitle`, and other such commands) as in the following examples:

Bourne, Bash, or Korn shell:

```
if [ -t 0 ]; then
    stty intr ^C
fi
```

C shell:

```
test -t 0
if ($status == 0) then
    stty intr ^C
endif
```

> **Note:**
>
> If the remote shell can load hidden files that contain `stty` commands, then OUI indicates an error and stops the installation.

Enabling Intelligent Platform Management Interface (IPMI)

Intelligent Platform Management Interface (IPMI) provides a set of common interfaces to computer hardware and firmware that system administrators can use to monitor system health and manage the system.

Oracle Clusterware can integrate IPMI to provide failure isolation support and to ensure cluster integrity. You can configure node-termination with IPMI during installation by selecting IPMI from the Failure Isolation Support screen. You can also configure IPMI after installation with `crsctl` commands.

- Requirements for Enabling IPMI
  You must have the following hardware and software configured to enable cluster nodes to be managed with IPMI:

- Configuring the IPMI Management Network
  You can configure the BMC for DHCP, or for static IP addresses. Oracle recommends that you configure the BMC for dynamic IP address assignment using DHCP. To use this option, you must have a DHCP server configured to assign the BMC IP addresses.

- Configuring the BMC
  On each node, complete the following steps to configure the BMC to support IPMI-based node fencing.

- Configuring the iLO Processor on HP-UX
  Review this procedure to configure IPMI in the iLO processor on HP platforms, and to set a password for the null (noname user).

Related Topics

- *Oracle Clusterware Administration and Deployment Guide*

Requirements for Enabling IPMI

You must have the following hardware and software configured to enable cluster nodes to be managed with IPMI:

- Each cluster member node requires a Baseboard Management Controller (BMC) running firmware compatible with IPMI version 1.5 or greater, which supports IPMI over LANs, and configured for remote control using LAN.

- The cluster requires a management network for IPMI. This can be a shared network, but Oracle recommends that you configure a dedicated network.

- Each cluster member node's port used by BMC must be connected to the IPMI management network.

- Each cluster member must be connected to the management network.

- Some server platforms put their network interfaces into a power saving mode when they are powered off. In this case, they may operate only at a lower link

speed (for example, 100 MB, instead of 1 GB). For these platforms, the network switch port to which the BMC is connected must be able to auto-negotiate down to the lower speed, or IPMI will not function properly.

> **Note:**
>
> IPMI operates on the physical hardware platform through the network interface of the baseboard management controller (BMC). Depending on your system configuration, an IPMI-initiated restart of a server can affect all virtual environments hosted on the server. Contact your hardware and OS vendor for more information.

Configuring the IPMI Management Network

You can configure the BMC for DHCP, or for static IP addresses. Oracle recommends that you configure the BMC for dynamic IP address assignment using DHCP. To use this option, you must have a DHCP server configured to assign the BMC IP addresses.

For Oracle Clusterware to communicate with the BMC, the IPMI driver must be installed permanently on each node, so that it is available on system restarts. The IPMI driver is available on the Asianux Linux, Oracle Linux, Red Hat Enterprise Linux, and SUSE Linux Enterprise Server distributions supported with this release.

Configuring the BMC with dynamic addresses (DHCP) is not supported on HP-UX and AIX platforms.

> **Note:**
>
> If you configure IPMI, and you use Grid Naming Service (GNS) you still must configure separate addresses for the IPMI interfaces. As the IPMI adapter is not seen directly by the host, the IPMI adapter is not visible to GNS as an address on the host.

Configuring the BMC

On each node, complete the following steps to configure the BMC to support IPMI-based node fencing.

*   Configure a static IP address for the BMC.

*   Set a password for the null (noname) account for the BMC.

On HP-UX, the BMC and the iLO share network configuration. They have the same IP address, the same hardware MAC address, and the same default gateway address.

If your iLO processor is already configured for network access using a static address, then you have already established the required BMC network configuration. If you have not set up a static address for the iLO, then you must set a static address, and note the address so that you can enter it into the Oracle Clusterware local registry after installation.

You must also set up a password for the null (noname) user account, as this is the single account that the BMC uses. Administration accounts in the iLO are unrelated to the BMC. For security reasons, Oracle requires that you set a password for the BMC account.

Refer to your HP-UX documentation for more information about how to configure the BMC. Ensure that the BMC is configured on each cluster member node.

Configuring the iLO Processor on HP-UX

Review this procedure to configure IPMI in the iLO processor on HP platforms, and to set a password for the null (noname user).

1. Log in to the iLO web interface, and configure the required network settings for the iLO and BMC under Administration, then Network Settings to obtain the IP address, and to check the netmask and default gateway.

2. Start a terminal session from a device with network connectivity to the BMC of the node to configure.

3. From the terminal session, set the IPMI password for the anonymous user (noname) over the network, using an IPMI administration tool, such as ipmitool from a client or server connected to the note whose user password you want to change. F

    or example, where `ipmiaddr` is the IPMI address, and password is the password:

    ```
    % ipmitool -H ipmiaddr -U "" user set password 1 "password"
    ```

    In this example, the anonymous username is being provided explicitly using -U "" for clarity, but it is implied if the username argument is missing. This command prompts for the current password, which can be the initial null password. If the password is successfully changed, then the command prints an error similar to: `"Close session command failed."`. This message is printed because the command attempts to terminate the IPMI network session using the previous password.

4. After you configure the BMC on each cluster member node, and complete Oracle Grid Infrastructure installation, you must store the IPMI administrator credentials and the BMC static IP address in the Oracle Local Registry (OLR) on each cluster member node. Use crsctl to do this. However, when you store the IPMI credentials in the OLR, you must have the anonymous user specified explicitly, or a parsing error will be reported:

    ```
    % crsctl set css ipmiadmin ""
    ```

    When prompted, provide the password you set with the IPMI administrator tool.

Granting MLOCK Privilege to OSDBA

Grant the MLOCK privilege to `dba` to avoid Oracle Database installation errors.

The asynchronous Input-Output pseudo-driver on HP-UX enables Oracle Database to perform Input-Output to raw disk partitions using an asynchronous method, resulting in less Input-Output overhead and higher throughput. To permit Oracle Database to process asynchronous Input-Output operations, assign the MLOCK privilege to the OSDBA group (`dba`) before you start the Oracle Database installation.

To assign the MLOCK privilege:

1. Log in as the `root` user.

2. Using any text editor, open the `/etc/privgroup` file, or create it, if necessary.

3. Add or edit the following line, which begins with the name of the OSDBA group, specifying the privilege MLOCK:

```
dba RTPRIO RTSCHED MLOCK
```

> **Note:**
>
> You must use only one line to specify the privileges for a particular group in this file. If the file already contains a line for the `dba` group, then add the MLOCK privilege in the same line.

4. Save the file and quit the text editor.

5. Enter the following command to grant the privileges to the OSDBA group:

```
# /usr/sbin/setprivgrp -f /etc/privgroup
```

6. Enter the following command to verify if the privileges are set correctly:

```
# /usr/bin/getprivgrp dba
```

6

Supported Storage Options for Oracle Database and Oracle Grid Infrastructure

Review supported storage options as part of your installation planning process.

- **Supported Storage Options for Oracle Grid Infrastructure**
  The following table shows the storage options supported for Oracle Grid Infrastructure binaries and files:

- **Storage Considerations for Oracle Grid Infrastructure and Oracle RAC**
  For all installations, you must choose the storage option to use for Oracle Grid Infrastructure (Oracle Clusterware and Oracle ASM), and Oracle Real Application Clusters (Oracle RAC) databases.

- **Guidelines for Using Oracle ASM Disk Groups for Storage**
  Plan how you want to configure Oracle ASM disk groups for deployment.

- **Guidelines for Configuring Oracle ASM Disk Groups on NFS**
  Configuration guidelines for Automatic Storage Management (Oracle ASM) on NFS file systems.

- **Using Logical Volume Managers with Oracle Grid Infrastructure and Oracle RAC**
  Oracle Grid Infrastructure and Oracle RAC only support cluster-aware volume managers.

- **Using a Cluster File System for Oracle Clusterware Files**
  Starting with Oracle Grid Infrastructure 19c, you can use Oracle Automatic Storage Management (Oracle ASM) or certified shared file system to store OCR files and voting files.

- **About NFS Storage for Data Files**
  Review this section for NFS storage configuration guidelines.

- **About Direct NFS Client Mounts to NFS Storage Devices**
  Direct NFS Client integrates the NFS client functionality directly in the Oracle software to optimize the I/O path between Oracle and the NFS server. This integration can provide significant performance improvements.

Supported Storage Options for Oracle Grid Infrastructure

The following table shows the storage options supported for Oracle Grid Infrastructure binaries and files:

Table 6-1   Supported Storage Options for Oracle Grid Infrastructure

Storage Option	OCR and Voting Files	Oracle Clusterware Binaries	Oracle RAC Database Binaries	Oracle RAC Database Data Files	Oracle RAC Database Recovery Files
Oracle Automatic Storage Management (Oracle ASM) **Note**: Loopback devices are not supported for use with Oracle ASM	Yes	No	No	Yes	Yes
Local file system	No	Yes	Yes	No	No
Network file system (NFS) on a certified network-attached storage (NAS) filer **Note:** Direct NFS Client does not support Oracle Clusterware files	Yes	Yes	Yes	Yes	Yes
Direct-attached storage (DAS)	No	No	Yes	Yes	Yes
Shared disk partitions (block devices or raw devices)	No	No	No	No	No

Guidelines for Storage Options

Use the following guidelines when choosing storage options:

- You can choose any combination of the supported storage options for each file type provided that you satisfy all requirements listed for the chosen storage options.

- You can use Oracle ASM or shared file system to store Oracle Clusterware files.

- Direct use of raw or block devices is not supported. You can only use raw or block devices under Oracle ASM.

Related Topics

- *Oracle Database Upgrade Guide*

Storage Considerations for Oracle Grid Infrastructure and Oracle RAC

For all installations, you must choose the storage option to use for Oracle Grid Infrastructure (Oracle Clusterware and Oracle ASM), and Oracle Real Application Clusters (Oracle RAC) databases.

Storage Considerations for Oracle Clusterware

Oracle Clusterware voting files are used to monitor cluster node status, and Oracle Cluster Registry (OCR) files contain configuration information about the cluster. You can store Oracle Cluster Registry (OCR) and voting files in Oracle ASM disk groups or a shared file system. You can also store a backup of the OCR file in a disk group. Storage must be shared; any node that does not have access to an absolute majority of voting files (more than half) is restarted.

If you use Oracle ASM disk groups created on Network File System (NFS) for storage, then ensure that you follow the recommendations for mounting NFS described in the topic *Guidelines for Configuring Oracle ASM Disk Groups on NFS*.

Storage Considerations for Oracle RAC

Oracle ASM is a supported storage option for database and recovery files. For all installations, Oracle recommends that you create at least two separate Oracle ASM disk groups: One for Oracle Database data files, and one for recovery files. Oracle recommends that you place the Oracle Database disk group and the recovery files disk group in separate failure groups.

- If you do not use Oracle ASM for database files, then Oracle recommends that you place the data files and the Fast Recovery Area in shared storage located outside of the Oracle home, in separate locations, so that a hardware failure does not affect availability.

- You can choose any combination of the supported storage options for each file type provided that you satisfy all requirements listed for the chosen storage options.

- To use Oracle ASM with Oracle RAC, and if you are configuring a new Oracle ASM instance, then your system must meet the following conditions:

  – All nodes on the cluster have Oracle Clusterware and Oracle ASM 21c installed as part of an Oracle Grid Infrastructure for a cluster installation.

  – Any existing Oracle ASM instance on any node in the cluster is shut down.

  – To provide voting file redundancy, one Oracle ASM disk group is sufficient. The Oracle ASM disk group provides three or five copies.

You can use NFS, with or without Direct NFS, to store Oracle Database data files. You cannot use NFS as storage for Oracle Clusterware files.

Guidelines for Using Oracle ASM Disk Groups for Storage

Plan how you want to configure Oracle ASM disk groups for deployment.

During Oracle Grid Infrastructure installation, you can create one or two Oracle ASM disk groups. After the Oracle Grid Infrastructure installation, you can create additional disk groups using Oracle Automatic Storage Management Configuration Assistant (ASMCA), SQL*Plus, or Automatic Storage Management Command-Line Utility (ASMCMD).

Choose to create a second disk group during Oracle Grid Infrastructure installation. The first disk group stores the Oracle Cluster Registry (OCR), voting files, and the Oracle ASM password file. The second disk group stores the Grid Infrastructure Management Repository (GIMR) data files and Oracle Cluster Registry (OCR) backup files. Oracle strongly recommends that you store the OCR backup files in a different disk group from the disk group where you store OCR files. In addition, having a second disk group for GIMR is advisable for performance, availability, sizing, and manageability of storage.

> **✏ Note:**
>
> You must specify the Grid Infrastructure Management Repository (GIMR) location at the time of installing Oracle Grid Infrastructure. You cannot migrate the GIMR from one disk group to another later.

If you install Oracle Database or Oracle RAC after you install Oracle Grid Infrastructure, then you can either use the same disk group for database files, OCR, and voting files, or you can use different disk groups. If you create multiple disk groups before installing Oracle RAC or before creating a database, then you can do one of the following:

- Place the data files in the same disk group as the Oracle Clusterware files.

- Use the same Oracle ASM disk group for data files and recovery files.

- Use different disk groups for each file type.

If you create only one disk group for storage, then the OCR and voting files, database files, and recovery files are contained in the one disk group. If you create multiple disk groups for storage, then you can place files in different disk groups.

> **✏ See Also:**
>
> *Oracle Automatic Storage Management Administrator's Guide* for information about creating disk groups

Guidelines for Configuring Oracle ASM Disk Groups on NFS

Configuration guidelines for Automatic Storage Management (Oracle ASM) on NFS file systems.

You can create Oracle ASM disk groups on block devices or NFS (Network File System) on a supported Network Attached Storage (NAS) device. Understand the following guidelines for configuring Oracle ASM disk groups on NFS:

> **✏ Note:**
>
> All storage products must be supported by both your server and storage vendors.

Guidelines for Deploying Oracle ASM Disk Groups Without Quorum Disks

- To use an NFS file system, it must be on a supported NAS device. Log in to My Oracle Support at the following URL, and click Certifications to find the most current information about supported NAS devices:

  https://support.oracle.com/

- NFS file systems must be mounted and available over NFS mounts before you start installation. Refer to your vendor documentation to complete NFS configuration and mounting.

- Direct NFS requires hard mounts. Hard mounting NFS filers prevents corruption which could occur if the client connection were to time out. If an NFS filer hangs on an I/O operation to a mirrored file, then the database and Oracle ASM cannot failover to the surviving mirror copy. Therefore, Oracle recommends that you use external redundancy when you deploy Oracle ASM disk groups on NFS storage.

- The performance of Oracle software and databases stored on Oracle ASM disk groups on NFS depends on the performance of the network connection between the Oracle server and the NAS device. Oracle recommends that you connect the server to the NAS device using a private dedicated network connection, which should be Gigabit Ethernet or better.

- You can configure Oracle ASM on NFS when you deploy an Oracle Standalone Cluster configuration.

- You can specify separate NFS locations for Oracle ASM disk groups for Oracle Clusterware files and the Grid Infrastructure Management Repository (GIMR).

- The user account with which you perform the installation (`oracle` or `grid`) must have write permissions to create the files in the path that you specify.

Guidelines for Deploying Oracle ASM Disk Groups With Quorum Disks

- SAN-attached storage or iSCSI-attached devices are the preferred ways to connect to quorum disks. If your standard deployment requires NFS to be used as storage, then use soft mounts for NFS-based Oracle ASM quorum disks and hard mounts for other Oracle ASM disks.

- You can use Direct NFS (dNFS) for storage of Oracle Database data files. dNFS does not support soft mounts, so you cannot use dNFS for quorum failure groups. Alternatively, use kernel-based NFS with a soft mount for NFS storage residing in a quorum failure group.

- The quorum failure group feature in Oracle ASM enables use of NFS storage in an Oracle ASM disk group without requiring a hard mount for NFS storage in the quorum failure group. This capability is useful for Oracle Extended Clusters where a third site is required for establishing quorum.

Using Logical Volume Managers with Oracle Grid Infrastructure and Oracle RAC

Oracle Grid Infrastructure and Oracle RAC only support cluster-aware volume managers.

Oracle Grid Infrastructure and Oracle RAC only support cluster-aware volume managers. Some third-party volume managers are not cluster-aware, and so are not supported. To confirm that a volume manager you want to use is supported, click **Certifications** on My Oracle Support to determine if your volume manager is certified for Oracle RAC. My Oracle Support is available at the following URL:

https://support.oracle.com

Using a Cluster File System for Oracle Clusterware Files

Starting with Oracle Grid Infrastructure 19c, you can use Oracle Automatic Storage Management (Oracle ASM) or certified shared file system to store OCR files and voting files.

For new Oracle Standalone Cluster installations, you can use Oracle ASM or shared file system to store voting files and OCR files. For other cluster types, you must use Oracle Automatic Storage Management (Oracle ASM) to store voting files and OCR files. For Linux 86-64 (64-bit) and Linux Itanium platforms, Oracle provides a cluster file system, OCFS2. However, Oracle does not recommend using OCFS2 for Oracle Clusterware files.

About NFS Storage for Data Files

Review this section for NFS storage configuration guidelines.

Network-Attached Storage and NFS Protocol

Network-attached storage (NAS) systems use the network file system (NFS) protocol to to access files over a network, which enables client servers to access files over networks as easily as to storage devices attached directly to the servers. You can store data files on supported NFS systems. NFS is a shared file system protocol, so NFS can support both single instance and Oracle Real Application Clusters databases.

> **Note:**
>
> The performance of Oracle software and databases stored on NAS devices depends on the performance of the network connection between the servers and the network-attached storage devices.For better performance, Oracle recommends that you connect servers to NAS devices using private dedicated network connections. NFS network connections should use Gigabit Ethernet or better.

Refer to your vendor documentation to complete NFS configuration and mounting.

Requirements for Using NFS Storage

Before you start installation, NFS file systems must be mounted and available to servers.

About Direct NFS Client Mounts to NFS Storage Devices

Direct NFS Client integrates the NFS client functionality directly in the Oracle software to optimize the I/O path between Oracle and the NFS server. This integration can provide significant performance improvements.

Direct NFS Client supports NFSv3, NFSv4, NFSv4.1, and pNFS protocols to access the NFS server. Direct NFS Client also simplifies, and in many cases automates, the performance optimization of the NFS client configuration for database workloads.

Starting with Oracle Database 12c Release 2, when you enable Direct NFS, you can also enable the Direct NFS dispatcher. The Direct NFS dispatcher consolidates the number of TCP connections that are created from a database instance to the NFS server. In large database deployments, using Direct NFS dispatcher improves scalability and network performance. Parallel NFS deployments also require a large number of connections. Hence, the Direct NFS dispatcher is recommended with Parallel NFS deployments too.

Direct NFS Client can obtain NFS mount points either from the operating system mount entries, or from the `oranfstab` file.

Direct NFS Client Requirements

- NFS servers must have write size values (`wtmax`) of 32768 or greater to work with Direct NFS Client.

- NFS mount points must be mounted both by the operating system kernel NFS client and Direct NFS Client, even though you configure Direct NFS Client to provide file service.

  If Oracle Database cannot connect to an NFS server using Direct NFS Client, then Oracle Database connects to the NFS server using the operating system kernel NFS client. When Oracle Database fails to connect to NAS storage though Direct NFS Client, it logs an informational message about the Direct NFS Client connect error in the Oracle alert and trace files.

- Follow standard guidelines for maintaining integrity of Oracle Database files mounted by both operating system NFS and by Direct NFS Client.

Direct NFS Mount Point Search Order

Direct NFS Client searches for mount entries in the following order:

1. `ORACLE_BASE_HOME/dbs/oranfstab`

2. `ORACLE_BASE_CONFIG/dbs/oranfstab`

3. `/var/opt/oracle/oranfstab`

4. `/etc/mnttab`

Direct NFS Client uses the first matching entry as the mount point.

> **✏ Note:**
>
> You can have only one active NFS Client implementation for each instance. Enabling Direct NFS Client on an instance prevents you from using another NFS Client implementation, such as kernel NFS Client.

- To print the ORACLE_BASE_HOME path, run:

```
$ setenv ORACLE_HOME /u01/app/oracle/product/21.0.0/dbhome_1
$ cd $ORACLE_HOME/bin
$ ./orabasehome
```

- To print the ORACLE_BASE_CONFIG path, run:

```
$ setenv ORACLE_HOME /u01/app/oracle/product/21.0.0/dbhome_1
$ cd $ORACLE_HOME/bin
$ ./orabaseconfig
```

> **✏ See Also:**
>
> - *Oracle Database Reference* for information about setting the `enable_dnfs_dispatcher` parameter in the initialization parameter file to enable Direct NFS dispatcher
>
> - *Oracle Database Performance Tuning Guide* for performance benefits of enabling Parallel NFS and Direct NFS dispatcher
>
> - *Oracle Automatic Storage Management Administrator's Guide* for guidelines about managing Oracle Database data files created with Direct NFS Client or kernel NFS

ORACLE®

7

Configuring Storage for Oracle Grid Infrastructure

Complete these procedures to configure Oracle Automatic Storage Management (Oracle ASM) for Oracle Grid Infrastructure for a cluster.

Oracle Grid Infrastructure for a cluster provides system support for Oracle Database. Oracle ASM is a volume manager and a file system for Oracle database files that supports single-instance Oracle Database and Oracle Real Application Clusters (Oracle RAC) configurations. Oracle Automatic Storage Management also supports a general purpose file system for your application needs, including Oracle Database binaries. Oracle Automatic Storage Management is Oracle's recommended storage management solution. It provides an alternative to conventional volume managers and file systems.

> **Note:**
>
> Oracle ASM and shared file system are the supported storage management solutions for Oracle Cluster Registry (OCR) and Oracle Clusterware voting files. The OCR is a file that contains the configuration information and status of the cluster. The installer automatically initializes the OCR during the Oracle Clusterware installation. Database Configuration Assistant uses the OCR for storing the configurations for the cluster databases that it creates.

- **Configuring Storage for Oracle Automatic Storage Management**
  Identify storage requirements and Oracle ASM disk group options.

- **Using Disk Groups with Oracle Database Files on Oracle ASM**
  Review this information to configure Oracle Automatic Storage Management (Oracle ASM) storage for Oracle Clusterware and Oracle Database Files.

- **Configuring File System Storage for Oracle Database**
  Complete these procedures to use file system storage for Oracle Database.

- **Creating and Using Oracle ASM Credentials File**
  Review this information to create an Oracle ASM credentials file.

- **Creating a Remote GIMR Client Data File for Oracle Standalone Clusters**
  Create a remote GIMR client data file to specify the Oracle Standalone Cluster configuration for the remote Grid Infrastructure Management Repository (GIMR).

Configuring Storage for Oracle Automatic Storage Management

Identify storage requirements and Oracle ASM disk group options.

- **Identifying Storage Requirements for Oracle Automatic Storage Management**
  To identify the storage requirements for using Oracle ASM, you must determine the number of devices and the amount of free disk space that you require.

- **Oracle Clusterware Storage Space Requirements**
  Use this information to determine the minimum number of disks and the minimum disk space requirements based on the redundancy type, for installing Oracle Clusterware files for Oracle Standalone Cluster deployments.

- **About the Grid Infrastructure Management Repository**
  Every Oracle Domain Services Cluster contains a Grid Infrastructure Management Repository (GIMR), but GIMR configuration is optional for Oracle Standalone Cluster.

- **Using an Existing Oracle ASM Disk Group**
  Use Oracle Enterprise Manager Cloud Control or the Oracle ASM command line tool (`asmcmd`) to identify existing disk groups, and to determine if sufficient space is available in the disk group.

- **About Upgrading Existing Oracle Automatic Storage Management Instances**
  Oracle Automatic Storage Management (Oracle ASM) upgrades are carried out during an Oracle Grid Infrastructure upgrade.

- **Selecting Disks to use with Oracle ASM Disk Groups**
  If you are sure that a suitable disk group does not exist on the system, then install or identify appropriate disk devices to add to a new disk group.

- **Specifying the Oracle ASM Disk Discovery String**
  When an Oracle ASM instance is initialized, Oracle ASM discovers and examines the contents of all of the disks that are in the paths that you designated with values in the `ASM_DISKSTRING` initialization parameter.

- **Creating Files on a NAS Device for Use with Oracle Automatic Storage Management**
  If you have a certified NAS storage device, then you can create zero-padded files in an NFS mounted directory and use those files as disk devices in an Oracle ASM disk group.

Related Topics

- *Oracle Automatic Storage Management Administrator's Guide*

Identifying Storage Requirements for Oracle Automatic Storage Management

To identify the storage requirements for using Oracle ASM, you must determine the number of devices and the amount of free disk space that you require.

To complete this task, follow these steps:

1. Determine whether you want to use Oracle ASM for Oracle Database files, recovery files, and Oracle Database binaries. Oracle Database files include data files, control files, redo log files, the server parameter file, and the password file.

> **Note:**
>
> - You do not have to use the same storage mechanism for Oracle Database files and recovery files. You can use a shared file system for one file type and Oracle ASM for the other.
>
> - There are two types of Oracle Clusterware files: OCR files and voting files. You must use Oracle ASM to store OCR and voting files.
>
> - If your database files are stored on a shared file system, then you can continue to use the same for database files, instead of moving them to Oracle ASM storage.

2. Choose the Oracle ASM redundancy level to use for the Oracle ASM disk group.

Except when using external redundancy, Oracle ASM mirrors all Oracle Clusterware files in separate failure groups within a disk group. A quorum failure group, a special type of failure group, contains mirror copies of voting files when voting files are stored in normal or high redundancy disk groups. The disk groups that contain Oracle Clusterware files (OCR and voting files) have a higher minimum number of failure groups than other disk groups because the voting files are stored in quorum failure groups in the Oracle ASM disk group.

A quorum failure group is a special type of failure group that is used to store the Oracle Clusterware voting files. The quorum failure group is used to ensure that a quorum of the specified failure groups are available. When Oracle ASM mounts a disk group that contains Oracle Clusterware files, the quorum failure group is used to determine if the disk group can be mounted in the event of the loss of one or more failure groups. Disks in the quorum failure group do not contain user data, therefore a quorum failure group is not considered when determining redundancy requirements in respect to storing user data.

The redundancy levels are as follows:

- High redundancy

  In a high redundancy disk group, Oracle ASM uses three-way mirroring to increase performance and provide the highest level of reliability. A high redundancy disk group requires a minimum of three disk devices (or three failure groups). The effective disk space in a high redundancy disk group is one-third the sum of the disk space in all of its devices.

  For Oracle Clusterware files, a high redundancy disk group requires a minimum of five disk devices and provides five voting files and one OCR (one primary and two secondary copies). For example, your deployment may consist of three regular failure groups and two quorum failure groups. Note that not all failure groups can be quorum failure groups, even though voting files need all five disks. With high redundancy, the cluster can survive the loss of two failure groups.

  While high redundancy disk groups do provide a high level of data protection, you should consider the greater cost of additional storage devices before deciding to select high redundancy disk groups.

- Normal redundancy

  In a normal redundancy disk group, to increase performance and reliability, Oracle ASM by default uses two-way mirroring. A normal redundancy disk group requires a minimum of two disk devices (or two failure groups). The effective disk space in a normal redundancy disk group is half the sum of the disk space in all of its devices.

For Oracle Clusterware files, a normal redundancy disk group requires a minimum of three disk devices and provides three voting files and one OCR (one primary and one secondary copy). For example, your deployment may consist of two regular failure groups and one quorum failure group. With normal redundancy, the cluster can survive the loss of one failure group.

If you are not using a storage array providing independent protection against data loss for storage, then Oracle recommends that you select normal redundancy.

- External redundancy

  An external redundancy disk group requires a minimum of one disk device. The effective disk space in an external redundancy disk group is the sum of the disk space in all of its devices.

  Because Oracle ASM does not mirror data in an external redundancy disk group, Oracle recommends that you use external redundancy with storage devices such as RAID, or other similar devices that provide their own data protection mechanisms.

- Flex redundancy

  A flex redundancy disk group is a type of redundancy disk group with features such as flexible file redundancy, mirror splitting, and redundancy change. A flex disk group can consolidate files with different redundancy requirements into a single disk group. It also provides the capability for databases to change the redundancy of its files. A disk group is a collection of file groups, each associated with one database. A quota group defines the maximum storage space or quota limit of a group of databases within a disk group.

  In a flex redundancy disk group, Oracle ASM uses three-way mirroring of Oracle ASM metadata to increase performance and provide reliability. For database data, you can choose no mirroring (unprotected), two-way mirroring (mirrored), or three-way mirroring (high). A flex redundancy disk group requires a minimum of three disk devices (or three failure groups).

  > **See Also:**
  >
  > *Oracle Automatic Storage Management Administrator's Guide* for more information about file groups and quota groups for flex disk groups

  > **Note:**
  >
  > You can alter the redundancy level of the disk group after a disk group is created. For example, you can convert a normal or high redundancy disk group to a flex redundancy disk group. Within a flex redundancy disk group, file redundancy can change among three possible values: unprotected, mirrored, or high.

3. Determine the total amount of disk space that you require for Oracle Clusterware files, and for the database files and recovery files.

If an Oracle ASM instance is running on the system, then you can use an existing disk group to meet these storage requirements. If necessary, you can add disks to an existing disk group during the database installation.

See *Oracle Clusterware Storage Space Requirements* to determine the minimum number of disks and the minimum disk space requirements for installing Oracle Clusterware files, and installing the starter database, where you have voting files in a separate disk group.

4. Determine an allocation unit size.

   Every Oracle ASM disk is divided into allocation units (AU). An allocation unit is the fundamental unit of allocation within a disk group. You can select the AU Size value from 1, 2, 4, 8, 16, 32 or 64 MB, depending on the specific disk group compatibility level. For flex disk groups, the default value for AU size is set to 4 MB. For external, normal, and high redundancies, the default AU size is 1 MB.

5. For Oracle Clusterware installations, you must also add additional disk space for the Oracle ASM metadata. You can use the following formula to calculate the disk space requirements (in MB) for OCR and voting files, and the Oracle ASM metadata:

   ```
   total = [2 * ausize * disks] + [redundancy * (ausize *
   (all_client_instances + nodes + disks + 32) + (64 * nodes) + clients +
   543)]
   ```

   redundancy = Number of mirrors: external = 1, normal = 2, high = 3, flex = 3.

   ausize = Metadata AU size in megabytes

   all_client_instance = Sum of all database clients

   nodes = Number of nodes in cluster.

   clients - Number of database instances for each node.

   disks - Number of disks in disk group.

   For example, for a four-node Oracle RAC installation, using three disks in a normal redundancy disk group, you require an additional 5293 MB of space: `[2 * 4 * 3] + [2 * (4 * (4 * (4 + 1)+ 30)+ (64 * 4)+ 533)] = 5293 MB`

6. Optionally, identify failure groups for the Oracle ASM disk group devices.

   If you intend to use a normal or high redundancy disk group, then you can further protect the database against hardware failure by associating a set of disk devices in a custom failure group. By default, each device is included in its failure group. However, if two disk devices in a normal redundancy disk group are attached to the same Host Bus Adapter (HBA), then the disk group becomes unavailable if the adapter fails. The HBA in this example is a single point of failure.

   For instance, to avoid failures of this type, you can use two HBA fabric paths, each with two disks, and define a failure group for the disks attached to each adapter. This configuration would enable the disk group to tolerate the failure of one HBA fabric path.

> **Note:**
>
> You can define custom failure groups during installation of Oracle Grid Infrastructure. You can also define failure groups after installation using the GUI tool ASMCA, the command line tool asmcmd, or SQL commands. If you define custom failure groups, then you must specify a minimum of two failure groups for normal redundancy disk groups and three failure groups for high redundancy disk groups.

7. If you are sure that a suitable disk group does not exist on the system, then install or identify appropriate disk devices to add to a new disk group. Use the following guidelines when identifying appropriate disk devices:

   - The disk devices must be owned by the user performing Oracle Grid Infrastructure installation.

   - All the devices in an Oracle ASM disk group must be the same size and have the same performance characteristics.

   - Do not specify multiple partitions on a single physical disk as a disk group device. Oracle ASM expects each disk group device to be on a separate physical disk.

   - Although you can specify a logical volume as a device in an Oracle ASM disk group, Oracle does not recommend their use because it adds a layer of complexity that is unnecessary with Oracle ASM. Oracle recommends that if you choose to use a logical volume manager, then use the logical volume manager to represent a single logical unit number (LUN) without striping or mirroring, so that you can minimize the effect on storage performance of the additional storage layer.

8. If you use Oracle ASM disk groups created on Network File System (NFS) for storage, then ensure that you follow recommendations described in *Guidelines for Configuring Oracle ASM Disk Groups on NFS*.

Oracle Clusterware Storage Space Requirements

Use this information to determine the minimum number of disks and the minimum disk space requirements based on the redundancy type, for installing Oracle Clusterware files for Oracle Standalone Cluster deployments.

Total Oracle Clusterware Storage Space Required by Oracle Standalone Cluster

During installation of an Oracle Standalone Cluster, if you create the `MGMT` disk group for Grid Infrastructure Management Repository (GIMR), then the installer requires that you use a disk group with at least 35 GB of available space.

> **Note:**
>
> Starting with Oracle Grid Infrastructure 19c, configuring GIMR is optional for Oracle Standalone Cluster deployments. When upgrading to Oracle Grid Infrastructure 19c or a later release, a new GIMR is created only if the source Grid home has a GIMR configured.

Oracle Clusterware space requirements vary for different redundancy levels. The following tables list the space requirements for each redundancy level.

> **Note:**
>
> The `DATA` disk group stores OCR and voting files, and the `MGMT` disk group stores GIMR and Oracle Clusterware backup files.

Table 7-1   Minimum Available Space Requirements for Oracle Standalone Cluster With GIMR Configuration

Redundancy Level	`DATA` Disk Group	`MGMT` Disk Group	Total Storage
External	1 GB	28 GB Each node beyond four: 5 GB	29 GB
Normal	2 GB	56 GB Each node beyond four: 5 GB	58 GB

- Oracle recommends that you use a separate disk group, other than `DATA`, for GIMR and Oracle Clusterware backup files.
- The initial GIMR sizing for the Oracle Standalone Cluster is for up to four nodes. You must add additional storage space to the disk group containing the GIMR and Oracle Clusterware backup files for each new node added to the cluster.

Table 7-2   Minimum Available Space Requirements for Oracle Standalone Cluster Without GIMR Configuration

Redundancy Level	`DATA` Disk Group	Total Storage
External	1 GB	1 GB
Normal	2 GB	2 GB
High/Flex/Extended	3 GB	3 GB

- Oracle recommends that you use a separate disk group, other than `DATA`, for Oracle Clusterware backup files.
- The initial sizing for the Oracle Standalone Cluster is for up to four nodes. You must add additional storage space to the disk group containing Oracle Clusterware backup files for each new node added to the cluster.

About the Grid Infrastructure Management Repository

Every Oracle Domain Services Cluster contains a Grid Infrastructure Management Repository (GIMR), but GIMR configuration is optional for Oracle Standalone Cluster.

The Grid Infrastructure Management Repository (GIMR), or the data repository is a multitenant database with a pluggable database (PDB) for the GIMR of each cluster. The GIMR stores the following information about the cluster:

- Real time performance data that Cluster Health Monitor collects

- Fault, diagnosis, and metric data that Cluster Health Advisor collects

- Cluster-wide events about all resources that Oracle Clusterware collects

- Workload performance and CPU architecture data that Quality of Service Management (QoS) collects

- Metadata required for Oracle Fleet Patching and Provisioning

Starting with Oracle Grid Infrastructure 19c, configuring GIMR is optional for Oracle Standalone Cluster deployments. The Oracle Standalone Cluster locally hosts the GIMR on an Oracle ASM disk group or a shared file system; this GIMR is a multitenant database with a single pluggable database (PDB).

The global GIMR runs in an Oracle Domain Services Cluster. Oracle Domain Services Cluster locally hosts the GIMR in a separate Oracle ASM disk group. Oracle Standalone Cluster with remote GIMR uses the remote GIMR located on the Oracle Domain Services Cluster. Hosting the GIMR on a remote cluster reduces the overhead of running an extra infrastructure repository on a cluster. The GIMR for an Oracle Domain Services Cluster is a multitenant database with one PDB.

When you configure an Oracle Domain Services Cluster, the installer prompts to configure a separate Oracle ASM disk group for the GIMR, with the default name as `MGMT`.

Using an Existing Oracle ASM Disk Group

Use Oracle Enterprise Manager Cloud Control or the Oracle ASM command line tool (`asmcmd`) to identify existing disk groups, and to determine if sufficient space is available in the disk group.

1. Connect to the Oracle ASM instance and start the instance if necessary:

   ```
   $ $ORACLE_HOME/bin/asmcmd
   ASMCMD> startup
   ```

2. Enter one of the following commands to view the existing disk groups, their redundancy level, and the amount of free disk space in each one:

   ```
   ASMCMD> lsdg
   ```

   or

   ```
   $ORACLE_HOME/bin/asmcmd -p lsdg
   ```

   The `lsdg` command lists information about mounted disk groups only.

3. From the output, identify a disk group with the appropriate redundancy level and note the free space that it contains.

4. If necessary, install or identify the additional disk devices required to meet the storage requirements for your installation.

> **✎ Note:**
>
> If you are adding devices to an existing disk group, then Oracle recommends that you use devices that have the same size and performance characteristics as the existing devices in that disk group.

About Upgrading Existing Oracle Automatic Storage Management Instances

Oracle Automatic Storage Management (Oracle ASM) upgrades are carried out during an Oracle Grid Infrastructure upgrade.

If you are upgrading from Oracle ASM 12c Release 2 (12.2) or later, then Oracle ASM is always upgraded with Oracle Grid Infrastructure as part of the upgrade, and Oracle Automatic Storage Management Configuration Assistant (Oracle ASMCA) is started by the root scripts during upgrade. Subsequently, you can use Oracle ASMCA (located in `Grid_home/bin`) to configure failure groups and Oracle ASM volumes.

Oracle ASMCA cannot perform a separate upgrade of Oracle ASM from a prior release to the current release.

Upgrades of Oracle ASM from releases prior to 12c Release 2 (12.2) are not supported.

Related Topics

- *Oracle Automatic Storage Management Administrator's Guide*
- *Oracle Database Upgrade Guide*

Selecting Disks to use with Oracle ASM Disk Groups

If you are sure that a suitable disk group does not exist on the system, then install or identify appropriate disk devices to add to a new disk group.

Use the following guidelines when identifying appropriate disk devices:

- All of the devices in an Oracle ASM disk group should be the same size and have the same performance characteristics.

- Do not specify multiple partitions on a single physical disk as a disk group device. Oracle ASM expects each disk group device to be on a separate physical disk.

- Nonshared logical partitions are not supported with Oracle RAC. To use logical partitions for your Oracle RAC database, you must use shared logical volumes created by a logical volume manager such as `fdisk`.

- Although you can specify a logical volume as a device in an Oracle ASM disk group, Oracle does not recommend their use because it adds a layer of complexity that is unnecessary with Oracle ASM. In addition, Oracle RAC requires a cluster logical volume manager in case you decide to use a logical volume with Oracle ASM and Oracle RAC.

Specifying the Oracle ASM Disk Discovery String

When an Oracle ASM instance is initialized, Oracle ASM discovers and examines the contents of all of the disks that are in the paths that you designated with values in the `ASM_DISKSTRING` initialization parameter.

The value for the `ASM_DISKSTRING` initialization parameter is an operating system–dependent value that Oracle ASM uses to limit the set of paths that the discovery process uses to search for disks. The exact syntax of a discovery string depends on the platform and whether Oracle Exadata disks are used. The path names that an operating system accepts are always usable as discovery strings.

The default value of `ASM_DISKSTRING` might not find all disks in all situations. If your installation uses multipathing software, then the software might place pseudo-devices in a path that is different from the operating system default.

> **✎ See Also:**
>
> - *Oracle Automatic Storage Management Administrator's Guide* for more information about the initialization parameter `ASM_DISKSTRING`
>
> - See "Oracle ASM and Multipathing" in *Oracle Automatic Storage Management Administrator's Guide* for information about configuring Oracle ASM to work with multipathing, and consult your multipathing vendor documentation for details.

Creating Files on a NAS Device for Use with Oracle Automatic Storage Management

If you have a certified NAS storage device, then you can create zero-padded files in an NFS mounted directory and use those files as disk devices in an Oracle ASM disk group.

Ensure that you specify the ASM discovery path for Oracle ASM disks.

During installation of Oracle Grid Infrastructure, Oracle Universal Installer (OUI) can create files in the NFS mounted directory you specify. The following procedure explains how to manually create files in an NFS mounted directory to use as disk devices in an Oracle ASM disk group:

1. If necessary, create an exported directory for the disk group files on the NAS device.

2. Switch user to `root`.

3. Create a mount point directory on the local system.

   For example:

   ```
   # mkdir -p /mnt/oracleasm
   ```

4. To ensure that the NFS file system is mounted when the system restarts, add an entry for the file system in the mount file `/etc/fstab`.

5. Enter a command similar to the following to mount the NFS on the local system:

```
# mount /mnt/oracleasm
```

6. Choose a name for the disk group to create, and create a directory for the files on the NFS file system, using the disk group name as the directory name.

   For example, if you want to set up a disk group for a sales database:

```
# mkdir /mnt/oracleasm/sales1
```

7. Use commands similar to the following to create the required number of zero-padded files in this directory:

```
# dd if=/dev/zero
of=/mnt/oracleasm/sales1/disk1 bs=1024k
count=1000
```

   This example creates 1 GB files on the NFS file system. You must create one, two, or three files respectively to create an external, normal, or high redundancy disk group.

   > **✎ Note:**
   >
   > Creating multiple zero-padded files on the same NAS device does not guard against NAS failure. Instead, create one file for each NAS device and mirror them using the Oracle ASM technology.

8. Enter commands similar to the following to change the owner, group, and permissions on the directory and files that you created:

```
# chown -R grid:asmadmin /mnt/oracleasm
# chmod -R 660 /mnt/oracleasm
```

   In this example, the installation owner is `grid` and the OSASM group is `asmadmin`.

9. During Oracle Database installations, edit the Oracle ASM disk discovery string to specify a regular expression that matches the file names you created.

   For example:

```
/mnt/oracleasm/sales1/
```

Related Topics

- My Oracle Support Note 359515.1

Using Disk Groups with Oracle Database Files on Oracle ASM

Review this information to configure Oracle Automatic Storage Management (Oracle ASM) storage for Oracle Clusterware and Oracle Database Files.

- Identifying and Using Existing Oracle Database Disk Groups on Oracle ASM
  Identify existing disk groups and determine the free disk space that they contain.
  Optionally, identify failure groups for the Oracle ASM disk group devices.

- Configuring Disk Devices for Oracle ASM on HP-UX Itanium
  Complete these tasks to configure disk devices for use with Oracle Automatic
  Storage Management (Oracle ASM).

- Guidelines for Creating Disk Groups for Oracle Database Data Files
  If you are sure that a suitable disk group does not exist on the system, then install
  or identify appropriate disk devices to add to a new disk group.

- Creating Directories for Oracle Database Files
  You can store Oracle Database and recovery files on a separate file system from
  the configuration files.

Identifying and Using Existing Oracle Database Disk Groups on Oracle ASM

Identify existing disk groups and determine the free disk space that they contain.
Optionally, identify failure groups for the Oracle ASM disk group devices.

If you intend to use a normal or high redundancy disk group, then you can further
protect your database against hardware failure by associating a set of disk devices in
a custom failure group. By default, each device comprises its own failure group.
However, if two disk devices in a normal redundancy disk group are attached to the
same Host Bus Adapter (HBA), then the disk group becomes unavailable if the
adapter fails. The adapter in this example is a single point of failure.

To protect against failures of this type, you could use two HBAs, each with two disks,
and define a failure group for the disks attached to each adapter. This configuration
would enable the disk group to tolerate the failure of one HBA.

> **Note:**
>
> If you define custom failure groups, then you must specify a minimum of two
> failure groups for normal redundancy and three failure groups for high
> redundancy.

Related Topics

- *Oracle Automatic Storage Management Administrator's Guide*

Configuring Disk Devices for Oracle ASM on HP-UX Itanium

Complete these tasks to configure disk devices for use with Oracle Automatic Storage
Management (Oracle ASM).

1. If necessary, install the disks that you intend to use for the disk group and restart
   the system.

2. Ensure the disks are available:

This command displays information about each disk attached to the system, including the block device name (`/dev/dsk/cxtydz`).

```
# /usr/sbin/ioscan -fun -C disk
```

The output from this command is similar to the following:

```
Class   I  H/W Path    Driver S/W State   H/W Type     Description
============================================================================
disk    0  0/0/1/0.6.0 sdisk  CLAIMED       DEVICE       HP   DVD 6x/32x
                       /dev/dsk/c0t6d0   /dev/rdsk/c0t6d0
disk    1  0/0/1/1.2.0 sdisk  CLAIMED       DEVICE       SEAGATE ST39103LC
                       /dev/dsk/c1t2d0   /dev/rdsk/c1t2d0
```

On HP-UX 11i v.3, you can also use agile view to review mass storage devices, including block devices (`/dev/disk/diskxyz`), or character raw devices (`/dev/rdisk/diskxyz`). For example:

```
#>/usr/sbin/ioscan -funN -C disk
Class I    H/W Path       Driver  S/W State H/W Type   Desc
====================================================================
disk  4  64000/0xfa00/0x1 esdisk  CLAIMED    DEVICE  HP73.4GST373454LC
         /dev/disk/disk4    /dev/rdisk/disk4
disk 907 64000/0xfa00/0x2f esdisk  CLAIMED    DEVICE  COMPAQ MSA1000 VOLUME
         /dev/disk/disk907  /dev/rdisk/disk907
```

3. If the `ioscan` command does not display the device name information for a device that you want to use, then enter the following command to install the special device files for any new devices:

```
# /usr/sbin/insf -e
```

4. For each disk that you want to add to a disk group, verify that it is not already part of an LVM volume group:

```
# /sbin/pvdisplay /dev/dsk/cxtydz
```

If this command displays volume group information, then the disk is already part of a volume group. The disks that you choose must not be part of an LVM volume group.

> **✎ Note:**
>
> If you are using different volume management software, for example VERITAS Volume Manager, then refer to the appropriate documentation for information about verifying that a disk is not in use.

5. Change the owner, group, and permissions on the character file for each disk that you want to add to a disk group:

```
# chown oracle:dba /dev/rdsk/cxtydz
# chmod 660 /dev/rdsk/cxtydz
```

> ✎ **Note:**
>
> If you are using a multi-pathing disk driver with Oracle Automatic Storage Management, then ensure that you set the permissions only on the correct logical device name for the disk.

Guidelines for Creating Disk Groups for Oracle Database Data Files

If you are sure that a suitable disk group does not exist on the system, then install or identify appropriate disk devices to add to a new disk group.

Use the following guidelines when identifying appropriate disk devices:

- All of the devices in an Oracle ASM disk group should be the same size and have the same performance characteristics.

- Do not specify multiple partitions on a single physical disk as a disk group device. Oracle ASM expects each disk group device to be on a separate physical disk.

- Although you can specify a logical volume as a device in an Oracle ASM disk group, Oracle does not recommend their use because it adds a layer of complexity that is unnecessary with Oracle ASM. In addition, Oracle RAC requires a cluster logical volume manager in case you decide to use a logical volume with Oracle ASM and Oracle RAC.

Creating Directories for Oracle Database Files

You can store Oracle Database and recovery files on a separate file system from the configuration files.

Perform this procedure to place the Oracle Database or recovery files on a separate file system from the Oracle base directory:

1. Use the following command to determine the free disk space on each mounted file system:

   ```
   # df -h
   ```

2. Identify the file systems to use, from the display:

Option	Description
Database Files	Select one of the following: • A single file system with at least 1.5 GB of free disk space • Two or more file systems with at least 3.5 GB of free disk space in total
Recovery Files	Choose a file system with at least 2 GB of free disk space

   If you are using the same file system for multiple file types, then add the disk space requirements for each type to determine the total disk space requirement.

3. Note the names of the mount point directories for the file systems that you identified.

4. If the user performing installation has permissions to create directories on the disks where you plan to install Oracle Database, then DBCA creates the Oracle Database file directory, and the Recovery file directory. If the user performing installation does not have write access, then you must create these directories manually.

For example, given the user `oracle` and Oracle Inventory Group `oinstall`, and using the paths `/u03/oradata/wrk_area` for Oracle Database files, and `/u01/oradata/rcv_area` for the recovery area, these commands create the recommended subdirectories in each of the mount point directories and set the appropriate owner, group, and permissions on them:

- Database file directory:

```
# mkdir -p /u01/oradata/
# chown oracle:oinstall /u01/oradata/
# chmod 775 /u01/oradata
```

The default location for the database file directory is `$ORACLE_BASE/oradata`.

- Recovery file directory (fast recovery area):

```
# mkdir -p /u01/oradata/rcv_area
# chown oracle:oinstall /u01/oradata/rcv_area
# chmod 775 /u01/oradata/rcv_area
```

The default fast recovery area is `$ORACLE_BASE/fast_recovery_area`.

Oracle recommends that you keep the fast recovery area on a separate physical disk than that of the database file directory. This method enables you to use the fast recovery area to retrieve data if the disk containing oradata is unusable for any reason.

Configuring File System Storage for Oracle Database

Complete these procedures to use file system storage for Oracle Database.

If you plan to place storage on Network File System (NFS) protocol devices, then Oracle recommends that you use Oracle Direct NFS (dNFS) to take advantage of performance optimizations built into the Oracle Direct NFS client.

For optimal database organization and performance, Oracle recommends that you install data files and the Oracle Database software in different disks.

- Configuring NFS Buffer Size Parameters for Oracle Database
  Set the values for the NFS buffer size parameters `rsize` and `wsize` to at least 32768.

- Checking TCP Network Protocol Buffer for Direct NFS Client
  Check your TCP network buffer size to ensure that it is adequate for the speed of your servers.

- Creating an oranfstab File for Direct NFS Client
  Direct NFS uses a configuration file, `oranfstab`, to determine the available mount points.

- Enabling and Disabling Direct NFS Client Control of NFS
  Use these commands to enable or disable Direct NFS Client Oracle Disk Manager Control of NFS.

- • Enabling Hybrid Columnar Compression on Direct NFS Client
  Perform these steps to enable Hybrid Columnar Compression (HCC) on Direct
  NFS Client:

Configuring NFS Buffer Size Parameters for Oracle Database

Set the values for the NFS buffer size parameters `rsize` and `wsize` to at least 32768.

For example, to use `rsize` and `wsize` buffer settings with the value 32768 for an Oracle Database data files mount point, set mount point parameters to values similar to the following:

```
nfs_server:/vol/DATA/oradata /home/oracle/netapp nfs\
rw,bg,hard,nointr,rsize=32768,wsize=32768,tcp,actimeo=0,vers=3,timeo=600
```

Direct NFS Client issues writes at `wtmax` granularity to the NFS server.

Related Topics

- • My Oracle Support note 359515.1

Checking TCP Network Protocol Buffer for Direct NFS Client

Check your TCP network buffer size to ensure that it is adequate for the speed of your servers.

By default, the network buffer size is set to 1 MB for TCP, and 2 MB for UDP. The TCP buffer size can set a limit on file transfers, which can negatively affect performance for Direct NFS Client users.

To check the current TCP buffer size:

```
bash-4.0$ ndd -get /dev/tcp tcp_xmit_hiwater_max
bash-4.0$ ndd -get /dev/tcp tcp_recv_hiwater_max
```

Oracle recommends that you set the value based on the link speed of your servers. For example:

```
bash-4.0# ndd -set /dev/tcp tcp_xmit_hiwater_max 10485760
bash-4.0# ndd -set /dev/tcp tcp_recv_hiwater_max 10485760
```

Creating an oranfstab File for Direct NFS Client

Direct NFS uses a configuration file, `oranfstab`, to determine the available mount points.

Create an `oranfstab` file with the following attributes for each NFS server that you want to access using Direct NFS Client:

- • `server`

  The NFS server name.

  For NFS setup with Kerberos authentication, the `server` attribute name must be the fully-qualified name of the NFS server. This `server` attribute name is used to

create service principal for Ticket Granting Service (TGS) request from the Kerberos server. If you are configuring external storage snapshot cloning, then the NFS `server` name should be a valid host name. For all other scenarios, the NFS `server` name can be any unique name.

- `local`

    Up to four paths on the database host, specified by IP address or by name, as displayed using the `ifconfig` command run on the database host.

- `path`

    Up to four network paths to the NFS server, specified either by IP address, or by name, as displayed using the `ifconfig` command on the NFS server.

- `export`

    The exported path from the NFS server.

- `mount`

    The corresponding local mount point for the exported volume.

- `mnt_timeout`

    Specifies (in seconds) the time Direct NFS Client should wait for a successful mount before timing out. This parameter is optional. The default timeout is 10 minutes (600).

- `nfs_version`

    Specifies the NFS protocol version used by Direct NFS Client. Possible values are NFSv3, NFSv4, NFSv4.1, and pNFS. The default version is NFSv3. If you select NFSv4.x, then you must configure the value in `oranfstab` for `nfs_version`.

    Specify `nfs_version` as `pNFS`, if you want to use Direct NFS with Parallel NFS. Direct NFS supports only the default `sys` security authentication with Parallel NFS. Direct NFS does not support Parallel NFS when combined with any of the Kerberos authentication parameters.

- `security_default`

    Specifies the default security mode applicable for all the exported NFS server paths for a server entry. This parameter is optional. `sys` is the default value. See the description of the `security` parameter for the supported security levels for the `security_default` parameter.

- `security`

    Specifies the security level, to enable security using Kerberos authentication protocol with Direct NFS Client. This optional parameter can be specified per export-mount pair. The supported security levels for the `security_default` and `security` parameters are:

    `sys`: UNIX level security AUTH_UNIX authentication based on user identifier (UID) and group identifier (GID) values. This is the default value for security parameters.
    `krb5`: Direct NFS runs with plain Kerberos authentication. Server is authenticated as the real server which it claims to be.
    `krb5i`: Direct NFS runs with Kerberos authentication and NFS integrity. Server is authenticated and each of the message transfers is checked for integrity.
    `krb5p`: Direct NFS runs with Kerberos authentication and NFS privacy. Server is authenticated, and all data is completely encrypted.

The `security` parameter, if specified, takes precedence over the `security_default` parameter. If neither of these parameters are specified, then `sys` is the default authentication.

For NFS server Kerberos security setup, review the relevant NFS server documentation. For Kerberos client setup, review the relevant operating system documentation.

- `dontroute`

  Specifies that outgoing messages should not be routed by the operating system, but instead sent using the IP address to which they are bound.

  > **✎ Note:**
  >
  > The `dontroute` option is a POSIX option, which sometimes does not work on Linux systems with multiple paths in the same subnet.

- `management`

  Enables Direct NFS Client to use the management interface for SNMP queries. You can use this parameter if SNMP is running on separate management interfaces on the NFS server. The default value is the server parameter value.

- `community`

  Specifies the community string for use in SNMP queries. Default value is `public`.

The following examples show three possible NFS server entries in `oranfstab`. A single `oranfstab` can have multiple NFS server entries.

Example 7-1   Using Local and Path NFS Server Entries

The following example uses both local and path. Because they are in different subnets, you do not have to specify `dontroute`.

```
server: MyDataServer1
local: 192.0.2.0
path: 192.0.2.1
local: 192.0.100.0
path: 192.0.100.1
export: /vol/oradata1 mount: /mnt/oradata1
```

Example 7-2   Using Local and Path in the Same Subnet, with dontroute

Local and path in the same subnet, where `dontroute` is specified:

```
server: MyDataServer2
local: 192.0.2.0
path: 192.0.2.128
local: 192.0.2.1
path: 192.0.2.129
dontroute
export: /vol/oradata2 mount: /mnt/oradata2
```

Example 7-3   Using Names in Place of IP Addresses, with Multiple Exports, management and community

```
server: MyDataServer3
local: LocalPath1
path: NfsPath1
local: LocalPath2
path: NfsPath2
local: LocalPath3
path: NfsPath3
local: LocalPath4
path: NfsPath4
dontroute
export: /vol/oradata3 mount: /mnt/oradata3
export: /vol/oradata4 mount: /mnt/oradata4
export: /vol/oradata5 mount: /mnt/oradata5
export: /vol/oradata6 mount: /mnt/oradata6
management: MgmtPath1
community: private
```

Example 7-4   Using Kerberos Authentication with Direct NFS Export

The `security` parameter overrides `security_default`:

```
server:  nfsserver
 local: 192.0.2.0
 path: 192.0.2.2
 local: 192.0.2.3
 path: 192.0.2.4
 export: /private/oracle1/logs  mount: /logs  security: krb5
 export: /private/oracle1/data  mount: /data  security: krb5p
 export: /private/oracle1/archive mount: /archive security: sys
 export: /private/oracle1/data1 mount: /data1
 security_default: krb5i
```

Enabling and Disabling Direct NFS Client Control of NFS

Use these commands to enable or disable Direct NFS Client Oracle Disk Manager Control of NFS.

By default, Direct NFS Client is installed in an enabled state. However, if Direct NFS Client is disabled and you want to enable it, complete the following steps on each node. If you use a shared Grid home for the cluster, then complete the following steps in the shared Grid home:

1. Log in as the Oracle Grid Infrastructure installation owner.

2. Change directory to `Grid_home/rdbms/lib`.

3. Enter the following command:

   ```
   $ make -f ins_rdbms.mk dnfs_on
   ```

> **Note:**
>
> If you remove an NFS path that an Oracle Database is using, then you must restart the database for the change to take effect.

Enabling Hybrid Columnar Compression on Direct NFS Client

Perform these steps to enable Hybrid Columnar Compression (HCC) on Direct NFS Client:

1. Ensure that SNMP is enabled on the ZFS storage server. For example:

    ```
    $ snmpget -v1 -c public server_name .1.3.6.1.4.1.42.2.225.1.4.2.0
    SNMPv2-SMI::enterprises.42.2.225.1.4.2.0 = STRING: "Sun Storage
    7410"
    ```

2. If SNMP is enabled on an interface other than the NFS server, then configure `oranfstab` using the management parameter.

3. If SNMP is configured using a community string other than public, then configure `oranfstab` file using the community parameter.

4. Ensure that `libnetsnmp.so` is installed by checking if `snmpget` is available.

Creating and Using Oracle ASM Credentials File

Review this information to create an Oracle ASM credentials file.

An Oracle ASM Storage Client does not have Oracle ASM running on the nodes and uses Oracle ASM storage services in a different client cluster.

The Oracle ASM Client Cluster requires Grid Naming Service (GNS) to be configured in the Oracle ASM Server Cluster.

1. Connect to any Oracle ASM instance as SYSASM user and execute the query:

    ```
    ALTER DISKGROUP data SET ATTRIBUTE 'access_control.enabled' =
    'true';
    ```

2. From the *Grid_home*/bin directory on the Storage Server, run the following command on one of the member nodes, where `credential_file` is the name and path location of the Oracle ASM credentials file you create:

    ```
    Grid_home/bin/asmcmd mkcc client_cluster_name credential_file
    ```

    For example:

    ```
    Grid_home/bin/asmcmd mkcc clientcluster1 /home/grid/
    clientcluster1_credentials.xml
    ```

3. Copy the Oracle ASM credentials file to a secure path on the client cluster node where you run the client cluster installation.

The Oracle Installation user must have permissions to access that file. Oracle recommends that no other user is granted permissions to access the Oracle ASM credentials file. During installation, you are prompted to provide a path to the file.

> **✎ Note:**
>
> - The Oracle ASM credentials file can be used only once. If an Oracle ASM Storage Client is configured and deconfigured, you must create a new Oracle ASM credentials file.
>
> - If the Oracle ASM credentials file is used to configure the client cluster, then it cannot be shared or reused to configure another client cluster.

Creating a Remote GIMR Client Data File for Oracle Standalone Clusters

Create a remote GIMR client data file to specify the Oracle Standalone Cluster configuration for the remote Grid Infrastructure Management Repository (GIMR).

1. As the `grid` user, run the following command from the Oracle Domain Services Cluster with remote GIMR or Oracle Standalone Cluster with remote GIMR to create the GIMR client data file:

   ```
   $ /u01/app/21.0.0/gimr/bin/mgmtca createRepos -clientDataFile
   client_data_file_location -clusterName standalone_cluster_name
   [-version standalone_cluster_version]
   ```

   `-clientDataFile` is the full path of the xml file to export the client data, `-clusterName` is the name of the Oracle Standalone Cluster, and `-version` is the five digit Oracle Standalone Cluster version of the cluster for which you want to add a remote GIMR, for example, `21.0.0.0.0`.

   This command creates a GIMR client data file containing configuration details about Grid Infrastructure Management Repository (GIMR) for the Oracle Standalone Cluster with remote GIMR.

2. Copy the client data file to a location on the Oracle Standalone Cluster, and select the file during the installation and configuration of the Oracle Standalone Cluster with remote GIMR.

8

Installing Oracle Grid Infrastructure

Review this information for installation and deployment options for Oracle Grid Infrastructure.

Oracle Database and Oracle Grid Infrastructure installation software is available in multiple media, and can be installed using several options. The Oracle Grid Infrastructure software is available as an image, available for download from the Oracle Technology Network website, or the Oracle Software Delivery Cloud portal. In most cases, you use the graphical user interface (GUI) provided by Oracle Universal Installer to install the software. You can also use Oracle Universal Installer to complete silent mode installations, without using the GUI.

- About Image-Based Oracle Grid Infrastructure Installation
  Installation and configuration of Oracle Grid Infrastructure software is simplified with image-based installation.

- Setup Wizard Installation Options for Creating Images
  Gold image-creation options to use with your Oracle Database or Oracle Grid Infrastructure installation setup wizards.

- Understanding Cluster Configuration Options
  Review these topics to understand the cluster configuration options available in Oracle Grid Infrastructure.

- Installing Oracle Grid Infrastructure for a New Cluster
  Review these procedures to install the cluster configuration options available in this release of Oracle Grid Infrastructure.

- Installing Oracle Grid Infrastructure Using a Cluster Configuration File
  During installation of Oracle Grid Infrastructure, you have the option of either of providing cluster configuration information manually, or of using a cluster configuration file.

- Installing Only the Oracle Grid Infrastructure Software
  This installation option requires manual postinstallation steps to enable the Oracle Grid Infrastructure software.

- Confirming Oracle Clusterware Function
  After Oracle Grid Infrastructure installation, confirm that your Oracle Clusterware installation is installed and running correctly.

- Confirming Oracle ASM Function for Oracle Clusterware Files
  Confirm Oracle ASM is running after installing Oracle Grid Infrastructure.

- Understanding Offline Processes in Oracle Grid Infrastructure
  After the installation of Oracle Grid Infrastructure, some components may be listed as OFFLINE. Oracle Grid Infrastructure activates these resources when you choose to add them.

About Image-Based Oracle Grid Infrastructure Installation

Installation and configuration of Oracle Grid Infrastructure software is simplified with image-based installation.

To install Oracle Grid Infrastructure, create the new Grid home with the necessary user group permissions, and then extract the image file into the newly-created Grid home, and run the setup wizard to register the Oracle Grid Infrastructure product.

Using image-based installation, you can do the following:

- Configure Oracle Grid Infrastructure for a new cluster.

- Configure Oracle Grid Infrastructure for a standalone server (Oracle Restart).

- Upgrade Oracle Grid Infrastructure.

- Setup software only.

- Add or remove nodes from your existing cluster, if the Oracle Grid Infrastructure software is already installed or configured.

This installation feature streamlines the installation process and supports automation of large-scale custom deployments. You can also use this installation method for deployment of customized images, after you patch the base-release software with the necessary Release Updates (RUs) or Release Update Revisions (RURs).

> **Note:**
>
> You must extract the image software into the directory where you want your Grid home to be located, and then run the `%ORACLE_HOME%\gridSetup.sh` script to start the Oracle Grid Infrastructure Setup Wizard. Ensure that the Grid home directory path you create is in compliance with the Oracle Optimal Flexible Architecture recommendations.

Setup Wizard Installation Options for Creating Images

Gold image-creation options to use with your Oracle Database or Oracle Grid Infrastructure installation setup wizards.

Options

In image-based installations, you can start your Oracle Database installation or Oracle Grid Infrastructure installations by running the setup wizards `runInstaller` and `gridSetup.sh` respectively. Both these wizards come with the following image-creation options.

Table 8-1   Image-Creation Options for Setup Wizard

Option	Description
-createGoldImage	Creates a gold image from the current Oracle home.
-destinationLocation	Specify the complete path, or location, where the gold image will be created.
-exclFiles	Specify the complete paths to the files to be excluded from the newly created gold image.
—help	Displays help for all the available options.

For example:

```
./runInstaller -createGoldImage -destinationLocation /tmp/my_db_images -
exclFiles /u01/app/oracle/product/21.0.0/dbhome_1/relnotes
```

```
./gridSetup.sh -createGoldImage -destinationLocation /tmp/my_grid_images -
exclFiles /u01/app/oracle/product/21.0.0/dbhome_1/relnotes
```

Where:

`/tmp/my_db_images` is a temporary file location where the image zip file is created.

`/tmp/my_grid_images` is a temporary file location where the image zip file is created.

`/u01/app/oracle/product/21.0.0/dbhome_1/relnotes` is the file to be excluded in the newly created gold image.

Understanding Cluster Configuration Options

Review these topics to understand the cluster configuration options available in Oracle Grid Infrastructure.

- About Oracle Standalone Clusters
  An Oracle Standalone Cluster hosts all Oracle Grid Infrastructure services and Oracle ASM locally and requires direct access to shared storage.
- About Oracle Extended Clusters
  An Oracle Extended Cluster consists of nodes that are located in multiple locations called sites.

About Oracle Standalone Clusters

An Oracle Standalone Cluster hosts all Oracle Grid Infrastructure services and Oracle ASM locally and requires direct access to shared storage.

Oracle Standalone Clusters contain multiple cluster nodes. The number of nodes in an Oracle Standalone Cluster can be as many as 64. The cluster nodes can host different types of applications. Oracle Standalone Cluster nodes are tightly connected, and have direct access to shared storage. Shared storage is locally mounted on each of the Oracle Standalone Cluster nodes, with an Oracle ASM instance or a shared file system location available to all the nodes.

Oracle Standalone Clusters host Grid Infrastructure Management Repository (GIMR) locally, if GIMR was configured during the installation. The GIMR is a multitenant database, which stores information about the cluster. This information includes the real time performance data the Cluster Health Monitor collects.

When you deploy an Oracle Standalone Cluster, you can also choose to configure it as an Oracle Extended cluster. An Oracle Extended Cluster consists of nodes that are located in multiple locations or sites.

About Oracle Extended Clusters

An Oracle Extended Cluster consists of nodes that are located in multiple locations called sites.

When you deploy an Oracle Standalone Cluster, you can also choose to configure the cluster as an Oracle Extended Cluster. You can extend an Oracle RAC cluster across two, or more, geographically separate sites, each equipped with its own storage. In the event that one of the sites fails, the other site acts as an active standby.

Both Oracle ASM and the Oracle Database stack, in general, are designed to use enterprise-class shared storage in a data center. Fibre Channel technology, however, enables you to distribute compute and storage resources across two or more data centers, and connect them through Ethernet cables and Fibre Channel, for compute and storage needs, respectively.

You can configure an Oracle Extended Cluster when you install Oracle Grid Infrastructure. You can also do so post installation using the `ConvertToExtended` script. You manage your Oracle Extended Cluster using CRSCTL.

You can assign nodes and failure groups to sites. Sites contain failure groups, and failure groups contain disks.

The following conditions apply when you select redundancy levels for Oracle Extended Clusters:

Table 8-2   Oracle ASM Disk Group Redundancy Levels for Oracle Extended Clusters with 2 Data Sites

Redundancy Level	Number of Failure Groups for OCR and Voting Files Disk Groups	Number of Failure Groups for OCR Backup and GIMR Disk Groups
Normal redundancy	1 failure group per data site, 1 quorum failure group	1 failure group per data site
Flex redundancy	1 failure group per data site, 1 quorum failure group	1 failure group per data site, 1 quorum failure group
Extended redundancy	3 failure groups each for 2 data sites, 1 quorum failure group outside the 2 data sites	3 failure groups each for 2 data sites, 1 quorum failure group outside the 2 data sites
High redundancy	Not supported	Not supported

Related Topics

- *Oracle Clusterware Administration and Deployment Guide*

Installing Oracle Grid Infrastructure for a New Cluster

Review these procedures to install the cluster configuration options available in this release of Oracle Grid Infrastructure.

- About Oracle Grid Infrastructure Installation
  You can install Oracle Grid Infrastructure as Oracle Standalone Cluster or Oracle Domain Services Cluster.

- **Installing Oracle Standalone Cluster**
  Complete this procedure to install Oracle Grid Infrastructure software for Oracle Standalone Cluster.

About Oracle Grid Infrastructure Installation

You can install Oracle Grid Infrastructure as Oracle Standalone Cluster or Oracle Domain Services Cluster.

Starting with Oracle Grid Infrastructure 12c Release 2 (12.2), the installation media is replaced with a zip file for the Oracle Grid Infrastructure installer. Run the installation wizard after extracting the zip file into the target home path.

At any time during installation, if you have a question about what you are being asked to do, or what input you are required to provide during installation, click the **Help** button on the installer window.

You should have your network information, storage information, and operating system users and groups available to you before you start installation, and you should be prepared to run root scripts.

As the user that owns the software for Oracle Grid Infrastructure for a cluster (`grid`) on the first node, install Oracle Grid Infrastructure for a cluster. Note that the installer uses Secure Shell (SSH) to copy the binary files from this node to the other nodes during the installation. During installation, in the Cluster Node Information window, when you specify the nodes in your cluster, you can click **SSH Connectivity** and the installer configures SSH connectivity between the specified nodes for you.

> **Note:**
>
> These installation instructions assume you do not already have any Oracle software installed on your system.

Installing Oracle Standalone Cluster

Complete this procedure to install Oracle Grid Infrastructure software for Oracle Standalone Cluster.

1. As the `grid` user, download the Oracle Grid Infrastructure image files and extract the files into the Grid home. For example:

   ```
   mkdir -p /u01/app/21.0.0/grid
   chown grid:oinstall /u01/app/21.0.0/grid
   cd /u01/app/21.0.0/grid
   unzip -q download_location/grid.zip
   ```

   `grid.zip` is the name of the Oracle Grid Infrastructure image zip file.

> **Note:**
>
> - You must extract the zip image software into the directory where you want your Grid home to be located.
>
> - Download and copy the Oracle Grid Infrastructure image files to the local node only. During installation, the software is copied and installed on all other nodes in the cluster.

2. Log in as the `grid` user, and start the Oracle Grid Infrastructure installer by running the following command:

   ```
   /u01/app/21.0.0/grid/gridSetup.sh
   ```

   The installer starts and the Select Configuration Option window appears.

3. Choose the option **Configure Grid Infrastructure for a New Cluster**, then click **Next**.

   The Select Cluster Configuration window appears.

4. Choose the option **Configure an Oracle Standalone Cluster**, then click **Next**.

   Select the **Configure as Extended Cluster** option to extend an Oracle RAC cluster across two or more separate sites, each equipped with its own storage.

   The Grid Plug and Play Information window appears.

5. In the **Cluster Name** and **SCAN Name** fields, enter the names for your cluster and cluster scan that are unique throughout your entire enterprise network.

   You can select **Configure GNS** if you have configured your domain name server (DNS) to send to the GNS virtual IP address name resolution requests for the subdomain GNS serves, as explained in this guide.

   For cluster member node public and VIP network addresses, provide the information required depending on the kind of cluster you are configuring:

   - If you plan to use automatic cluster configuration with DHCP addresses configured and resolved through GNS, then you only need to provide the GNS VIP names as configured on your DNS.

   - If you plan to use manual cluster configuration, with fixed IP addresses configured and resolved on your DNS, then provide the SCAN names for the cluster, and the public names, and VIP names for each cluster member node. For example, you can choose a name that is based on the node names' common prefix. The cluster name can be `mycluster` and the cluster SCAN name can be `mycluster-scan`.

   Click **Next**.

   The Cluster Node Information window appears.

6. In the Public Hostname column of the table of cluster nodes, you should see your local node, for example `node1.example.com`.

   The following is a list of additional information about node IP addresses:

   - For the local node only, OUI automatically fills in public and VIP fields. If your system uses vendor clusterware, then OUI may fill additional fields.

- Host names and virtual host names are not domain-qualified. If you provide a domain in the address field during installation, then OUI removes the domain from the address.

- Interfaces identified as private for private IP addresses should not be accessible as public interfaces. Using public interfaces for Cache Fusion can cause performance problems.

- When you enter the public node name, use the primary host name of each node. In other words, use the name displayed by the `/bin/hostname` command.

  a. Click **Add** to add another node to the cluster.

  b. Enter the second node's public name (`node2`), and virtual IP name (`node2-vip`), then click **OK**. Provide the virtual IP (VIP) host name for all cluster nodes, or none.

     You are returned to the Cluster Node Information window. You should now see all nodes listed in the table of cluster nodes.

  c. Make sure all nodes are selected, then click the **SSH Connectivity** button at the bottom of the window.

     The bottom panel of the window displays the SSH Connectivity information.

  d. Enter the operating system user name and password for the Oracle software owner (`grid`). If you have configured SSH connectivity between the nodes, then select the **Reuse private and public keys existing in user home** option. Click **Setup**.

     A message window appears, indicating that it might take several minutes to configure SSH connectivity between the nodes. After a short period, another message window appears indicating that passwordless SSH connectivity has been established between the cluster nodes. Click **OK** to continue.

  e. When returned to the Cluster Node Information window, click **Next** to continue.

     The Specify Network Interface Usage window appears.

7. Select the usage type for each network interface displayed.

   Verify that each interface has the correct interface type associated with it. If you have network interfaces that should not be used by Oracle Clusterware, then set the network interface type to **Do Not Use**. For example, if you have only two network interfaces, then set the public interface to have a Use for value of **Public** and set the private network interface to have a Use for value of **ASM & Private**.

   Click **Next**. The Storage Option Information window appears.

8. Select storage option for Oracle Cluster Registry (OCR) and voting files:

   a. Select **Use Oracle Flex ASM for storage** to store OCR and voting files on an Oracle ASM disk group.

   b. Select **Use Shared File System** to store OCR and voting files on a shared file system, and then click Next. The Create Grid Infrastructure Management Repository Option window appears.

9. Select an option to create a Grid Infrastructure Management Repository for your Oracle Standalone Cluster installation, then click **Next**.

   a. Select **Use a Local GIMR database** to create a local GIMR by installing Oracle RACsoftware and creating GIMR container database after completing Oracle Grid Infrastructure installation.

   b. Select **Use an Existing remote GIMR database** to use a remote GIMR from an Oracle Standalone Cluster, and specify a GIMR client data file. Refer to the *Creating*

a Remote GIMR Client Data File for Oracle Standalone Clusters for steps to create a GIMR client data file.

   c. Select **Do Not use a GIMR database** to do not create a GIMR.

If you choose **Yes** on this window, then the Grid Infrastructure Management Repository Option window appears. Otherwise, the Create ASM Disk Group or ASM Client Storage Option window appears, based on your storage option.

> **Note:**
>
> If you choose **No** on this window and you create a GIMR later, then create a separate Oracle ASM disk group for GIMR using Oracle ASMCA.

10. Choose whether you want to store the Grid Infrastructure Management Repository in a separate Oracle ASM disk group, then click **Next**.

The Create ASM Disk Group window appears.

11. Provide the name and specifications for the Oracle ASM disk group.

   a. In the **Disk Group Name** field, enter a name for the disk group, for example DATA.

   b. Choose the Redundancy level for this disk group. Normal is the recommended option.

   c. In the Add Disks section, choose the disks to add to this disk group.

When you have finished providing the information for the disk group, click **Next**.

12. If you selected to use a different disk group for the GIMR, then the Grid Infrastructure Management Repository Option window appears. Provide the name and specifications for the GIMR disk group.

   a. In the **Disk Group Name** field, enter a name for the disk group, for example DATA.

   b. Choose the Redundancy level for this disk group. Normal is the recommended option.

   c. In the Add Disks section, choose the disks to add to this disk group.

When you have finished providing the information for the disk group, click **Next**.

The Specify ASM Password window appears.

13. Choose the same password for the Oracle ASM SYS and ASMSNMP account, or specify different passwords for each account, then click **Next**.

The Failure Isolation Support window appears.

14. Select the option **Do not use Intelligent Platform Management Interface (IPMI)**, then click **Next**.

The Specify Management Options window appears.

15. If you have Enterprise Manager Cloud Control installed in your enterprise, then choose the option **Register with Enterprise Manager (EM) Cloud Control** and provide the EM configuration information. If you do not have Enterprise Manager Cloud Control installed in your enterprise, then click **Next** to continue.

The Privileged Operating System Groups window appears.

16. Accept the default operating system group names for Oracle ASM administration and click **Next**.

    The Specify Install Location window appears.

17. Specify the directory to use for the Oracle base for the Oracle Grid Infrastructure installation, then click **Next**. The Oracle base directory must be different from the Oracle home directory.

    If you copied the Oracle Grid Infrastructure installation files into the Oracle Grid home directory as directed in Step 1, then the default location for the Oracle base directory should display as `/u01/app/grid`.

    If you have not installed Oracle software previously on this computer, then the Create Inventory window appears.

18. Change the path for the inventory directory, if required. Then, click **Next**.

    If you are using the same directory names as the examples in this book, then it should show a value of `/u01/app/oraInventory`. The group name for the `oraInventory` directory should show `oinstall`.

    The Root Script Execution Configuration window appears.

19. Select the option to **Automatically run configuration scripts.** Enter the credentials for the root user or a sudo account, then click **Next**.

    Alternatively, you can Run the scripts manually as the `root` user at the end of the installation process when prompted by the installer.

    The Perform Prerequisite Checks window appears.

20. If any of the checks have a status of Failed and are not Fixable, then you must manually correct these issues. After you have fixed the issue, you can click the **Check Again** button to have the installer recheck the requirement and update the status. Repeat as needed until all the checks have a status of Succeeded. Click **Next**.

    The Summary window appears.

21. Review the contents of the Summary window and then click **Install**.

    The installer displays a progress indicator enabling you to monitor the installation process.

22. If you did not configure automation of the root scripts, then you are required to run certain scripts as the `root` user, as specified in the Execute Configuration Scripts window. Do *not* click **OK** until you have run all the scripts. Run the scripts on all nodes as directed, in the order shown.

    For example, on Oracle Linux you perform the following steps (note that for clarity, the examples show the current user, node and directory in the prompt):

    a. As the `grid` user on `node1`, open a terminal window, and enter the following commands:

    ```
    [grid@node1 grid]$ cd /u01/app/oraInventory
    [grid@node1 oraInventory]$ su
    ```

    b. Enter the password for the `root` user, and then enter the following command to run the first script on `node1`:

    ```
    [root@node1 oraInventory]# ./orainstRoot.sh
    ```

c. After the `orainstRoot.sh` script finishes on `node1`, open another terminal window, and as the `grid` user, enter the following commands:

```
[grid@node1 grid]$ ssh node2
[grid@node2 grid]$ cd /u01/app/oraInventory
[grid@node2 oraInventory]$ su
```

d. Enter the password for the `root` user, and then enter the following command to run the first script on `node2`:

```
[root@node2 oraInventory]#./orainstRoot.sh
```

e. After the `orainstRoot.sh` script finishes on `node2`, go to the terminal window you opened in part a of this step. As the `root` user on `node1`, enter the following commands to run the second script, `root.sh`:

```
[root@node1 oraInventory]# cd /u01/app/21.0.0/grid
[root@node1 grid]# ./root.sh
```

Press **Enter** at the prompt to accept the default value.

> **Note:**
>
> You must run the `root.sh` script on the first node and wait for it to finish. You can run `root.sh` scripts concurrently on all other nodes except for the last node on which you run the script. Like the first node, the `root.sh` script on the last node must be run separately.

f. After the `root.sh` script finishes on `node1`, go to the terminal window you opened in part c of this step. As the `root` user on `node2`, enter the following commands:

```
[root@node2 oraInventory]#cd /u01/app/21.0.0/grid
[root@node2 grid]#./root.sh
```

After the `root.sh` script completes, return to the Oracle Universal Installer window where the Installer prompted you to run the `orainstRoot.sh` and `root.sh` scripts. Click **OK**.

The software installation monitoring window reappears.

23. Continue monitoring the installation until the Finish window appears. Then click **Close** to complete the installation process and exit the installer.

> ⚠ **Caution:**
>
> After installation is complete, do not remove manually or run `cron` jobs that remove `/tmp/.oracle` or `/var/tmp/.oracle` directories or their files while Oracle software is running on the server. If you remove these files, then the Oracle software can encounter intermittent hangs. Oracle Clusterware installations can fail with the error:
>
> ```
> CRS-0184: Cannot communicate with the CRS daemon.
> ```

After your Oracle Grid Infrastructure installation is complete, you can install Oracle Database on a cluster node for high availability, or install Oracle RAC.

> ✎ **See Also:**
>
> *Oracle Real Application Clusters Installation Guide* or *Oracle Database Installation Guide* for your platform for information on installing Oracle Database

Installing Oracle Grid Infrastructure Using a Cluster Configuration File

During installation of Oracle Grid Infrastructure, you have the option of either of providing cluster configuration information manually, or of using a cluster configuration file.

A cluster configuration file is a text file that you can create before starting `gridSetup.sh`, which provides the installer with cluster node addresses that it requires to configure the cluster.
Oracle recommends that you consider using a cluster configuration file if you intend to perform repeated installations on a test cluster, or if you intend to perform an installation on many nodes. A sample cluster configuration file is available in the directory *Grid_home*/`install/response/sample.ccf`.

To create a cluster configuration file manually, start a text editor, and create a file that provides the name of the public and virtual IP addresses for each cluster member node, in the following format:

```
node1 node1-vip /node-role
node2 node2-vip /node-role
.
.
.
```

Specify the different nodes, separating them with either spaces or colon (:).

For example:

```
mynode1 mynode1-vip
mynode2 mynode2-vip
```

Or, for example:

```
mynode1:mynode1-vip
mynode2:mynode2-vip
```

Example 8-1   Sample Cluster Configuration File

The following sample cluster configuration file is available in the directory
Grid_home/install/response/sample.ccf:

```
#
# Cluster nodes configuration specification file
#
# Format:
# node [vip] [role-identifier] [site-name]
#
# node            - Node's public host name
# vip             - Node's virtual host name
# site-name       - Node's assigned site
#
# Specify details of one node per line.
# Lines starting with '#' will be skipped.
#
# (1) vip and role are not required for Oracle Grid Infrastructure
software only
#     installs
# (2) vip should be specified as AUTO if Node Virtual host names are
Dynamically
#     assigned
# (3) site-name should be specified only when configuring Oracle Grid
Infrastructure with "Extended Cluster" option
#
# Examples:
# --------
# For installing GI software only on a cluster:
# ^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^
# node1
# node2
#
# For Standalone Cluster:
# ^^^^^^^^^^^^^^^^^^^^^^^
# node1 node1-vip
# node2 node2-vip
#
# For Standalone Extended Cluster:
# ^^^^^^^^^^^^^^^^^^^^^^^^
# node1 node1-vip sitea
# node2 node2-vip siteb
#
```

Installing Only the Oracle Grid Infrastructure Software

This installation option requires manual postinstallation steps to enable the Oracle Grid Infrastructure software.

If you use the **Set Up Software Only** option during installation, then Oracle Universal Installer (OUI) installs the software binaries on multiple nodes. You can then perform the additional steps of configuring Oracle Clusterware and Oracle ASM.

- Installing Software Binaries for Oracle Grid Infrastructure for a Cluster
  You can install Oracle Grid Infrastructure software binaries for a cluster software on multiple nodes at a time.

- Configuring Software Binaries for Oracle Grid Infrastructure for a Cluster
  Configure the software binaries by starting Oracle Grid Infrastructure configuration wizard in GUI mode.

- Configuring the Software Binaries Using a Response File
  When you install or copy Oracle Grid Infrastructure software on any node, you can defer configuration for a later time. Review this procedure for completing configuration after the software is installed or copied on nodes, using the configuration wizard (`gridSetup.sh`).

- Setting Ping Targets for Network Checks
  Receive notification about network status by setting the `Ping_Targets` parameter during the Oracle Grid Infrastructure installation.

> **✎ See Also:**
>
> *Oracle Clusterware Administration and Deployment Guide* for information about cloning an Oracle Grid Infrastructure installation to other nodes that were not included in the initial installation of Oracle Grid Infrastructure, and then adding them to the cluster

Installing Software Binaries for Oracle Grid Infrastructure for a Cluster

You can install Oracle Grid Infrastructure software binaries for a cluster software on multiple nodes at a time.

1. Download the Grid home image files.

2. Run the `gridSetup.sh` command and select the Configuration Option as **Set Up Software Only**.

> **✎ Note:**
>
> You can use the `gridSetup.sh` command with the `-applyRU` and `-applyOneOffs` flags to install Release Updates (RUs) and one-off patches during an Oracle Grid Infrastructure installation or upgrade.

3. Complete installation of Oracle Grid Infrastructure software on one or more nodes by providing information in the installer screens in response to your configuration selection. You can install Oracle Grid Infrastructure software on multiple nodes at a time.

4. When the software is configured, run the `orainstRoot.sh` script on all nodes, when prompted.

5. On all nodes, the `root.sh` script output provides information about how to proceed, depending on the configuration you plan to complete in this installation. Make note of this information.

6. Ensure that you have completed all storage and server preinstallation requirements.

7. Verify that all of the cluster nodes meet the installation requirements:

    ```
    runcluvfy.sh stage -pre crsinst -n node_list
    ```

8. Configure the cluster using the Oracle Universal Installer (OUI) configuration wizard or response files.

Configuring Software Binaries for Oracle Grid Infrastructure for a Cluster

Configure the software binaries by starting Oracle Grid Infrastructure configuration wizard in GUI mode.

1. Log in on a cluster node as the Oracle Grid Infrastructure installation owner, and change directory to *Grid_home*.

2. Start the Oracle Grid Infrastructure configuration wizard:

    ```
    $ ./gridSetup.sh
    ```

3. Provide information as needed for configuration. OUI validates the information and configures the installation on all cluster nodes.

4. When you complete providing information, OUI shows you the Summary page, listing the information you have provided for the cluster. Verify that the summary has the correct information for your cluster, and click **Install** to start configuration of the local node.

    When configuration of the local node is complete, OUI copies the Oracle Grid Infrastructure configuration file to other cluster member nodes.

5. When prompted, run `root` scripts.

6. When you confirm that all `root` scripts are run, OUI checks the cluster configuration status, and starts other configuration tools as needed.

Configuring the Software Binaries Using a Response File

When you install or copy Oracle Grid Infrastructure software on any node, you can defer configuration for a later time. Review this procedure for completing configuration after the software is installed or copied on nodes, using the configuration wizard (`gridSetup.sh`).

To configure the Oracle Grid Infrastructure software binaries using a response file:

1. As the Oracle Grid Infrastructure installation owner (`grid`), start Oracle Universal Installer in Oracle Grid Infrastructure configuration wizard mode from the Oracle

Grid Infrastructure software-only home using the following syntax, where *filename* is the response file name:

```
/u01/app/21.0.0/grid/gridSetup.sh [-debug] [-silent -responseFile
filename]
```

For example:

```
$ cd /u01/app/21.0.0/grid
$ ./gridSetup.sh -responseFile /u01/app/21.0.0/grid/install/response/
response_file.rsp
```

2. When you complete configuring values, OUI shows you the Summary page, listing all information you have provided for the cluster. Verify that the summary has the correct information for your cluster, and click **Install** to start configuration of the local node.

    When configuration of the local node is complete, OUI copies the Oracle Grid Infrastructure configuration file to other cluster member nodes.

3. When prompted, run root scripts.

4. When you confirm that all root scripts are run, OUI checks the cluster configuration status, and starts other configuration tools as needed.

Setting Ping Targets for Network Checks

Receive notification about network status by setting the `Ping_Targets` parameter during the Oracle Grid Infrastructure installation.

For environments where the network link status is not correctly returned when the network cable is disconnected, for example, in a virtual machine, you can receive notification about network status by setting the `Ping_Targets` parameter during the Oracle Grid Infrastructure installation.

Run the installer:

```
./gridSetup.sh oracle_install_crs_Ping_Targets=Host1|IP1,Host2|IP2
```

The ping utility contacts the comma-separated list of host names or IP addresses `Host1| IP1,Host2|IP2` to determine whether the public network is available. If none of the hosts respond, then the network is considered to be offline. Addresses outside the cluster, like of a switch or router, should be used.

For example:

```
/gridSetup.sh oracle_install_crs_Ping_Targets=192.0.2.1,192.0.2.2
```

Confirming Oracle Clusterware Function

After Oracle Grid Infrastructure installation, confirm that your Oracle Clusterware installation is installed and running correctly.

After installation, log in as `root`, and use the following command syntax to confirm that your Oracle Clusterware installation is installed and running correctly:

```
crsctl check cluster -all
```

For example:

```
$ crsctl check cluster -all

**************************************************************
node1:
CRS-4537: Cluster Ready Services is online
CRS-4529: Cluster Synchronization Services is online
CRS-4533: Event Manager is online
**************************************************************
node2:
CRS-4537: Cluster Ready Services is online
CRS-4529: Cluster Synchronization Services is online
CRS-4533: Event Manager is online
**************************************************************
node3:
CRS-4537: Cluster Ready Services is online
CRS-4529: Cluster Synchronization Services is online
CRS-4533: Event Manager is online
**************************************************************
```

> **Note:**
>
> After installation is complete, do not remove manually or run `cron` jobs that remove `/tmp/.oracle` or `/var/tmp/.oracle` or its files while Oracle Clusterware is up. If you remove these files, then Oracle Clusterware could encounter intermittent hangs, and you will encounter error `CRS-0184: Cannot communicate with the CRS daemon.`

Confirming Oracle ASM Function for Oracle Clusterware Files

Confirm Oracle ASM is running after installing Oracle Grid Infrastructure.

After Oracle Grid Infrastructure installation, Oracle Clusterware files are stored on Oracle ASM. Use the following command syntax as the Oracle Grid Infrastructure installation owner (`grid`) to confirm that your Oracle ASM installation is running:

```
srvctl status asm
```

For example:

```
srvctl status asm
ASM is running on node1,node2, node3, node4
```

> **✎ Note:**
>
> To manage Oracle ASM or Oracle Net, use the `srvctl` binary in the Oracle Grid Infrastructure home for a cluster (Grid home). If you have Oracle Real Application Clusters or Oracle Database installed, then you cannot use the `srvctl` binary in the database home to manage Oracle ASM or Oracle Net.

Understanding Offline Processes in Oracle Grid Infrastructure

After the installation of Oracle Grid Infrastructure, some components may be listed as OFFLINE. Oracle Grid Infrastructure activates these resources when you choose to add them.

Oracle Grid Infrastructure provides required resources for various Oracle products and components. Some of those products and components are optional, so you can install and enable them after installing Oracle Grid Infrastructure. To simplify postinstall additions, Oracle Grid Infrastructure preconfigures and registers all required resources for all products available for these products and components, but only activates them when you choose to add them. As a result, some components may be listed as OFFLINE after the installation of Oracle Grid Infrastructure. Run the following command to view status of any resource:

```
$ crsctl status resource resource_name -t
```

Resources listed as TARGET:OFFLINE and STATE:OFFLINE do not need to be monitored. They represent components that are registered, but not enabled, so they do not use any system resources. If an Oracle product or component is installed on the system, and it requires a particular resource to be online, then the software prompts you to activate the required offline resource.

9

Oracle Grid Infrastructure Postinstallation Tasks

Complete configuration tasks after you install Oracle Grid Infrastructure.

You are required to complete some configuration tasks after Oracle Grid Infrastructure is installed. In addition, Oracle recommends that you complete additional tasks immediately after installation. You must also complete product-specific configuration tasks before you use those products.

> **Note:**
>
> This chapter describes basic configuration only. Refer to product-specific administration and tuning guides for more detailed configuration and tuning information.

- Required Postinstallation Tasks
  Download and apply required patches for your software release after completing your initial installation.

- Recommended Postinstallation Tasks
  Oracle recommends that you complete these tasks after installation.

- Installing Grid Infrastructure Management Repository
  Install and configure Grid Infrastructure Management Repository (GIMR) into a separate Oracle home in this release of Oracle Grid Infrastructure.

- Using Earlier Oracle Database Releases with Oracle Grid Infrastructure
  Review the following topics for information about using earlier Oracle Database releases with Oracle Grid Infrastructure 21c installations:

- Modifying Oracle Clusterware Binaries After Installation
  After installation, if you need to modify the Oracle Clusterware configuration, then you must unlock the Grid home. Review this information about unlocking the Grid home.

Required Postinstallation Tasks

Download and apply required patches for your software release after completing your initial installation.

- Downloading Release Update Patches
  Download and install Release Updates (RU) and Release Update Revisions (RUR) patches for your Oracle software after you complete installation.

- Setting External Jobs Ownership for Installations on HP-UX
  On HP-UX platforms only, complete the following procedure to set external jobs ownership to the low-privilege user `extjob`:

Downloading Release Update Patches

Download and install Release Updates (RU) and Release Update Revisions (RUR) patches for your Oracle software after you complete installation.

Starting with Oracle Database 18c, Oracle provides quarterly updates in the form of Release Updates (RU) and Release Update Revisions (RUR). Oracle no longer releases patch sets. For more information, see My Oracle Support Note 2285040.1.

Check the My Oracle Support website for required updates for your installation.

1. Use a web browser to view the My Oracle Support website:

   https://support.oracle.com

2. Log in to My Oracle Support website.

   > **✎ Note:**
   >
   > If you are not a My Oracle Support registered user, then click **Register for My Oracle Support** and register.

3. On the main My Oracle Support page, click **Patches & Updates**.

4. In the **Patch Search** region, select **Product or Family (Advanced)**.

5. On the **Product or Family (Advanced)** display, provide information about the product, release, and platform for which you want to obtain patches, and click **Search**.

   The Patch Search pane opens, displaying the results of your search.

6. Select the patch number and click **ReadMe**.

   The README page is displayed. It contains information about the patch and how to apply the patches to your installation.

7. Uncompress the Oracle patch updates that you downloaded from My Oracle Support.

Related Topics

• My Oracle Support note 2285040.1

Setting External Jobs Ownership for Installations on HP-UX

On HP-UX platforms only, complete the following procedure to set external jobs ownership to the low-privilege user `extjob`:

1. Log on as `root`.

2. Change directory to the Oracle Database Oracle home:

   ```
   # cd $ORACLE_HOME/rdbms/admin/
   ```

3. Open `externaljob.ora` with a text editor, and find the parameters `run_user` and `run_group`.

4. Set `run_user` to the external jobs user (`extjob`), and set `run_group` to a low-privileged group, such as `other`. For example:

```
run_user=extproc
run_group=other
```

5. Save the file.

> 📝 **Note:**
>
> Modify `externaljob.ora` only as `root`.

Recommended Postinstallation Tasks

Oracle recommends that you complete these tasks after installation.

- Creating a Backup of the root.sh Script
  Oracle recommends that you back up the `root.sh` script after you complete an installation.

- About Installing Oracle Autonomous Health Framework
  Oracle Autonomous Health Framework is pre-installed with Oracle Database.

- Creating a Fast Recovery Area
  During an Oracle Restart installation, you can create only one disk group. During an Oracle Clusterware installation, you can create multiple disk groups. If you plan to add an Oracle Database for a standalone server or an Oracle RAC database, then you should create the fast recovery area for database files.

- Checking the SCAN Configuration
  The Single Client Access Name (SCAN) is a name that is used to provide service access for clients to the cluster.

- Setting Resource Limits for Oracle Clusterware and Associated Databases and Applications
  After you have completed Oracle Grid Infrastructure installation, you can set resource limits in the $Grid\_home$/crs/install/s_crsconfig_$nodename$_env.txt file.

Creating a Backup of the root.sh Script

Oracle recommends that you back up the `root.sh` script after you complete an installation.

If you install other products in the same Oracle home directory subsequent to this installation, then Oracle Universal Installer updates the contents of the existing `root.sh` script during the installation. If you require information contained in the original `root.sh` script, then you can recover it from the backed up `root.sh` file.

About Installing Oracle Autonomous Health Framework

Oracle Autonomous Health Framework is pre-installed with Oracle Database.

Oracle Autonomous Health Framework performs proactive heath checks and collects diagnostics data for the Oracle software stack. Oracle Autonomous Health Framework updates are available on My Oracle Support and also shipped with Oracle Database Release Updates.

Oracle recommends that you update to the latest version of Oracle Autonomous Health Framework by either applying the latest Database Release Updates or by downloading and installing it from My Oracle Support Note 2550798.1:

https://support.oracle.com/epmos/faces/DocContentDisplay?
id=2550798.1&parent=DOCUMENTATION&sourceId=USERGUIDE

Oracle Autonomous Health Framework includes the functionality from Oracle ORAchk, Oracle EXAchk, and Oracle Trace File Analyzer (TFA). Oracle Autonomous Health Framework extends health check coverage to the entire Oracle software stack, based on critical and reoccurring problems. Oracle Autonomous Health Framework proactively scans for known problems with Oracle products and deployments, including the following:

- Single-instance Oracle Database

- Oracle Grid Infrastructure

- Oracle Real Application Clusters

- Maximum Availability Architecture (MAA) Validation

- Upgrade Readiness Validations

- Oracle GoldenGate

- Oracle Exadata

- Oracle SuperCluster

- Oracle Zero Data Loss Recovery Appliance

Related Topics

- *Oracle Autonomous Health Framework Checks and Diagnostics User's Guide*

Creating a Fast Recovery Area

During an Oracle Restart installation, you can create only one disk group. During an Oracle Clusterware installation, you can create multiple disk groups. If you plan to add an Oracle Database for a standalone server or an Oracle RAC database, then you should create the fast recovery area for database files.

- About the Fast Recovery Area and the Fast Recovery Area Disk Group
  The fast recovery area is a unified storage location for all Oracle Database files related to recovery. Enabling rapid backups for recent data can reduce requests to system administrators to retrieve backup tapes for recovery operations.

- Creating the Fast Recovery Area Disk Group
  Procedure to create the fast recovery area disk group.

About the Fast Recovery Area and the Fast Recovery Area Disk Group

The fast recovery area is a unified storage location for all Oracle Database files related to recovery. Enabling rapid backups for recent data can reduce requests to system administrators to retrieve backup tapes for recovery operations.

Database administrators can define the DB_RECOVERY_FILE_DEST parameter to the path for the fast recovery area to enable on disk backups and rapid recovery of data. When you enable fast recovery in the init.ora file, Oracle Database writes all RMAN backups, archive logs, control file automatic backups, and database copies to the fast

recovery area. RMAN automatically manages files in the fast recovery area by deleting obsolete backups and archiving files no longer required for recovery.

Oracle recommends that you create a fast recovery area disk group. Oracle Clusterware files and Oracle Database files can be placed on the same disk group, and you can also place fast recovery files in the same disk group. However, Oracle recommends that you create a separate fast recovery disk group to reduce storage device contention.

The fast recovery area is enabled by setting the `DB_RECOVERY_FILE_DEST` parameter. The size of the fast recovery area is set with `DB_RECOVERY_FILE_DEST_SIZE`. As a general rule, the larger the fast recovery area, the more useful it becomes. For ease of use, Oracle recommends that you create a fast recovery area disk group on storage devices that can contain at least three days of recovery information. Ideally, the fast recovery area is large enough to hold a copy of all of your data files and control files, the online redo logs, and the archived redo log files needed to recover your database using the data file backups kept under your retention policy.

Multiple databases can use the same fast recovery area. For example, assume you have created a fast recovery area disk group on disks with 150 GB of storage, shared by 3 different databases. You can set the size of the fast recovery for each database depending on the importance of each database. For example, if database1 is your least important database, database2 is of greater importance, and database3 is of greatest importance, then you can set different `DB_RECOVERY_FILE_DEST_SIZE` settings for each database to meet your retention target for each database: 30 GB for database1, 50 GB for database2, and 70 GB for database3.

Creating the Fast Recovery Area Disk Group

Procedure to create the fast recovery area disk group.

1. Go to the Oracle Grid Infrastructure home `bin` directory, and start Oracle ASM Configuration Assistant (ASMCA).

   For example:

   ```
   $ cd /u01/app/21.0.0/grid/bin
   $ ./asmca
   ```

   ASMCA opens the home window.

2. Click **Disk Groups** in the left panel to open the Disk Groups tab.

3. Click **Create** to create a new disk group.

   The Create Disk Group window opens.

4. Provide configuration information for the fast recovery area as prompted:

   In the Disk Group Name field, enter a descriptive name for the fast recovery area group. For example: `FRA`.

   In the Redundancy section, select the level of redundancy you want to use. For example: `Normal`

   In the Select Member Disks field, select eligible disks you want to add to the fast recovery area, and click **OK**. If you do not see the disks, click the **Change Disk Discovery Path** button and provide a path and pattern match for the disk. For example, `/dev/sd*` for local Oracle ASM disks.

5. When the Fast Recovery Area disk group creation is complete, click **Exit** and click **Yes** to confirm closing the ASMCA application.

Checking the SCAN Configuration

The Single Client Access Name (SCAN) is a name that is used to provide service access for clients to the cluster.

The SCAN makes it possible to add or remove nodes from the cluster without needing to reconfigure clients because the SCAN is associated with the cluster as a whole, rather than to a particular node.

The Single Client Access Name (SCAN) also adds location independence for the databases, so that client configuration does not have to depend on which nodes are running a particular database instance. Clients can continue to access the cluster in the same way as with previous releases, but Oracle recommends that clients accessing the cluster use the SCAN.

You can use the command `cluvfy comp scan` (located in *Grid home*/bin) to confirm that the DNS is correctly associating the SCAN with the addresses. For example:

```
$cluvfy comp scan

Verifying Single Client Access Name (SCAN) ...
  Verifying DNS/NIS name service 'rws127064-clu-scan.rws127064-
clu.rws12706410644.example.com' ...
    Verifying Name Service Switch Configuration File Integrity ...PASSED
  Verifying DNS/NIS name service 'rws127064-clu-scan.rws127064-
clu.rws12706410644.example.com' ...PASSED
Verifying Single Client Access Name (SCAN) ...PASSED

Verification of SCAN was successful.

CVU operation performed: SCAN
Date: Jul 29, 2016 1:42:41 AM
CVU home: /u01/crshome/
User: crsusr
```

After installation, when a client sends a request to the cluster, the Oracle Clusterware SCAN listeners redirect client requests to servers in the cluster.

> **✎ See Also:**
>
> *Oracle Clusterware Administration and Deployment Guide* for more information about system checks and configurations

Setting Resource Limits for Oracle Clusterware and Associated Databases and Applications

After you have completed Oracle Grid Infrastructure installation, you can set resource limits in the *Grid_home*/crs/install/s_crsconfig_*nodename*_env.txt file.

The resource limits apply to all Oracle Clusterware processes and Oracle databases managed by Oracle Clusterware. For example, to set a higher number of processes limit, edit the file and set the CRS_LIMIT_NPROC parameter to a high value.

```
---
#Do not modify this file except as documented above or under the
#direction of Oracle Support Services.
######################################################################
TZ=PST8PDT
NLS_LANG=AMERICAN_AMERICA.WE8ISO8859P1
CRS_LIMIT_STACK=2048
CRS_LIMIT_OPENFILE=65536
CRS_LIMIT_NPROC=65536
TNS_ADMIN=
```

Installing Grid Infrastructure Management Repository

Install and configure Grid Infrastructure Management Repository (GIMR) into a separate Oracle home in this release of Oracle Grid Infrastructure.

- About Grid Infrastructure Management Repository Installation
  You can install Grid Infrastructure Management Repository (GIMR) for Oracle Standalone Cluster and you must install GIMR for Oracle Domain Services Cluster.

- Creating the GIMR Disk Group
  Create an Oracle Automatic Storage Management (Oracle ASM) disk group to host the Grid Infrastructure Management Repository (GIMR) for a cluster.

- Installing a Local GIMR for a New Cluster
  Complete this procedure to install a local Grid Infrastructure Management Repository (GIMR) for Oracle Standalone Cluster and Oracle Domain Services Cluster.

- Adding a Remote Oracle Domain Service Cluster GIMR to an Existing Cluster
  Complete this procedure to add a remote Oracle Domain Services Cluster Grid Infrastructure Management Repository (GIMR) to an Oracle Standalone Cluster that was initially configured without a GIMR.

- Converting a Local Oracle Standalone Cluster GIMR to a Remote GIMR
  Complete this procedure to convert a local Oracle Standalone Cluster GIMR to a remote GIMR on an Oracle Standalone Cluster or an Oracle Domain Services Cluster.

About Grid Infrastructure Management Repository Installation

You can install Grid Infrastructure Management Repository (GIMR) for Oracle Standalone Cluster and you must install GIMR for Oracle Domain Services Cluster.

Starting with Oracle Grid Infrastructure 21c, the Grid Infrastructure Management Repository will be configured in a separate Oracle home, instead of an Oracle Grid Infrastructure home (Grid home). This option is available if you choose to configure GIMR during a fresh installation or you add a GIMR to an existing Oracle Grid Infrastructure deployment.

> **✎ Note:**
>
> Installing GIMR is optional for Oracle Standalone Clusters, but GIMR installation is mandatory for Oracle Domain Services Clusters.

Install and configure Grid Infrastructure Management Repository in a new Oracle home after installing and configuring Oracle Grid infrastructure. After you install Oracle Grid Infrastructure and choose to install GIMR, you need to download and install Oracle Database software of the same version as Oracle Grid Infrastructure.

If you choose to install GIMR during Oracle Grid Infrastructure installation, then you are prompted to create a new Oracle ASM disk group for GIMR and your decision is recorded in Oracle Cluster Registry (OCR). Cluster Verification Utility (CVU) creates a log to remind you to install GIMR. However, if you do not choose to install GIMR during Oracle Grid Infrastructure installation, then CVU does not remind you to create a GIMR.

Creating the GIMR Disk Group

Create an Oracle Automatic Storage Management (Oracle ASM) disk group to host the Grid Infrastructure Management Repository (GIMR) for a cluster.

If you have not created an Oracle ASM disk group for GIMR during Oracle Grid Infrastructure installation, then follow these steps to create a new Oracle ASM disk group using Oracle ASMCA.

1. Go to the Oracle Grid Infrastructure home `bin` directory, and start Oracle ASM Configuration Assistant (ASMCA).

   ```
   $ cd $ORACLE_HOME/bin
   $ ./asmca
   ```

   ASMCA home window opens.

2. Click **Disk Groups** in the left panel to open the Disk Groups tab.

3. Click **Create** to create a new disk group.

   The Create Disk Groups window opens.

4. Provide configuration information for the GIMR disk group as prompted:

   In the Disk Group Name field, enter a descriptive name for the GIMR disk group. For example: `MGMT`.

   In the Redundancy section, select the level of redundancy you want to use. For example: `External`

   In the Select Member Disks field, select eligible disks you want to add to the MGMT disk group, and click **OK**.

> **Note:**
>
> Determine the minimum storage space that you require for the GIMR and make sure the disk group meets the minimum storage space requirements for your GIMR.

5. When the GIMR disk group creation is complete, click **Exit** and click **Yes** to confirm closing the ASMCA application.

Installing a Local GIMR for a New Cluster

Complete this procedure to install a local Grid Infrastructure Management Repository (GIMR) for Oracle Standalone Cluster and Oracle Domain Services Cluster.

If you want to place GIMR on an Oracle ASM disk group, then create an Oracle ASM disk group for GIMR either during Oracle Grid Infrastructure installation or using Oracle ASMCA before installing the GIMR.

1. Ensure that you can access other nodes with SSH. The installer requires that the `grid` user is permitted to set up passwordless SSH. The installer can set this up for you automatically, or your system administrator can set this up for you before installation is started.

2. As the `grid` user, create GIMR home directory on all cluster nodes:

```
$ mkdir -p /u01/app/21.0.0/gimr
$ chown grid:oinstall /u01/app/21.0.0/gimr
```

3. Download the Oracle Database image files and extract the files into the GIMR home. For example:

```
$ cd /u01/app/21.0.0/gimr
$ unzip -q download_location/db_home.zip
```

`db_home.zip` is the name of the Oracle Database image zip file.

> **Note:**
>
> • You must extract the zip image software into the directory where you want your GIMR home to be located.
>
> • Download and copy the Oracle Database image files to the local node only. During installation, the software is copied and installed on all other nodes in the cluster.

4. From the GIMR home directory, start the Oracle Database software installation:

```
$ cd /u01/app/21.0.0/gimr
$ ./runInstaller
```

> **✎ Note:**
>
> Run the `runInstaller` command from the GIMR home directory only. Do not run `runInstaller` command that resides at `$GIMR_HOME/oui/bin/`, or any other location, to install Oracle Real Application Clusters.

5. In the Select Configuration Option screen, select the **Setup Software Only** option. Click **Next**.

6. In the Select Database Installation Option screen, select the **Oracle Real Application Clusters Database Installation** option. Click **Next**.

7. In the Node Selection screen, select all the nodes to install Oracle RAC or Oracle RAC One Node software. Click **Next**.

8. In the Select Database Edition screen, select the **Enterprise Edition** option. Click **Next**.

9. In the Specify Install Location screen, provide the location of Oracle base for Oracle RAC software installation.

   The Oracle base directory must be different from the GIMR home directory. Click **Next**.

10. In the Privileged Operating System Groups screen, accept the default operating system group names for Oracle Database administration and click **Next**.

11. Select the option to **Automatically run configuration scripts**. Enter the credentials for the `root` user or a `sudo` account, then click **Next**.

    Alternatively, you can run the scripts manually as the `root` user at the end of the installation process when prompted by the installer.

12. The Perform Prerequisite Checks screen displays the results of the prerequisites checks. If any of the checks have a status of Failed and are not Fixable, then you must manually correct these issues. After you have fixed the issue, you can click the Check Again button to have the installer recheck the requirement and update the status. Repeat as needed until all the checks have a status of Succeeded. Click **Next**.

13. Review the contents of the Summary screen and then click **Install**.

    The installer displays a progress indicator enabling you to monitor the installation process.

14. If you did not configure automation of the root scripts, then you are required to run `root.sh` script as the `root` user, as specified in the Execute Configuration Scripts window. Do *not* click **OK** until you have run the `root.sh` script on all nodes as directed.

    You can run the `root.sh` script simultaneously on all nodes in the cluster for Oracle RAC installations or upgrades.

15. Create GIMR container database using the following command:

```
$ /u01/app/21.0.0/gimr/bin/mgmtca createGIMRContainer [-
storageDiskLocation disk_location]
```

Adding a Remote Oracle Domain Service Cluster GIMR to an Existing Cluster

Complete this procedure to add a remote Oracle Domain Services Cluster Grid Infrastructure Management Repository (GIMR) to an Oracle Standalone Cluster that was initially configured without a GIMR.

1. As the `grid` user on the Oracle Standalone Cluster, get the cluster Globally Unique Identifier (GUID) of the Oracle Standalone Cluster to which you want to add a GIMR.

   ```
   $ Grid_home/bin/crsctl get css clusterguid
   ```

2. As the `grid` user on the Oracle Domain Services Cluster, create a GIMR pluggable database (PDB) for the Oracle Standalone Cluster and generate the client data file.

   ```
   $ /u01/app/21.0.0/gimr/bin/mgmtca createRepos
   -clientDataFile client_data_file_location -clusterName
   standalone_cluster_name
   -guid standalone_cluster_guid [-version standalone_cluster_version]
   ```

3. Copy the client data file from the Oracle Domain Services Cluster to the Oracle Standalone Cluster.

4. As the `grid` user on the Oracle Standalone Cluster, configure the remote GIMR.

   ```
   $ Grid_home/bin/mgmtca configRepos deploySchema –clientDataFile
   client_data_file_location
   ```

Converting a Local Oracle Standalone Cluster GIMR to a Remote GIMR

Complete this procedure to convert a local Oracle Standalone Cluster GIMR to a remote GIMR on an Oracle Standalone Cluster or an Oracle Domain Services Cluster.

Ensure that the Oracle Standalone Cluster or Oracle Domain Services Cluster server version is equal to or greater than the Oracle Standalone Cluster client version.

1. As the `grid` user on the Oracle Standalone Cluster client with a local Grid Infrastructure Management Repository (GIMR), unplug the GIMR pluggable database (PDB) and export the cluster property information to a file.

   ```
   $ /u01/app/21.0.0/gimr/bin/mgmtca unplugRepos -local
   -serverVersion remote_GIMR_server_version -clusterInfoDir local_directory
   ```

   Where *local_directory* is a directory to store cluster property information file, GIMR PDB backup, and GIMR PDB metadata information.

2. Copy the contents of the `-clusterInfoDir` directory from the Oracle Standalone Cluster client to the Oracle Domain Services Cluster or the Oracle Standalone Cluster to which you want to migrate the GIMR.

3. As the `grid` user on the Oracle Domain Services Cluster or the Oracle Standalone Cluster to which you want to migrate the GIMR, plug the Oracle Standalone Cluster GIMR PDB and generate the client data file.

```
$ /u01/app/21.0.0/gimr/bin/mgmtca plugRepos -clientDataFile
client_data_file_location
-clusterInfoDir cluster.properties_file_location
```

If this step fails, then run the `/u01/app/21.0.0/gimr/bin/mgmtca plugRepos -local -clusterInfoDir cluster.properties_file_location` command on the Oracle Standalone Cluster client to rollback.

4. Copy the client data file from the Oracle Domain Services Cluster or the Oracle Standalone Cluster to which you have migrated the GIMR to the Oracle Standalone Cluster client.

5. As the `grid` user on the Oracle Standalone Cluster client, import credentials from the client data file.

```
$ Grid_home/bin/mgmtca configRepos importCred -clientDataFile
    client_data_file_location
```

6. As the `grid` user on the Oracle Standalone Cluster client, delete the local GIMR.

```
$ /u01/app/21.0.0/gimr/dbca -deleteDatabase -sourceDB -MGMTDB -
silent -skipDBExistencecheck
```

7. As the `grid` user on the Oracle Standalone Cluster client, delete the management listener.

```
$ Grid_home/bin/srvctl disable mgmtlsnr
$ Grid_home/bin/srvctl stop mgmtlsnr
$ Grid_home/bin/srvctl remove mgmtlsnr
```

Using Earlier Oracle Database Releases with Oracle Grid Infrastructure

Review the following topics for information about using earlier Oracle Database releases with Oracle Grid Infrastructure 21c installations:

- General Restrictions for Using Earlier Oracle Database Releases
  You can use Oracle Database 21c, 19c, 18c, and Oracle Database 12c release 2 with Oracle Grid Infrastructure 21c.

- Making Oracle ASM Available to Earlier Oracle Database Releases
  To use Oracle ASM with Oracle Database releases earlier than Oracle Database 12c Release 2 (12.2), you must pin all the cluster nodes.

- Using ASMCA to Administer Disk Groups for Earlier Database Releases
  Use Oracle ASM Configuration Assistant (ASMCA) to create and modify disk groups when you install earlier Oracle databases and Oracle RAC databases on Oracle Grid Infrastructure installations.

- • Using the Correct LSNRCTL Commands
  To administer Oracle Database 21c local and scan listeners using the `lsnrctl` command, set your `$ORACLE_HOME` environment variable to the path for the Oracle Grid Infrastructure home (Grid home).

General Restrictions for Using Earlier Oracle Database Releases

You can use Oracle Database 21c, 19c, 18c, and Oracle Database 12c release 2 with Oracle Grid Infrastructure 21c.

Do not use the versions of `srvctl`, `lsnrctl`, or other Oracle Grid infrastructure home tools to administer earlier version databases. Only administer earlier Oracle Database releases using the tools in the earlier Oracle Database homes. To ensure that the versions of the tools you are using are the correct tools for those earlier release databases, run the tools from the Oracle home of the database or object you are managing.

Making Oracle ASM Available to Earlier Oracle Database Releases

To use Oracle ASM with Oracle Database releases earlier than Oracle Database 12c Release 2 (12.2), you must pin all the cluster nodes.

1. After you install Oracle Grid Infrastructure 18c or later release, if you want to use Oracle ASM to provide storage service for Oracle Database releases that are earlier than Oracle Database 12c Release 2 (12.2), then you must use the following command to pin the nodes:

   ```
   $ crsctl pin css -n node1 node2
   ```

   This setting updates the `oratab` file for Oracle ASM entries.

2. Check the pinned nodes using the following command:

   $ ./olsnodes -t -n

   > ✎ **Note:**
   >
   > Restart Oracle ASM to load the updated `oratab` file.

Using ASMCA to Administer Disk Groups for Earlier Database Releases

Use Oracle ASM Configuration Assistant (ASMCA) to create and modify disk groups when you install earlier Oracle databases and Oracle RAC databases on Oracle Grid Infrastructure installations.

Starting with Oracle Database 11g Release 2, Oracle ASM is installed as part of an Oracle Grid Infrastructure installation, with Oracle Clusterware. You can no longer use Database Configuration Assistant (DBCA) to perform administrative tasks on Oracle ASM.

> **See Also:**
>
> *Oracle Automatic Storage Management Administrator's Guide* for details
> about configuring disk group compatibility for databases using Oracle
> Database 11g or earlier software with Oracle Grid Infrastructure 21c.

Using the Correct LSNRCTL Commands

To administer Oracle Database 21c local and scan listeners using the `lsnrctl`
command, set your `$ORACLE_HOME` environment variable to the path for the Oracle Grid
Infrastructure home (Grid home).

Do not attempt to use the `lsnrctl` commands from Oracle home locations for previous
releases, as they cannot be used with the new release.

Modifying Oracle Clusterware Binaries After Installation

After installation, if you need to modify the Oracle Clusterware configuration, then you
must unlock the Grid home. Review this information about unlocking the Grid home.

For example, if you want to apply a one-off patch, or if you want to modify an Oracle
Exadata configuration to run IPC traffic over RDS on the interconnect instead of using
the default UDP, then you must unlock the Grid home.

> **⚠ Caution:**
>
> Before relinking executables, you must shut down all executables that run in
> the Oracle home directory that you are relinking. In addition, shut down
> applications linked with Oracle shared libraries.

1. Go to the Oracle Grid Infrastructure home directory, and unlock the Grid home
   using the command `rootcrs.sh -unlock`:

   ```
   # cd /u01/app/21.0.0/grid/crs/install
   # ./rootcrs.sh -unlock
   ```

2. Change user to the Oracle Grid Infrastructure software owner, and relink binaries
   using the command syntax `make -f /u01/app/21.0.0/grid/rdbms/lib/`
   `ins_rdbms.mk` *target*, where *target* is the binaries that you want to relink. For
   example, where you are updating the interconnect protocol from UDP to IPC, enter
   the following command:

   ```
   # su grid
   $ make -f /u01/app/21.0.0/grid/rdbms/lib/ins_rdbms.mk ipc_rds
   ioracle
   ```

> **Note:**
>
> To relink binaries, you can also change to the grid installation owner and run the command `/u01/app/21.0.0/grid/bin/relink`.

3. Relock the Grid home and restart the cluster as follows:

```
# ./rootcrs.sh -lock
# crsctl start crs
```

Repeat steps 1 through 3 on each cluster member node.

> **Note:**
>
> Do not delete directories in the Grid home. For example, do not delete the directory *Grid_home*/OPatch. If you delete the directory, then the Grid infrastructure installation owner cannot use OPatch to patch the Grid home, and OPatch displays the error message "`checkdir error: cannot create` *Grid_home*`/OPatch`".

10
Upgrading Oracle Grid Infrastructure

Oracle Grid Infrastructure upgrade consists of upgrade of Oracle Clusterware and Oracle Automatic Storage Management (Oracle ASM).

Oracle Grid Infrastructure upgrades can be rolling upgrades, in which a subset of nodes are brought down and upgraded while other nodes remain active.

- Understanding Out-of-Place Upgrade
  With an out-of-place upgrade, the installer installs the newer version in a separate Oracle Clusterware home.

- About Oracle Grid Infrastructure Upgrade and Downgrade
  You have the ability to upgrade or downgrade Oracle Grid Infrastructure to a supported release.

- Options for Oracle Grid Infrastructure Upgrades
  Review these upgrade options when you upgrade to Oracle Grid Infrastructure 21c.

- Restrictions for Oracle Grid Infrastructure Upgrades
  Review the following information for restrictions and changes for upgrades to Oracle Grid Infrastructure installations, which consists of Oracle Clusterware and Oracle Automatic Storage Management (Oracle ASM).

- Preparing to Upgrade an Existing Oracle Clusterware Installation
  If you have an existing Oracle Clusterware installation, then you upgrade your existing cluster by performing an out-of-place upgrade. You cannot perform an in-place upgrade.

- Understanding Rolling Upgrades Using Batches
  You can perform rolling upgrades of Oracle Grid Infrastructure in batches.

- Performing Rolling Upgrade of Oracle Grid Infrastructure
  Review this information to perform rolling upgrade of Oracle Grid Infrastructure.

- Upgrading Grid Infrastructure Management Repository from an Earlier Release
  Complete this procedure to upgrade Grid Infrastructure Management Repository (GIMR) from an earlier release.

- Applying Patches to Oracle Grid Infrastructure
  After you have upgraded Oracle Grid Infrastructure 21c, you can install individual software patches by downloading them from My Oracle Support.

- Applying Patches to Grid Infrastructure Management Repository
  You can patch Grid Infrastructure Management Repository by installing in-place or out-of-place patches using the OPatch utility.

- Updating Oracle Enterprise Manager Cloud Control Target Parameters
  After upgrading Oracle Grid Infrastructure, upgrade the Enterprise Manager Cloud Control target.

- Unlocking and Deinstalling the Previous Release Grid Home
  After upgrading from previous releases, if you want to deinstall the previous release Oracle Grid Infrastructure home, then you must first change the permission and ownership of the previous release Grid home.

- • Checking Cluster Health Monitor Repository Size After Upgrading
  If you are upgrading Oracle Grid Infrastructure from a prior release using IPD/OS to the current release, then review the Cluster Health Monitor repository size (the CHM repository).

- • Downgrading Oracle Clusterware to an Earlier Release
  After a successful or a failed upgrade, you can downgrade Oracle Clusterware to the previous release.

- • Completing Failed or Interrupted Installations and Upgrades
  If Oracle Universal Installer (OUI) exits on the node from which you started the upgrade, or the node reboots before you confirm that the `rootupgrade.sh` script was run on all nodes, then the upgrade remains incomplete.

- • Converting to Oracle Extended Cluster After Upgrading Oracle Grid Infrastructure
  Review this information to convert to an Oracle Extended Cluster after upgrading Oracle Grid Infrastructure. Oracle Extended Cluster enables you to deploy Oracle RAC databases on a cluster, in which some of the nodes are located in different sites.

Understanding Out-of-Place Upgrade

With an out-of-place upgrade, the installer installs the newer version in a separate Oracle Clusterware home.

Rolling upgrade avoids downtime and ensure continuous availability while the software is upgraded to a new version. Both versions of Oracle Clusterware are on each cluster member node, but only one version is active.

If you have separate Oracle Clusterware homes on each node, then you can perform an out-of-place upgrade on all nodes, or perform an out-of-place rolling upgrade, so that some nodes are running Oracle Clusterware from the earlier version Oracle Clusterware home, and other nodes are running Oracle Clusterware from the new Oracle Clusterware home.

An in-place upgrade of Oracle Grid Infrastructure is not supported.

About Oracle Grid Infrastructure Upgrade and Downgrade

You have the ability to upgrade or downgrade Oracle Grid Infrastructure to a supported release.

You can upgrade Oracle Grid Infrastructure in any of the following ways:

- • Rolling Upgrade which involves upgrading individual nodes without stopping Oracle Grid Infrastructure on other nodes in the cluster.

- • Non-rolling Upgrade which involves bringing down all the nodes except one. A complete cluster outage occurs while the root script stops the old Oracle Clusterware stack and starts the new Oracle Clusterware stack on the node where you initiate the upgrade. After upgrade is completed, the new Oracle Clusterware is started on all the nodes.

Note that some services are disabled when one or more nodes are in the process of being upgraded. All upgrades are out-of-place upgrades, meaning that the software binaries are placed in a different Grid home from the Grid home used for the prior release.

You can downgrade from Oracle Grid Infrastructure 21c to Oracle Grid Infrastructure 19c, Oracle Grid Infrastructure 18c, and Oracle Grid Infrastructure 12c Release 2 (12.2). Be aware that if you downgrade to a prior release, then your cluster must conform with the configuration requirements for that prior release, and the features available for the cluster consist only of the features available for that prior release of Oracle Clusterware and Oracle ASM.

You can perform out-of-place upgrades to an Oracle ASM instance using Oracle ASM Configuration Assistant (ASMCA). In addition to running ASMCA using the graphical user interface, you can run ASMCA in non-interactive (silent) mode.

> **Note:**
>
> You must complete an upgrade before attempting to use cluster backup files. You cannot use backups for a cluster that has not completed upgrade.

> **See Also:**
>
> *Oracle Database Upgrade Guide* and *Oracle Automatic Storage Management Administrator's Guide* for additional information about upgrading existing Oracle ASM installations

Options for Oracle Grid Infrastructure Upgrades

Review these upgrade options when you upgrade to Oracle Grid Infrastructure 21c.

Supported upgrade paths for Oracle Grid Infrastructure for this release are:

- Oracle Grid Infrastructure upgrade from 12c Release 2 (12.2) to Oracle Grid Infrastructure 21c.

- Oracle Grid Infrastructure upgrade from 18c release to Oracle Grid Infrastructure 21c.

- Oracle Grid Infrastructure upgrade from 19c release to Oracle Grid Infrastructure 21c.

Upgrade options from Oracle Grid Infrastructure 12c Release 2 (12.2), Oracle Grid Infrastructure 18c, and Oracle Grid Infrastructure 19c to Oracle Grid Infrastructure 21c include the following:

- Oracle Grid Infrastructure rolling upgrade which involves upgrading individual nodes without stopping Oracle Grid Infrastructure on other nodes in the cluster

- Oracle Grid Infrastructure non-rolling upgrade by bringing the cluster down and upgrading the complete cluster

> **✎ Note:**
>
> - You can use either Oracle ASM or a shared file system to store OCR and voting files on Oracle Standalone Cluster deployments. If storage for OCR and voting files is other than Oracle ASM on other cluster types, then you need to migrate OCR and voting files to Oracle ASM before upgrading to Oracle Grid Infrastructure 21c.

Restrictions for Oracle Grid Infrastructure Upgrades

Review the following information for restrictions and changes for upgrades to Oracle Grid Infrastructure installations, which consists of Oracle Clusterware and Oracle Automatic Storage Management (Oracle ASM).

- Oracle Grid Infrastructure upgrades are always out-of-place upgrades. You cannot perform an in-place upgrade of Oracle Grid Infrastructure to existing homes.

- The same user that owned the earlier release Oracle Grid Infrastructure software must perform the Oracle Grid Infrastructure 21c upgrade.

- Oracle ASM and Oracle Clusterware both run in the Oracle Grid Infrastructure home.

- When you upgrade to Oracle Grid Infrastructure 12c Release 2 (12.2) or later, you upgrade to an Oracle Flex Cluster configuration.

- Do not delete directories in the Grid home. For example, do not delete the directory *Grid_home*/OPatch. If you delete the directory, then the Grid infrastructure installation owner cannot use OPatch utility to patch the grid home, and OPatch displays the error message "'checkdir' error: cannot create *Grid_home*/OPatch".

- To upgrade existing Oracle Grid Infrastructure installations to Oracle Grid Infrastructure 21c, you must first verify if you need to apply any mandatory patches for upgrade to succeed.

  Oracle recommends that you use the Cluster Verification Utility tool (CVU) to check if there are any patches required for upgrading your existing Oracle Grid Infrastructure or Oracle RAC database installations. See *Using CVU to Validate Readiness for Oracle Clusterware Upgrades* for steps to check readiness.

- The software in the 21c Oracle Grid Infrastructure home is not fully functional until the upgrade is completed. Running `srvctl`, `crsctl`, and other commands from the new Grid homes are not supported until the final `rootupgrade.sh` script is run and the upgrade is complete across all nodes.

  To manage databases in existing earlier release database homes during the Oracle Grid Infrastructure upgrade, use the `srvctl` from the existing database homes.

- To upgrade existing Oracle Clusterware installations to Oracle Grid Infrastructure 21c cluster, your release must be greater than or equal to Oracle Grid Infrastructure 12c Release 2 (12.2).

> ✎ **See Also:**
>
> *Oracle Database Upgrade Guide* for additional information about preparing for upgrades

About Storage Restrictions for Upgrade

If the Oracle Cluster Registry (OCR) and voting file locations for your current installation are on raw or block devices, or shared file systems, then you must migrate them to Oracle ASM disk groups, certified NAS device, or a shared file system before upgrading to Oracle Grid Infrastructure 21c.

About Upgrading Shared Grid Homes

- If the existing Oracle Clusterware home is a shared home, then you can use a non-shared home for the Oracle Grid Infrastructure for a cluster home for Oracle Clusterware and Oracle ASM 21c.

- You can perform upgrades on a shared Oracle Clusterware home.

About Single-Instance Oracle ASM Upgrade

- During Oracle Grid Infrastructure installation or upgrade, if there is a single instance Oracle ASM release on the local node, then it is converted to an Oracle Flex ASM 21c installation, and Oracle ASM runs in the Oracle Grid Infrastructure home on all nodes.

- If a single instance (non-clustered) Oracle ASM installation is on a remote node, which is a node other than the local node (the node on which the Oracle Grid Infrastructure installation or upgrade is being performed), then it remains a single instance Oracle ASM installation. However, during the installation or upgrade, when the OCR and voting files are placed on Oracle ASM, then an Oracle Flex ASM installation is created on all nodes in the cluster. The single instance Oracle ASM installation on the remote node becomes nonfunctional.

Preparing to Upgrade an Existing Oracle Clusterware Installation

If you have an existing Oracle Clusterware installation, then you upgrade your existing cluster by performing an out-of-place upgrade. You cannot perform an in-place upgrade.

The following topics list the steps you can perform before you upgrade Oracle Grid Infrastructure:

- Upgrade Checklist for Oracle Grid Infrastructure
  Review this checklist before upgrading an existing Oracle Grid Infrastructure. A cluster is being upgraded until all cluster member nodes are running the new installations, and the new clusterware becomes the active version.

- Checks to Complete Before Upgrading Oracle Grid Infrastructure
  Complete the following tasks before upgrading Oracle Grid Infrastructure.

- Moving Oracle Clusterware Files from NFS to Oracle ASM
  You can move Oracle Cluster Registry (OCR) and voting files from Network File System (NFS) to Oracle Automatic Storage Management (Oracle ASM) disk groups.

- **Running the Oracle ORAchk Upgrade Readiness Assessment**
  Download and run the Oracle ORAchk Upgrade Readiness Assessment before upgrading Oracle Grid Infrastructure.

- **Using CVU to Validate Readiness for Oracle Clusterware Upgrades**
  Oracle recommends that you use Cluster Verification Utility (CVU) to help to ensure that your upgrade is successful.

- **Using Dry-Run Upgrade Mode to Check System Upgrade Readiness**
  Use dry-run upgrade mode of Oracle Grid Infrastructure installation wizard, `gridSetup.sh`, to check readiness for Oracle Clusterware upgrades.

Upgrade Checklist for Oracle Grid Infrastructure

Review this checklist before upgrading an existing Oracle Grid Infrastructure. A cluster is being upgraded until all cluster member nodes are running the new installations, and the new clusterware becomes the active version.

Table 10-1   Upgrade Checklist for Oracle Grid Infrastructure Installation

Check	Task
Review Upgrade Guide for deprecation and desupport information that may affect upgrade planning.	*Oracle Database Upgrade Guide*
Patch set (recommended)	Install the latest patch set release for your existing installation. Review My Oracle Support note 2180188.1 for the list of latest patches before upgrading Oracle Grid Infrastructure.
Install user account	Confirm that the installation owner you plan to use is the same as the installation owner that owns the installation you want to upgrade.
Create a Grid home	Create a new Oracle Grid Infrastructure Oracle home (Grid home) where you can extract the image files. All Oracle Grid Infrastructure upgrades (upgrades of existing Oracle Clusterware and Oracle ASM installations) are out-of-place upgrades.
Instance names for Oracle ASM	Oracle Automatic Storage Management (Oracle ASM) instances must use standard Oracle ASM instance names. The default ASM SID for a single-instance database is +ASM.
Cluster names and Site names	Cluster names must have the following characteristics: • At least one character but no more than 15 characters in length. • Hyphens (-), and single-byte alphanumeric characters (a to z, A to Z, and 0 to 9). • It cannot begin with a numeric character. • It cannot begin or end with the hyphen (-) character.
Operating System	Confirm that you are using a supported operating system, kernel release, and all required operating system packages for the new Oracle Grid Infrastructure installation.
Network addresses for standard Oracle Grid Infrastructure	For standard Oracle Grid Infrastructure installations, confirm the following network configuration: • The private and public IP addresses are in unrelated, separate subnets. The private subnet should be in a dedicated private subnet. • The public and virtual IP addresses, including the SCAN addresses, are in the same subnet (the range of addresses permitted by the subnet mask for the subnet network). • Neither private nor public IP addresses use a link local subnet (169.254.*.*).

Table 10-1   (Cont.) Upgrade Checklist for Oracle Grid Infrastructure Installation

Check	Task
OCR on raw or block devices	Migrate OCR files from RAW or Block devices to Oracle ASM or a supported file system. Direct use of RAW and Block devices is not supported. Run the `ocrcheck` command to confirm Oracle Cluster Registry (OCR) file integrity. If this check fails, then repair the OCR before proceeding.
Check space for GIMR	When upgrading to Oracle Grid Infrastructure 19c or a later release, a new GIMR is created only if the source Grid home has a GIMR configured. Allocate additional storage space as described in *Oracle Clusterware Storage Space Requirements.* When upgrading from Oracle Grid Infrastructure 12c Release 2 (12.2), the GIMR is preserved with its contents.
Oracle ASM password file	When upgrading from Oracle Grid Infrastructure 12c Release 2 (12.2), Oracle Grid Infrastructure 18c, or Oracle Grid Infrastructure 19c to Oracle Grid Infrastructure 21c, move the Oracle ASM password file from file system to Oracle ASM, before proceeding with the upgrade. ✎ **Note:** Set `compatible.asm` to at least 12.2.0.1, before moving the password file.
CVU Upgrade Validation	Use Cluster Verification Utility (CVU) to assist you with system checks in preparation for starting an upgrade.
Shut Down Single-instance Database	If you have a single-instance database on a cluster that uses Oracle ASM, then you must shut down the single-instance database before upgrading.
Unset Environment variables	As the user performing the upgrade, unset the environment variables *$ORACLE_HOME* and *$ORACLE_SID*. Check that the *ORA_CRS_HOME* environment variable is not set. Do not use *ORA_CRS_HOME* as an environment variable, except under explicit direction from Oracle Support. Refer to the *Checks to Complete Before Upgrading Oracle Grid Infrastructure* for a complete list of environment variables to unset.
RACcheck Upgrade Readiness Assessment	Download and run the RACcheck Upgrade Readiness Assessment to obtain automated upgrade-specific health check for upgrades to Oracle Grid Infrastructure. See My Oracle Support note 1457357.1, which is available at the following URL: https://support.oracle.com/rs?type=doc&id=1457357.1
Back Up the Oracle Software Before Upgrades	Before you make any changes to the Oracle software, Oracle recommends that you create a backup of the Oracle software and databases.

Related Topics

- Moving Oracle Clusterware Files from NFS to Oracle ASM
  You can move Oracle Cluster Registry (OCR) and voting files from Network File System (NFS) to Oracle Automatic Storage Management (Oracle ASM) disk groups.

- Checks to Complete Before Upgrading Oracle Grid Infrastructure
  Complete the following tasks before upgrading Oracle Grid Infrastructure.

- My Oracle Support Note 2180188.1

Checks to Complete Before Upgrading Oracle Grid Infrastructure

Complete the following tasks before upgrading Oracle Grid Infrastructure.

1. For each node, use Cluster Verification Utility to ensure that you have completed preinstallation steps. It can generate Fixup scripts to help you to prepare servers. In addition, the installer helps you to ensure all required prerequisites are met.

   Ensure that you have the information you need during installation, including the following:

   • An Oracle base location for Oracle Clusterware.

   • An Oracle Grid Infrastructure home location that is different from your existing Oracle Clusterware location.

   • SCAN name and addresses, and other network addresses.

   • Privileged user operating system groups.

   • `root` user access, to run scripts as `root` during installation.

2. For the installation owner running the installation, if you have environment variables set for the existing installation, then unset the environment variables `$ORACLE_HOME` and `$ORACLE_SID`, as these environment variables are used during upgrade. For example, as the `grid` user, run the following commands on the local node:

   For bash shell:

   ```
   $ unset ORACLE_BASE
   $ unset ORACLE_HOME
   $ unset ORACLE_SID
   ```

   For C shell:

   ```
   $ unsetenv ORACLE_BASE
   $ unsetenv ORACLE_HOME
   $ unsetenv ORACLE_SID
   ```

3. If you have set `ORA_CRS_HOME` as an environment variable, following instructions from Oracle Support, then unset it before starting an installation or upgrade. You should never use `ORA_CRS_HOME` as an environment variable except under explicit direction from Oracle Support.

4. Ensure that the user profile for the installation user, for example, .profile or .cshrc, does not set any of these environment variables.

5. If you have an existing installation on your system, and you are using the same user account to install this installation, then unset the following environment variables: `ORA_CRS_HOME`, `ORACLE_HOME`, `ORA_NLS10`, `TNS_ADMIN` and any other environment variable set for the Oracle installation user that is connected with Oracle software homes.

6. Ensure that the `$ORACLE_HOME/bin` path is removed from your PATH environment variable.

Moving Oracle Clusterware Files from NFS to Oracle ASM

You can move Oracle Cluster Registry (OCR) and voting files from Network File System (NFS) to Oracle Automatic Storage Management (Oracle ASM) disk groups.

1. As the `grid` user, create the Oracle ASM disk group using ASMCA.

   ```
   $ ./asmca
   ```

   Follow the steps in the ASMCA wizard to create the Oracle ASM disk group, for example, `DATA`.

2. Move the voting files to the Oracle ASM disk group you created:

   ```
   $ crsctl replace votedisk +DATA
   ```

   The output of this command is as follows:

   ```
   CRS-4256: Updating the profile
   Successful addition of voting disk 24c6d682874a4f1ebf54f5ab0098b9e4.
   Successful deletion of voting disk 1b5044fa39684f86bfbe681f388e55fb.
   Successfully replaced voting disk group with +DATA_DG_OCR_VDSK.
   CRS-4256: Updating the profile
   CRS-4266: Voting file(s) successfully replaced
   ```

3. Check the status of Oracle Cluster Registry (OCR):

   ```
   $ ./ocrcheck
   ```

   The output of the command is as follows:

   ```
   Status of Oracle Cluster Registry is as follows :
            Version                  :          4
            Total space (kbytes)     :     409568
            Used space (kbytes)      :       1380
            Available space (kbytes) :     408188
            ID                       :  288871063
            Device/File Name         : /oradbocfs/storage/12101/ocr
            Device/File integrity check succeeded
            Cluster registry integrity check succeeded
   ```

4. As the `root` user, move the OCR files to the Oracle ASM disk group you created:

   ```
   # ./ocrconfig -add +DATA
   ```

5. Delete the Oracle Clusterware files from the NFS location:

   ```
   # ./ocrconfig -delete ocr_file_ path_previously_on_nfs
   ```

Running the Oracle ORAchk Upgrade Readiness Assessment

Download and run the Oracle ORAchk Upgrade Readiness Assessment before upgrading Oracle Grid Infrastructure.

Oracle ORAchk is an Oracle RAC configuration audit tool. Oracle ORAchk Upgrade Readiness Assessment can be used to obtain an automated upgrade-specific health check for upgrades to Oracle Grid Infrastructure 12.2, 18c, 19c, and 21c. You can run the Oracle ORAchk Upgrade Readiness Assessment tool and automate many of the manual pre-upgrade and post-upgrade checks.

Oracle recommends that you download and run the latest version of Oracle ORAchk from My Oracle Support. For information about downloading, configuring, and running Oracle ORAchk, refer to My Oracle Support note 1457357.1.

Related Topics

- *Oracle ORAchk and EXAchk User's Guide*

- https://support.oracle.com/rs?type=doc&id=1457357.1

Using CVU to Validate Readiness for Oracle Clusterware Upgrades

Oracle recommends that you use Cluster Verification Utility (CVU) to help to ensure that your upgrade is successful.

You can use CVU to assist you with system checks in preparation for starting an upgrade. CVU runs the appropriate system checks automatically, and either prompts you to fix problems, or provides a fixup script to be run on all nodes in the cluster before proceeding with the upgrade.

- About the CVU Upgrade Validation Command Options
  Review this information about running upgrade validations.

- Example of Verifying System Upgrade Readiness for Grid Infrastructure
  You can verify that the permissions required for installing Oracle Clusterware have been configured on the nodes `node1` and `node2` by running a command similar to the following.

About the CVU Upgrade Validation Command Options

Review this information about running upgrade validations.

About Cluster Verification Utility (CVU)

- Run Oracle Universal Installer (OUI), and allow the Cluster Verification Utility (CVU) validation built into OUI to perform system checks and generate fixup scripts.

- Run the CVU manual script `cluvfy.sh` to perform system checks and generate fixup scripts.

To use OUI to perform pre-install checks and generate fixup scripts, run the installation as you normally would. OUI starts CVU, and performs system checks as part of the installation process. Selecting OUI to perform these checks is particularly appropriate if you think you have completed preinstallation checks, and you want to confirm that your system configuration meets minimum requirements for installation.

To use the cluvfy.sh command-line script for CVU, navigate to the new Grid home where you extracted the image files for upgrade, that contains the `runcluvfy.sh` script, and run the command `runcluvfy.sh stage -pre crsinst -upgrade` to check the readiness of your Oracle Clusterware installation for upgrades. Running `runcluvfy.sh` with the `-pre crsinst -upgrade` options performs system checks to confirm if the cluster is in a correct state for upgrading from an existing clusterware installation.

Syntax

The command uses the following syntax, where variable content is indicated by italics:

```
runcluvfy.sh stage -pre crsinst -upgrade [-rolling]
-src_crshome src_Gridhome ]-dest_crshome dest_Gridhome -dest_version
dest_release
[-fixup][-fixupnoexec][-method sudo -user user_name [-location dir_path][-
method root][-verbose]
```

Options

The options are:

- `-rolling`

  Use this option to verify readiness for rolling upgrades.

- `-src_crshome` *src_Gridhome*

  Use this option to indicate the location of the source Oracle Clusterware or Grid home that you are upgrading, where *src_Gridhome* is the path to the home that you want to upgrade.

- `-dest_crshome` *dest_Gridhome*

  Use this option to indicate the location of the upgrade Grid home, where *dest_ Gridhome* is the path to the Grid home.

- `-dest_version` *dest_release*

  Use the `-dest_version` option to indicate the release number of the upgrade, including any patchset. The release number must include the five digits designating the release to the level of the platform-specific patch. For example: `21.0.0.0.0`.

- `-fixup [-method sudo -user` *user_name* `[-location` *dir_path*`][-method` root`]`

  Use the `-fixup` option to indicate that you want to generate instructions for any required steps you need to complete to ensure that your cluster is ready for an upgrade. The default location is the CVU work directory.

  The `-fixup -method` option defines the method by which `root` scripts are run. The `-method` flag requires one of the following options:

  - `sudo`: Run as a user on the sudoers list.

  - `root`: Run as the `root` user.

  If you select `sudo`, then enter the `-location` option to provide the path to Sudo on the server, and enter the `-user` option to provide the user account with Sudo privileges.

- `-fixupnoexec`

  If the option is specified, then on verification failure, the fix up data is generated and the instruction for manual execution of the generated fix ups is displayed.

- `-verbose`

  Use the `-verbose` flag to produce detailed output of individual checks.

Example of Verifying System Upgrade Readiness for Grid Infrastructure

You can verify that the permissions required for installing Oracle Clusterware have been configured on the nodes `node1` and `node2` by running a command similar to the following.

Example 10-1   Running `runcluvfy` Command

```
$ /u01/app/21.0.0/grid/runcluvfy.sh stage -pre crsinst -upgrade -
rolling -src_crshome
/u01/app/19.0.0/grid -dest_crshome /u01/app/21.0.0/grid -dest_version
21.0.0.0.0 -fixup -verbose
```

Related Topics

- *Oracle Database Upgrade Guide*

Using Dry-Run Upgrade Mode to Check System Upgrade Readiness

Use dry-run upgrade mode of Oracle Grid Infrastructure installation wizard, `gridSetup.sh`, to check readiness for Oracle Clusterware upgrades.

- [About Oracle Grid Infrastructure Dry-Run Upgrade Mode](#)
  Oracle Grid Infrastructure dry-run upgrade mode enables you to check system readiness for upgrade.

- [Performing Dry-Run Upgrade Using Oracle Universal Installer](#)
  Run the Oracle Grid Infrastructure installer in dry-run upgrade mode to determine if the system is ready for upgrade.

About Oracle Grid Infrastructure Dry-Run Upgrade Mode

Oracle Grid Infrastructure dry-run upgrade mode enables you to check system readiness for upgrade.

Starting with Oracle Grid Infrastructure 19c, the Oracle Grid Infrastructure installer enables you to perform a dry-run upgrade to check readiness of the system for upgrade. To perform Oracle Grid Infrastructure dry-run upgrade, create a new Grid home with the necessary user group permissions, extract the Oracle Grid Infrastructure 21c gold image to the new Grid home, and then start the installer with `-dryRunForUpgrade` flag.

Dry-run upgrades are not supported on Oracle Grid Infrastructure for a standalone server (Oracle Restart) configurations.

> **✎ Note:**
>
> The installer does not perform an actual upgrade in the dry-run upgrade mode. You can relaunch the installer, without any flag, from any of the cluster nodes to upgrade Oracle Grid Infrastructure if dry-run is successful.

The installer performs the following tasks in dry-run upgrade mode:

- Validates storage and network configuration for the new release

- Checks if the system meets the software and hardware requirements for the new release

- Checks for the patch requirements and apply necessary patches before starting the upgrade

- Writes system configuration issues or errors in the `gridSetupActions<timestamp>.log` log file

The Grid infrastructure dry-run upgrade flow is similar to a regular upgrade, but the installer does not run any configuration tool.

Related Topics

- Performing Dry-Run Upgrade Using Oracle Universal Installer
  Run the Oracle Grid Infrastructure installer in dry-run upgrade mode to determine if the system is ready for upgrade.

Performing Dry-Run Upgrade Using Oracle Universal Installer

Run the Oracle Grid Infrastructure installer in dry-run upgrade mode to determine if the system is ready for upgrade.

At any time during the dry-run upgrade, if you have a question about what you are being asked to do, or what input you are required to provide during dry-run upgrade, click the **Help** button on the installer page.

You should have your network information, storage information, and operating system users and groups available to you before you start dry-run upgrade.

1. As the `grid` user, download the Oracle Grid Infrastructure image file and extract the file to the Grid home.

   For example:

   ```
   $ mkdir -p /u01/app/21.0.0/grid
   $ chown grid:oinstall /u01/app/21.0.0/grid
   $ cd /u01/app/21.0.0/grid
   $ unzip -q download_location/grid_home.zip
   ```

   `download_location/grid_home.zip` is the path of the downloaded Oracle Grid Infrastructure image file.

   > **Note:**
   >
   > You must extract the image file into the directory where you want your Grid home to be located.

2. Start the Oracle Grid Infrastructure installation wizard in dry-run upgrade mode.

   ```
   $ /u01/app/21.0.0/grid/gridSetup.sh –dryRunForUpgrade
   ```

3. Select **Upgrade Oracle Grid Infrastructure** option to perform dry-run upgrade for Oracle Grid Infrastructure (Oracle Clusterware and Oracle ASM).

4. Select installation options as prompted. Oracle recommends that you configure `root` script automation, so that the `rootupgrade.sh` script can be run automatically during the dry-run upgrade.

5. Run the `root` scripts, either automatically or manually:

    • **Running root scripts automatically:**
    If you have configured `root` script automation, then the installer will run the `rootupgrade.sh` script automatically on the local node.

    • **Running root scripts manually**
    If you have not configured `root` script automation, then when prompted, run the `rootupgrade.sh` script on the local node.

    If you run the `root` scripts manually, then run the script only on the local node.

6. Check the `gridSetupActions<timestamp>.log` log file for errors and fix errors reported in the log file.

7. Exit the installer on the **Finish** screen.

> **Note:**
>
> The `gridSetup.sh -dryRunForUpgrade` command registers the new Oracle Grid Infrastructure home in the Oracle Inventory (oraInventory).

8. Relaunch the installer from the same Grid home without any flag to start an actual upgrade.

```
$ /u01/app/21.0.0/grid/gridSetup.sh
```

> **Note:**
>
> If you need to apply patches before performing the upgrade, then use the OPatch utility to apply patches on each cluster node.

Related Topics

• Upgrading Oracle Grid Infrastructure from an Earlier Release
Complete this procedure to upgrade Oracle Grid Infrastructure (Oracle Clusterware and Oracle Automatic Storage Management) from an earlier release.

Understanding Rolling Upgrades Using Batches

You can perform rolling upgrades of Oracle Grid Infrastructure in batches.

You can use `root` user automation to automate running the `rootupgrade.sh` script during the upgrade. When you use `root` user automation, you can divide the nodes into groups, or batches, and start upgrades of these batches. Between batches, you can move services from nodes running the previous release to the upgraded nodes, so that services are not affected by the upgrade. Oracle recommends that you use `root` automation, and allow the `rootupgrade.sh` script to stop and start instances automatically. You can also continue to run `root` scripts manually.

When you upgrade Oracle Grid Infrastructure without using `root` user automation, you upgrade the entire cluster. You cannot select or de-select individual nodes for upgrade. Oracle does not support attempting to add additional nodes to a cluster during a rolling upgrade. Oracle recommends that you leave Oracle RAC instances running when upgrading Oracle Clusterware. When you start the `root` script on each node, the database instances on that node are shut down and then the `rootupgrade.sh` script starts the instances again.

Restrictions for Selecting Nodes for Batch Upgrades

The following restrictions apply when selecting nodes in batches for upgrade:

- You can pool nodes in batches for upgrade, up to a maximum of three batches.

- The local node, where Oracle Universal Installer (OUI) is running, must be upgraded in batch one.

Performing Rolling Upgrade of Oracle Grid Infrastructure

Review this information to perform rolling upgrade of Oracle Grid Infrastructure.

- Upgrading Oracle Grid Infrastructure from an Earlier Release
  Complete this procedure to upgrade Oracle Grid Infrastructure (Oracle Clusterware and Oracle Automatic Storage Management) from an earlier release.

- Completing an Oracle Clusterware Upgrade when Nodes Become Unreachable
  If some nodes become unreachable in the middle of an upgrade, then you cannot complete the upgrade, because the upgrade script (`rootupgrade.sh`) did not run on the unreachable nodes. Because the upgrade is incomplete, Oracle Clusterware remains in the previous release.

- Joining Inaccessible Nodes After Forcing an Upgrade
  Use this procedure to join inaccessible nodes after a force cluster upgrade.

- Changing the First Node for Install and Upgrade
  If the first node becomes inaccessible, you can force another node to be the first node for installation or upgrade.

Upgrading Oracle Grid Infrastructure from an Earlier Release

Complete this procedure to upgrade Oracle Grid Infrastructure (Oracle Clusterware and Oracle Automatic Storage Management) from an earlier release.

At any time during the upgrade, if you have a question about what you are being asked to do, or what input you are required to provide during upgrade, click the **Help** button on the installer page.

> **Note:**
>
> Oracle ASM is upgraded as part of the Oracle Grid Infrastructure upgrade.

You should have your network information, storage information, and operating system users and groups available to you before you start upgrade, and you should be prepared to run root scripts.

1. As the `grid` user, download the Oracle Grid Infrastructure image files and extract the files to the Grid home.

For example:

```
$ mkdir -p /u01/app/21.0.0/grid
$ chown grid:oinstall /u01/app/21.0.0/grid
$ cd /u01/app/21.0.0/grid
$ unzip -q download_location/grid_home.zip
```

download_location/grid_home.zip is the path of the downloaded Oracle Grid Infrastructure image file.

> **Note:**
>
> - You must extract the image software into the directory where you want your Grid home to be located.
>
> - Download and copy the Oracle Grid Infrastructure image files to the local node only. During upgrade, the software is copied and installed on all other nodes in the cluster.

2. Start the Oracle Grid Infrastructure wizard by running the following command:

```
$ /u01/app/21.0.0/grid/gridSetup.sh
```

3. Select the **Upgrade Oracle Grid Infrastructure:** option to upgrade Oracle Grid Infrastructure (Oracle Clusterware and Oracle ASM).

4. On the Node Selection page, ensure that all cluster nodes are selected.

> **Note:**
>
> If any cluster node is unreachable, then select **Skip upgrade on unreachable nodes**. Otherwise, the upgrade will not complete until the unreachable nodes become reachable.

5. Select installation options as prompted. Oracle recommends that you configure `root` script automation, so that the `rootupgrade.sh` script can be run automatically during the upgrade.

6. Run the `root` scripts, either automatically or manually:

   - **Running root scripts automatically:**
     If you have configured `root` script automation, then use the pause between batches to relocate services from the nodes running the previous release to the new release.

   - **Running root scripts manually**
     If you have not configured `root` script automation, then when prompted, run the `rootupgrade.sh` script on each node in the cluster that you want to upgrade.

   If you run the `root` scripts manually, then run the script on the local node first. The script shuts down the earlier release installation, replaces it with the new Oracle Clusterware release, and starts the new Oracle Clusterware installation. After the

script completes successfully, you can run the script in parallel on all nodes except for one, which you select as the last node. When the script is run successfully on all the nodes except the last node, run the script on the last node.

7. Because the Oracle Grid Infrastructure home is in a different location than the former Oracle Clusterware and Oracle ASM homes, update any scripts or applications that use utilities, libraries, or other files that reside in the Oracle Clusterware and Oracle ASM homes.

8. Update the Oracle Enterprise Manager target parameters as described in the topic *Updating Oracle Enterprise Manager Cloud Control Target Parameters*.

> **Note:**
>
> * After upgrading to Oracle Grid Infrastructure 21c, remove the `ADR_BASE=/u01/app/grid` entry from `$ORACLE_HOME/network/admin/sqlnet.ora` file, if you use Oracle ASM for database storage.
>
> * At the end of the upgrade, if you set the Oracle Cluster Registry (OCR) backup location manually to the earlier release Oracle Clusterware home (CRS home), then you must change the OCR backup location to the new Oracle Grid Infrastructure home (Grid home). If you did not set the OCR backup location manually, then the backup location is changed for you during the upgrade.
>
> * Because upgrades of Oracle Clusterware are out-of-place upgrades, the previous release Oracle Clusterware home cannot be the location of the current release OCR backups. Backups in the old Oracle Clusterware home can be deleted.
>
> * If the cluster being upgraded has a single disk group that stores the OCR, OCR backup, Oracle ASM password, Oracle ASM password file backup, and the Grid Infrastructure Management Repository (GIMR), then Oracle recommends that you create a separate disk group or use another existing disk group and store the OCR backup, the GIMR and Oracle ASM password file backup in that disk group.

Related Topics

* *Oracle Clusterware Administration and Deployment Guide*

Completing an Oracle Clusterware Upgrade when Nodes Become Unreachable

If some nodes become unreachable in the middle of an upgrade, then you cannot complete the upgrade, because the upgrade script (`rootupgrade.sh`) did not run on the unreachable nodes. Because the upgrade is incomplete, Oracle Clusterware remains in the previous release.

1. Confirm that the upgrade is incomplete by entering the following command:

```
$ crsctl query crs activeversion
```

The active release should be the source release.

2. To resolve this problem, run the `rootupgrade.sh` command with the `-force` flag using the following syntax:

   Grid_home`/rootupgrade -force`

   For example:

   ```
   # /u01/app/21.0.0/grid/rootupgrade -force
   ```

   This command forces the upgrade to complete. The force cluster upgrade has the following limitations:

   - All active nodes must be upgraded to the newer release.

   - All inactive nodes (accessible or inaccessible) may be either upgraded or not upgraded.

   - For inaccessible nodes, after patch set upgrades, you can delete the node from the cluster. If the node becomes accessible later, and the patch version upgrade path is supported, then you can upgrade it to the new patch version.

3. Verify that the upgrade has completed by using the following command:

   ```
   $ crsctl query crs activeversion -f
   Oracle Clusterware active version on the cluster is [21.0.0.0.0].
   The cluster upgrade state is [FORCED]. The cluster active patch
   level is [0].
   ```

   The active release should be the upgraded release. The upgrade state `FORCED` indicates that Oracle Clusterware upgrade was completed forcibly on all reachable nodes, while one or more cluster nodes were unreachable. You can either upgrade or delete the unreachable cluster nodes to change the cluster upgrade state from `FORCED` to `NORMAL`.

Joining Inaccessible Nodes After Forcing an Upgrade

Use this procedure to join inaccessible nodes after a force cluster upgrade.

Starting with Oracle Grid Infrastructure 12c, after you complete a force cluster upgrade, you can use the procedure described here to join inaccessible nodes to the cluster as an alternative to deleting the nodes, which was required in earlier releases. To use this option, you must already have Oracle Grid Infrastructure 21c software installed on the nodes.

1. Log in as the `root` user on the node that you want to join to the cluster.

2. Change directory to the Oracle Grid Infrastructure home directory:

   ```
   $ cd Grid_home
   ```

3. Run the following command, where *upgraded_node* is one of the cluster nodes that is upgraded successfully:

   ```
   $ rootupgrade.sh -join -existingnode upgraded_node
   ```

Changing the First Node for Install and Upgrade

If the first node becomes inaccessible, you can force another node to be the first node for installation or upgrade.

During installation, if `root.sh` fails to complete on the first node, run the following command on another node using the `-force` option:

```
root.sh -force -first
```

For upgrade:

```
rootupgrade.sh -force -first
```

Upgrading Grid Infrastructure Management Repository from an Earlier Release

Complete this procedure to upgrade Grid Infrastructure Management Repository (GIMR) from an earlier release.

You must upgrade Oracle Grid Infrastructure to the same version before upgrading the GIMR.

1. As the `grid` user, create GIMR home directory on all cluster nodes:

    ```
    $ mkdir -p /u01/app/21.0.0/gimr
    $ chown grid:oinstall /u01/app/21.0.0/gimr
    ```

2. Download the Oracle Database image files and extract the files to the GIMR home.

    For example:

    ```
    $ cd /u01/app/21.0.0/gimr
    $ unzip -q download_location/db_home.zip
    ```

    download_location/db_home.zip is the path of the downloaded Oracle Database image file.

    > **Note:**
    >
    > *   You must extract the image software into the directory where you want your GIMR home to be located.
    >
    > *   Download and copy the Oracle Database image files to the local node only. During upgrade, the software is copied and installed on all other nodes in the cluster.

3. From the GIMR home directory, start the Oracle Database software installation:

```
$ cd /u01/app/21.0.0/gimr
$ ./runInstaller
```

> **Note:**
>
> Run the `runInstaller` command from the GIMR home directory only. Do not run `runInstaller` command that resides at `$GIMR_HOME/oui/bin/`, or any other location, to install Oracle Real Application Clusters.

4. In the Select Configuration Option screen, select the **Setup Software Only** option. Click **Next**.

5. In the Select Database Installation Option screen, select the **Oracle Real Application Clusters Database Installation** option. Click **Next**.

6. In the Node Selection screen, select all the nodes where you want to install Oracle RAC or Oracle RAC One Node software. Click **Next**.

7. In the Select Database Edition screen, select the **Enterprise Edition** option. Click **Next**.

8. In the Specify Install Location screen, provide the location of Oracle base for Oracle RAC software installation.

    The Oracle base directory must be different from the GIMR home directory. Click **Next**.

9. In the Privileged Operating System Groups screen, accept the default operating system group names for Oracle Database administration and click **Next**.

10. Select the option to **Automatically run configuration scripts**. Enter the credentials for the `root` user or a `sudo` account, then click **Next**.

    Alternatively, you can run the scripts manually as the `root` user at the end of the installation process when prompted by the installer.

11. The Perform Prerequisite Checks screen displays the results of the prerequisites checks. If any of the checks have a status of Failed and are not Fixable, then you must manually correct these issues. After you have fixed the issue, you can click the Check Again button to have the installer recheck the requirement and update the status. Repeat as needed until all the checks have a status of Succeeded. Click **Next**.

12. Review the contents of the Summary screen and then click **Install**.

    The installer displays a progress indicator enabling you to monitor the installation process.

13. If you did not configure automation of the root scripts, then you are required to run `root.sh` script as the `root` user, as specified in the Execute Configuration Scripts window. Do *not* click **OK** until you have run the `root.sh` script on all nodes as directed.

    You can run the `root.sh` script simultaneously on all nodes in the cluster for Oracle RAC installations or upgrades.

14. Complete GIMR upgrade process using the following command:

```
$ /u01/app/21.0.0/gimr/bin/mgmtua upgrade gimrHome
```

Applying Patches to Oracle Grid Infrastructure

After you have upgraded Oracle Grid Infrastructure 21c, you can install individual software patches by downloading them from My Oracle Support.

- About Individual Oracle Grid Infrastructure Patches
  Download Oracle ASM individual (one-off) patch and apply it to Oracle Grid Infrastructure using the OPatchAuto utility.

- About Oracle Grid Infrastructure Software Patch Levels
  Review this topic to understand how to apply patches for Oracle ASM and Oracle Clusterware.

- Patching Oracle Grid Infrastructure
  Review these topics to patch Oracle Grid Infrastructure in different stages of installation or upgrade.

About Individual Oracle Grid Infrastructure Patches

Download Oracle ASM individual (one-off) patch and apply it to Oracle Grid Infrastructure using the OPatchAuto utility.

Individual patches are called **one-off** patches. An Oracle ASM one-off patch is available for a specific release of Oracle ASM. If a patch you want is available, then you can download the patch and apply it to Oracle ASM using the OPatchAuto Utility. The OPatch inventory keeps track of the patches you have installed for your release of Oracle ASM. If there is a conflict between the patches you have installed and patches you want to apply, then the OPatchAuto Utility advises you of these conflicts.

About Oracle Grid Infrastructure Software Patch Levels

Review this topic to understand how to apply patches for Oracle ASM and Oracle Clusterware.

The software patch level for Oracle Grid Infrastructure represents the set of all one-off patches and Release Updates (RUs) applied to the Oracle Grid Infrastructure software release, including Oracle ASM. The release is the release number, in the format of major and RU release number. For example, with the release number 21.3.0.0.0, the major release is 21 and the RU release number is 3. The RU release number changes for RU patches, but it does not change for one-off patches.

As with standard upgrades to Oracle Grid Infrastructure, at any given point in time for normal operation of the cluster, all the nodes in the cluster must have the same software release and patch level. Because one-off patches can be applied as rolling upgrades, all possible patch levels on a particular software release are compatible with each other.

Patching Oracle Grid Infrastructure

Review these topics to patch Oracle Grid Infrastructure in different stages of installation or upgrade.

- Applying Patches During an Oracle Grid Infrastructure Installation or Upgrade
  You can download and apply Release Updates (RUs) and one-off patches during an Oracle Grid Infrastructure installation or upgrade.

- Applying Patches After an Oracle Grid Infrastructure Installation or Upgrade
  You can apply patches to Oracle Grid Infrastructure deployment in rolling mode after a successful or a failed upgrade. This mode is similar to the "Rolling Upgrade" mode in terms of the Oracle ASM operations allowed in this quiesce state.

- Applying Patches when Oracle Clusterware Fails to Start
  If the Oracle Clusterware stack does not start because of any error, then you can patch Oracle Grid Infrastructure using the OPatch utility.

Applying Patches During an Oracle Grid Infrastructure Installation or Upgrade

You can download and apply Release Updates (RUs) and one-off patches during an Oracle Grid Infrastructure installation or upgrade.

1. Download the patches you want to apply from My Oracle Support:

   https://support.oracle.com

2. Select the **Patches and Updates** tab to locate the patch.

   Oracle recommends that you select **Recommended Patch Advisor**, and enter the product group, release, and platform for your software.

3. Move the patches to an accessible directory like `/tmp`.

4. Change to the Oracle Grid Infrastructure home directory:

   ```
   $ cd /u01/app/21.0.0/grid
   ```

5. Apply Release Updates (RUs) and any one-off patches during the installation or upgrade process:

   ```
   $ ./gridSetup.sh -applyRU patch_directory_location -applyOneOffs
   comma_seperated_list_of_patch_directory_locations
   ```

   > **Note:**
   >
   > You can apply RUs and one-off patches separately or together in the same command.

6. Complete the remaining steps in the Oracle Grid Infrastructure configuration wizard to complete the installation or upgrade.

Applying Patches After an Oracle Grid Infrastructure Installation or Upgrade

You can apply patches to Oracle Grid Infrastructure deployment in rolling mode after a successful or a failed upgrade. This mode is similar to the "Rolling Upgrade" mode in terms of the Oracle ASM operations allowed in this quiesce state.

1. Download patches you want to apply from My Oracle Support:

   https://support.oracle.com

Select the **Patches and Updates** tab to locate the patch.

Oracle recommends that you select **Recommended Patch Advisor**, and enter the product group, release, and platform for your software. My Oracle Support provides you with a list of the most recent Release Updates (RUs) and Release Update Revisions (RURs).

Place the patches in an accessible directory, such as `/tmp`.

2. Review the patch documentation for the patch you want to apply, and complete all required steps before starting the patch upgrade.

3. Follow the instructions in the patch documentation to apply the patch on the first node.

```
# Grid_home/OPatch/opatchauto apply patch_directory_location/patch_ID
```

Wait for the patching to finish on the first node before patching other cluster nodes.

4. Apply the patch on all other cluster nodes, one at a time.

```
# Grid_home/OPatch/opatchauto apply patch_directory_location/patch_ID
```

Related Topics

• *OPatch User's Guide*

Applying Patches when Oracle Clusterware Fails to Start

If the Oracle Clusterware stack does not start because of any error, then you can patch Oracle Grid Infrastructure using the OPatch utility.

1. Download patches you want to apply from My Oracle Support:

https://support.oracle.com

Select the **Patches and Updates** tab to locate the patch.

Oracle recommends that you select **Recommended Patch Advisor**, and enter the product group, release, and platform for your software. My Oracle Support provides you with a list of the most recent Release Updates (RUs) and Release Update Revisions (RURs).

Place the patches in an accessible directory, such as `/tmp`.

2. Review the patch documentation for the patch you want to apply, and complete all required steps before starting the patch upgrade.

3. As the `root` user, stop the Oracle Clusterware stack on the first node.

```
# $ORACLE_HOME/bin/crsctl stop crs
```

4. Unlock the Oracle Grid Infrastructure home on the first node.

```
# $ORACLE_HOME/crs/install/rootcrs.sh -unlock
```

5. As the `grid` user, follow the instructions in the patch documentation to apply the patch on the first node.

```
$ $ORACLE_HOME/OPatch/opatch apply patch_directory_location/patch_ID
```

Wait for the patching to finish on the first node before patching other cluster nodes.

6. As the `root` user, lock the Oracle Grid Infrastructure home on the first node.

```
# $ORACLE_HOME/rdbms/install/rootadd_rdbms.sh
# $ORACLE_HOME/crs/install/rootcrs.sh -lock
```

7. Update Oracle Local Registry (OLR) to the patch level on the first node.

```
# $ORACLE_HOME/bin/clscfg -localpatch
```

8. Start the Oracle Clusterware stack on the first node.

```
# $ORACLE_HOME/bin/crsctl start crs
```

9. Update Oracle Cluster Registry (OCR) to the patch level on the first node.

```
# $ORACLE_HOME/bin/clscfg -patch
```

10. Repeat steps 3 through 9 on all other cluster nodes.

You can patch all other cluster nodes in parallel.

Applying Patches to Grid Infrastructure Management Repository

You can patch Grid Infrastructure Management Repository by installing in-place or out-of-place patches using the OPatch utility.

- In-Place Grid Infrastructure Management Repository Patching
  Use this procedure on each cluster node to perform in-place patching of Grid Infrastructure Management Repository (GIMR).

- Out-of-Place Grid Infrastructure Management Repository Patching
  Use this procedure on each cluster node to perform out-of-place patching of Grid Infrastructure Management Repository (GIMR).

- In-Place Grid Infrastructure Management Repository Patch Rollback
  Use this procedure on each cluster node for in-place Grid Infrastructure Management Repository (GIMR) patch rollback, in case of a failure or an error.

- Out-of-Place Grid Infrastructure Management Repository Patch Rollback
  Use this procedure on each cluster node for out-of-place Grid Infrastructure Management Repository (GIMR) patch rollback, in case of a failure or an error.

In-Place Grid Infrastructure Management Repository Patching

Use this procedure on each cluster node to perform in-place patching of Grid Infrastructure Management Repository (GIMR).

1. Download patches you want to apply from My Oracle Support:

   https://support.oracle.com

   Select the Patches and Updates tab to locate the patch.

Oracle recommends that you select **Recommended Patch Advisor**, and enter the product group, release, and platform for your software. My Oracle Support provides you with a list of the most recent Release Updates (RUs) and Release Update Revisions (RURs).

Place the patches in an accessible directory, such as `/tmp`.

2. Review the patch documentation for the patch you want to apply, and complete all required steps before starting the patch upgrade.

3. As the `grid` user, run the following command to perform the prerequisite checks:

```
$ $ORACLE_HOME/bin/mgmtprepatchbootstrap.pl apply [inplace] --patchHome
bootstrap_location
--sourceHome source_GIMR_home --id patch_ID_list
```

For example:

```
$ $ORACLE_HOME/bin/mgmtprepatchbootstrap.pl apply [inplace] --patchHome /tmp
--sourceHome /u01/app/21.0.0/gimr --id 123456
```

Where:

- `patchHome` is the location of the unzipped patch file.

- `sourceHome` is the GIMR home that you are patching.

> ✎ **Note:**
>
> Do not use slash (/) at the end of the path.

4. Change directory to the `/OPatch` directory in the GIMR home. For example:

```
$ cd GIMR_home/OPatch
```

5. Follow the instructions in the patch documentation to apply the patch. For example:

```
$ ./opatch apply patch directory_location/patch_ID
```

6. Run the following command to apply the SQL patch:

```
$ $ORACLE_HOME/bin/mgmtca applyDatapatch -postPatchApply -id
patch_ID_list -sourceHome source_GIMR_home
```

Out-of-Place Grid Infrastructure Management Repository Patching

Use this procedure on each cluster node to perform out-of-place patching of Grid Infrastructure Management Repository (GIMR).

1. Download patches you want to apply from My Oracle Support:

   https://support.oracle.com

   Select the Patches and Updates tab to locate the patch.

Oracle recommends that you select **Recommended Patch Advisor**, and enter the product group, release, and platform for your software. My Oracle Support provides you with a list of the most recent Release Updates (RUs) and Release Update Revisions (RURs).

Place the patches in an accessible directory, such as `/tmp`.

2. Review the patch documentation for the patch you want to apply, and complete all required steps before starting the patch upgrade.

3. Clone the GIMR home.

4. Change directory to the `/OPatch` directory in the cloned GIMR home. For example:

   ```
   $ cd GIMR_home/OPatch
   ```

5. Follow the instructions in the patch documentation to apply the patch. For example:

   ```
   $ ./opatch apply patch directory_location/patch_ID -sourceHome
   source_GIMR_home
   ```

   > **Note:**
   >
   > Do not use slash (/) at the end of the source GIMR home path.

6. Run the following command to apply the SQL patch:

   ```
   $ $ORACLE_HOME/bin/mgmtca applyDatapatch -postPatchApply -id
   patch_ID_list -sourceHome source_GIMR_home
   ```

In-Place Grid Infrastructure Management Repository Patch Rollback

Use this procedure on each cluster node for in-place Grid Infrastructure Management Repository (GIMR) patch rollback, in case of a failure or an error.

1. Review the patch documentation for the patch you want to rollback, and complete all required steps before starting the rollback.

2. As the `grid` user, run the following command to perform the prerequisite checks:

   ```
   $ $ORACLE_HOME/bin/mgmtprepatchbootstrap.pl rollback [inplace] --
   sourceHome source_GIMR_home --id patch_ID_list
   ```

3. Change directory to the `/OPatch` directory in the GIMR home. For example:

   ```
   $ cd GIMR_home/OPatch
   ```

4. Follow the instructions in the patch documentation to rollback the patch. For example:

   ```
   # ./opatch rollback directory_location/patch_ID
   ```

5. Run the following command to rollback the SQL patch:

```
$ $ORACLE_HOME/bin/mgmtca applyDatapatch -postPatchRollback -id
patch_ID_list
```

Out-of-Place Grid Infrastructure Management Repository Patch Rollback

Use this procedure on each cluster node for out-of-place Grid Infrastructure Management Repository (GIMR) patch rollback, in case of a failure or an error.

1. Review the patch documentation for the patch you want to rollback, and complete all required steps before starting the rollback.

2. Change directory to the `/OPatch` directory in the cloned GIMR home. For example:

```
$ cd GIMR_home/OPatch
```

3. Follow the instructions in the patch documentation to rollback the patch. For example:

```
# ./opatch rollback directory_location/patch_ID
```

4. Run the following command to rollback the SQL patch:

```
$ $ORACLE_HOME/bin/mgmtca applyDatapatch -postPatchRollback -id
patch_ID_list -sourceHome source_GIMR_home
```

Updating Oracle Enterprise Manager Cloud Control Target Parameters

After upgrading Oracle Grid Infrastructure, upgrade the Enterprise Manager Cloud Control target.

Because Oracle Grid Infrastructure 21c is an out-of-place upgrade of the Oracle Clusterware home in a new location (the Oracle Grid Infrastructure for a cluster home, or Grid home), the path for the `CRS_HOME` parameter in some parameter files must be changed. If you do not change the parameter, then you encounter errors such as "cluster target broken" on Oracle Enterprise Manager Cloud Control.

To resolve the issue, update the Enterprise Manager Cloud Control target, and then update the Enterprise Manager Agent Base Directory on each cluster member node running an agent.

- Updating the Enterprise Manager Cloud Control Target After Upgrades
  After upgrading Oracle Grid Infrastructure, update the Enterprise Manager Target with the new Grid home path.

- Updating the Enterprise Manager Agent Base Directory After Upgrades
  After upgrading Oracle Grid Infrastructure, update the Enterprise Manager Agent Base Directory on each cluster member node running an agent.

- Registering Resources with Oracle Enterprise Manager After Upgrades
  After upgrading Oracle Grid Infrastructure, add the new resource targets to Oracle Enterprise Manager Cloud Control.

Updating the Enterprise Manager Cloud Control Target After Upgrades

After upgrading Oracle Grid Infrastructure, update the Enterprise Manager Target with the new Grid home path.

1. Log in to Enterprise Manager Cloud Control.

2. Navigate to the **Targets** menu, and then to the **Cluster** page.

3. Click a cluster target that was upgraded.

4. Click **Cluster**, then **Target Setup**, and then **Monitoring Configuration** from the menu.

5. Update the value for **Oracle Home** with the new Grid home path.

6. Save the updates.

Updating the Enterprise Manager Agent Base Directory After Upgrades

After upgrading Oracle Grid Infrastructure, update the Enterprise Manager Agent Base Directory on each cluster member node running an agent.

The Agent Base directory is a directory where the Management Agent home is created. The Management Agent home is in the path *Agent_Base_Directory*/core/ *EMAgent_Version*. For example, if the Agent Base directory is `/u01/app/emagent`, then the Management Agent home is created as `/u01/app/emagent/core/13.1.1.0`.

1. Navigate to the `bin` directory in the Management Agent home.

2. In the `/u01/app/emagent/core/13.1.1.0/bin` directory, open the file `emctl` with a text editor.

3. Locate the parameter `CRS_HOME`, and update the parameter to the new Grid home path.

4. Repeat steps 1 to 3 on each node of the cluster with an Enterprise Manager agent.

Registering Resources with Oracle Enterprise Manager After Upgrades

After upgrading Oracle Grid Infrastructure, add the new resource targets to Oracle Enterprise Manager Cloud Control.

Discover and add new resource targets in Oracle Enterprise Manager after Oracle Grid Infrastructure upgrade. The following procedure provides an example of discovering an Oracle ASM listener target after upgrading Oracle Grid Infrastructure:

1. Log in to Oracle Enterprise Manager Cloud Control.

2. From the **Setup** menu, select **Add Target**, and then select **Add Targets Manually**.

   The Add Targets Manually page is displayed.

3. In the Add Targets page, select the **Add Using Guided Process** option and **Target Type** as `Oracle Database, Listener and Automatic Storage Management`.

   For any other resource to be added, select the appropriate **Target Type** in Oracle Enterprise Manager discovery wizard.

4. Click **Add Using Guided Process**.

   The Target Discover wizard is displayed.

5. For the **Specify Host or Cluster** field, click on the **Search** icon and search for **Target Types** of `Hosts`, and select the corresponding Host.

6. Click **Next**.

7. In the Target Discovery: Results page, select the discovered Oracle ASM Listener target, and click **Configure**.

8. In the Configure Listener dialog box, specify the listener properties and click **OK**.

9. Click **Next** and complete the discovery process.

   The listener target is discovered in Oracle Enterprise Manager with the status as Down.

10. From the **Targets** menu, select the type of target.

11. Click the target name to navigate to the target home page.

12. From the host, database, middleware target, or application menu displayed on the target home page, select **Target Setup**, then select **Monitoring Configuration**.

13. In the Monitoring Configuration page for the listener, specify the host name in the **Machine Name** field and the password for the `ASMSNMP` user in the **Password** field.

14. Click **OK**.

Oracle ASM listener target is displayed with the correct status.
Similarly, you can add other clusterware resources to Oracle Enterprise Manager after an Oracle Grid Infrastructure upgrade.

Unlocking and Deinstalling the Previous Release Grid Home

After upgrading from previous releases, if you want to deinstall the previous release Oracle Grid Infrastructure home, then you must first change the permission and ownership of the previous release Grid home.

1. As the `root` user, change the permission and ownership of the previous release Grid home using the following command syntax, where *oldGH* is the previous release Grid home, *swowner* is the Oracle Grid Infrastructure installation owner, and *oldGHParent* is the parent directory of the previous release Grid home:

```
#chmod -R 755 oldGH
#chown -R swowner oldGH
#chown swowner oldGHParent
```

For example:

```
#chmod -R 755 /u01/app/19.0.0/grid
#chown -R grid /u01/app/19.0.0/grid
#chown grid /u01/app/19.0.0
```

2. After you change the permissions and ownership of the previous release Grid home, log in as the Oracle Grid Infrastructure installation owner (`grid`, in the preceding example), and use the `deinstall` command from previous release Grid home (*oldGH*) `$ORACLE_HOME/deinstall` directory.

> ⚠️ **Caution:**
>
> You must use the `deinstall` command from the same release to remove Oracle software. Do not run the `deinstall` command from a later release to remove Oracle software from an earlier release. For example, do not run the `deinstall` command from the 21.0.0.0.0 Oracle home to remove Oracle software from an existing 19.0.0.0.0 Oracle home.

Checking Cluster Health Monitor Repository Size After Upgrading

If you are upgrading Oracle Grid Infrastructure from a prior release using IPD/OS to the current release, then review the Cluster Health Monitor repository size (the CHM repository).

1. Review your CHM repository needs, and determine if you need to increase the repository size to maintain a larger CHM repository.

> ✏️ **Note:**
>
> Your previous IPD/OS repository is deleted when you install Oracle Grid Infrastructure.

By default, the CHM repository size is a minimum of either 1GB or 3600 seconds (1 hour), regardless of the size of the cluster.

2. To enlarge the CHM repository, use the following command syntax, where *RETENTION_TIME* is the size of CHM repository in number of seconds:

```
oclumon manage -repos changeretentiontime RETENTION_TIME
```

For example, to set the repository size to four hours:

```
oclumon manage -repos changeretentiontime 14400
```

The value for *RETENTION_TIME* must be more than 3600 (one hour) and less than 259200 (three days). If you enlarge the CHM repository size, then you must ensure that there is local space available for the repository size you select on each node of the cluster. If you do not have sufficient space available, then you can move the repository to shared storage.

Downgrading Oracle Clusterware to an Earlier Release

After a successful or a failed upgrade, you can downgrade Oracle Clusterware to the previous release.

- About Oracle Grid Infrastructure Downgrades
  Understand the process and impacts of Oracle Grid Infrastructure downgrades.

- Options for Oracle Grid Infrastructure Downgrades
  You can downgrade Oracle Grid Infrastructure to earlier releases after a successful or a failed upgrade.

- Restrictions for Oracle Grid Infrastructure Downgrades
  Review the following information for restrictions and changes for downgrading Oracle Grid Infrastructure installations.

- Downgrading Oracle Standalone Cluster
  Use this procedure to downgrade Oracle Standalone Cluster after a successful upgrade.

- Downgrading Oracle Grid Infrastructure when Upgrade Fails
  If upgrade of Oracle Grid Infrastructure fails before CVU post upgrade checks succeed, then you can run `gridSetup.sh` and downgrade Oracle Grid Infrastructure to the earlier release.

- Downgrading Oracle Grid Infrastructure Using Online Abort Upgrade
  If upgrade of Oracle Grid Infrastructure fails before setting the active version of Oracle Clusterware, then follow these steps to downgrade Oracle Grid Infrastructure to the earlier release.

About Oracle Grid Infrastructure Downgrades

Understand the process and impacts of Oracle Grid Infrastructure downgrades.

Downgrading Oracle Clusterware restores the Oracle Clusterware configuration to the state it was in before the Oracle Grid Infrastructure upgrade. Any configuration changes you performed during or after the Oracle Grid Infrastructure upgrade are removed and cannot be recovered.

To restore Oracle Clusterware to the previous release, use the downgrade procedure for the release to which you want to downgrade.

> **✎ Note:**
>
> - You can downgrade the cluster nodes in any sequence. You can downgrade all cluster nodes except one, in parallel. You must downgrade the last node after you downgrade all other nodes.
>
> - When downgrading after a failed upgrade, if the `rootcrs.sh` or `rootcrs.bat` file does not exist on a node, then instead of executing the script, use the command `perl rootcrs.pl`. Use the Perl interpreter located in the Oracle Home directory.

Options for Oracle Grid Infrastructure Downgrades

You can downgrade Oracle Grid Infrastructure to earlier releases after a successful or a failed upgrade.

Downgrade options include the following earlier releases:

- Oracle Grid Infrastructure downgrade to Oracle Grid Infrastructure 19c.

- Oracle Grid Infrastructure downgrade to Oracle Grid Infrastructure 18c.

- Oracle Grid Infrastructure downgrade to Oracle Grid Infrastructure 12c Release 2 (12.2).

> **Note:**
>
> When you downgrade Oracle Grid Infrastructure to an earlier release, for example from Oracle Grid Infrastructure 21c to Oracle Grid Infrastructure 19c, the later release RAC databases already registered with Oracle Grid Infrastructure will not start after the downgrade.

Related Topics

- [My Oracle Support Note 2180188.1](#)

Restrictions for Oracle Grid Infrastructure Downgrades

Review the following information for restrictions and changes for downgrading Oracle Grid Infrastructure installations.

- You can only downgrade to the Oracle Grid Infrastructure release you upgraded from. For example, if you upgraded from Oracle Grid Infrastructure 12c Release 2 (12.2) to Oracle Grid Infrastructure 21c, you can only downgrade to Oracle Grid Infrastructure 12c Release 2 (12.2).

Downgrading Oracle Standalone Cluster

Use this procedure to downgrade Oracle Standalone Cluster after a successful upgrade.

1. As the `grid` user, downgrade the GIMR using the command syntax `mgmtua downgrade gimrHome` from 21c GIMR home directory after a successful upgrade:

    ```
    $ $ORACLE_HOME/bin/mgmtua downgrade gimrHome
    ```

    > **Note:**
    >
    > Run this step only if the GIMR is configured locally and upgrade is failed after Oracle Grid Infrastructure is upgraded successfully. The GIMR home is not required after the downgrade is complete.

2. As the `root` user, use the command syntax `rootcrs.sh -downgrade` from 21c Grid home to downgrade Oracle Grid Infrastructure on all nodes, in any sequence. For example:

```
# $ORACLE_HOME/crs/install/rootcrs.sh -downgrade
```

Run this command from a directory that has write permissions for the Oracle Grid Infrastructure installation user. You can run the downgrade script in parallel on all cluster nodes, but one.

3. Downgrade the last node after you downgrade all other nodes:

```
# $ORACLE_HOME/crs/install/rootcrs.sh -downgrade
```

4. As the `grid` user, remove Oracle Grid Infrastructure 21c Grid home as the active Oracle Clusterware home:

   a. On any of the cluster member nodes where the `rootupgrade.sh` script has run successfully, log in as the Oracle Grid Infrastructure installation owner.

   b. Use the following command to start the installer, where `/u01/app/21.0.0/grid` is the location of the new (upgraded) Grid home:

```
$ cd $ORACLE_HOME/oui/bin
$ ./runInstaller -nowait -waitforcompletion -ignoreSysPrereqs -
updateNodeList -silent CRS=false
ORACLE_HOME=/u01/app/21.0.0/grid "CLUSTER_NODES=node1,node2,node3" -
doNotUpdateNodeList
```

Add the flag `-cfs` if the Grid home is a shared home.

5. Set previous release Grid home as the active Oracle Clusterware home:

   a. On any of the cluster member nodes where the `rootupgrade` script has run successfully, log in as the Oracle Grid Infrastructure installation owner.

   b. Use the following command to start the installer, where the path you provide for `ORACLE_HOME` is the location of the home directory from the earlier Oracle Clusterware installation.

```
$ cd $ORACLE_HOME/oui/bin
$ ./runInstaller -nowait -waitforcompletion -ignoreSysPrereqs
-updateNodeList -silent CRS=true
ORACLE_HOME=pre-upgrade-grid-home
"CLUSTER_NODES=node1,node2,node3"
```

pre-upgrade-grid-home is the previous release Grid home.

6. As the `root` user, start the Oracle Clusterware stack on all nodes from previous release Grid home.

```
# crsctl start crs
```

7. As the `grid` user, run the `GIMR_recreate.sh` script from 21c Oracle Base directory to recreate the GIMR container:

```
$ $ORACLE_BASE/admin/mgmtdb/GIMR_recreate.sh
```

> **✏ Note:**
>
> Run this step only if the GIMR is configured locally in the cluster and upgrade is failed after `rootupgade.sh` is run successfully on all of the cluster nodes.

Downgrading Oracle Grid Infrastructure when Upgrade Fails

If upgrade of Oracle Grid Infrastructure fails before CVU post upgrade checks succeed, then you can run `gridSetup.sh` and downgrade Oracle Grid Infrastructure to the earlier release.

Run this procedure to downgrade Oracle Clusterware only when the upgrade fails before CVU post upgrade checks succeed.

- From the later release Grid home, run `gridSetup.sh` in silent mode, to downgrade Oracle Clusterware:

```
$ /u01/app/21.0.0/grid/gridSetup.sh -silent -downgrade [-nodes]
nodes_to_be_downgraded
[-oldHome] pre-upgrade-grid-home
[-configmethod] root | sudo [-sudopath path_to_sudo_program]
[-sudousername sudoer_name]
```

On Windows systems, run `setup.exe` instead of `gridSetup.sh`.

> **✏ Note:**
>
> You can downgrade the cluster nodes in any sequence.

Downgrading Oracle Grid Infrastructure Using Online Abort Upgrade

If upgrade of Oracle Grid Infrastructure fails before setting the active version of Oracle Clusterware, then follow these steps to downgrade Oracle Grid Infrastructure to the earlier release.

Run this procedure to downgrade Oracle Clusterware only when the upgrade fails before root script runs the `crsctl set crs activeversion` command on the last node. Use this procedure for downgrading Oracle Grid Infrastructure if there is a need to avoid downtime of the whole cluster. This procedure downgrades the cluster to the previous release. Because Oracle ASM and database operations are limited in this state, it is recommended to move the cluster from this state as soon as possible. Complete the downgrade of Oracle Grid Infrastructure as per the procedure documented in *Downgrading Oracle Grid Infrastructure after Upgrade Fails*.

1. Shut down the Oracle Grid Infrastructure stack on the first node:

```
crsctl stop crs
```

2. From any node where the Grid Infrastructure stack from the earlier release is running, unset the Oracle ASM rolling migration mode as follows:

- As the `grid` user, and run the following command as SYSASM user on the Oracle ASM instance:

```
SQL> ALTER SYSTEM STOP ROLLING MIGRATION;
```

3. On any node running Oracle Grid Infrastructure other than the first node, from the Grid home of the earlier release, run the command:

```
$ clscfg -nodedowngrade -h hostname
```

hostname is the host name of the first node.

4. From the later release Grid home, run the command to downgrade Oracle Clusterware:

```
$ rootcrs.sh -downgrade -online
```

If `rootcrs.sh` is not present, then use `rootcrs.pl`.

5. Start Oracle Grid Infrastructure stack on the first node from the earlier release Grid home:

```
$ crsctl start crs
```

> **Note:**
>
> You can downgrade the cluster nodes in any sequence.

Completing Failed or Interrupted Installations and Upgrades

If Oracle Universal Installer (OUI) exits on the node from which you started the upgrade, or the node reboots before you confirm that the `rootupgrade.sh` script was run on all nodes, then the upgrade remains incomplete.

In an incomplete installation or upgrade, configuration assistants still need to run, and the new Grid home still needs to be marked as active in the central Oracle inventory. You must complete the installation or upgrade on the affected nodes manually.

- Completing Failed Installations and Upgrades
  Understand how to join nodes to the cluster after installation or upgrade fails on some nodes.

- Continuing Incomplete Upgrade of First Node
  Review this information to complete the upgrade, if upgrade of Oracle Grid Infrastructure fails on the first node.

- Continuing Incomplete Upgrade of First Node
  Review this information to complete the upgrade, if upgrade of Oracle Grid Infrastructure fails on the first node.

- Continuing Incomplete Installation on First Node
  Review this information to continue an incomplete installation of Oracle Grid Infrastructure, if installation fails on the first node.

- Continuing Incomplete Installation on Remote Nodes
  Review this information to continue incomplete installation on remote nodes.

Completing Failed Installations and Upgrades

Understand how to join nodes to the cluster after installation or upgrade fails on some nodes.

If installation or upgrade of Oracle Grid Infrastructure on some nodes fails, and the installation or upgrade completes with only successful nodes in the cluster, then follow this procedure to add the failed nodes to the cluster.

1. Remove the Oracle Grid Infrastructure software from the failed nodes:

   Grid_home/deinstall/deinstall -local

2. As the `root` user, from a node where Oracle Clusterware is installed, delete the failed nodes using the `delete node` command:

   Grid_home/bin/crsctl delete node -n *node_name*

   node_name is the node to be deleted.

3. Run the Oracle Grid Infrastructure installation wizard and follow the steps in the wizard to add the nodes:

   Grid_home/gridSetup.sh

   The nodes are added to the cluster.

Continuing Incomplete Upgrade of First Node

Review this information to complete the upgrade, if upgrade of Oracle Grid Infrastructure fails on the first node.

1. If the root script failure indicated a need to reboot, through the message `CLSRSC-400`, then reboot the first node (the node where the upgrade was started). Otherwise, manually fix or clear the error condition, as reported in the error output.

2. If necessary, log in as `root` to the first node. Change directory to the new Grid home on the first node, and run the `rootupgrade.sh` script on that node again. For example:

   ```
   [root@node1]# cd Grid_home
   [root@node1]# ./rootupgrade.sh
   ```

3. Complete the upgrade of all other nodes in the cluster.

   Complete the upgrade of all other nodes in the cluster.

   ```
   [root@node2]# ./rootupgrade.sh
   ```

4. Configure a response file, and provide passwords for the installation.

5. To complete the upgrade, log in as the Grid installation owner, and run `gridSetup.sh`, located in the `Grid_home`, specifying the response file that you created. For example, where the response file is `gridinstall.rsp`:

```
[grid@node1]$ Grid_home/gridSetup.sh -executeConfigTools -responseFile
Grid_home/install/response/gridinstall.rsp
```

> **✎ Note:**
>
> When you reupgrade Oracle Grid Infrastructure, you must use the `-all` flag with the `executeConfigTools` command to run all the configuration tools.

Continuing Incomplete Upgrade of First Node

Review this information to complete the upgrade, if upgrade of Oracle Grid Infrastructure fails on the first node.

1. If the root script failure indicated a need to reboot, through the message `CLSRSC-400`, then reboot the first node (the node where the upgrade was started). Otherwise, manually fix or clear the error condition, as reported in the error output.

2. If necessary, log in as `root` to the first node. Change directory to the new Grid home on the first node, and run the `rootupgrade.sh` script on that node again. For example:

```
[root@node1]# cd Grid_home
[root@node1]# ./rootupgrade.sh
```

3. Complete the upgrade of all other nodes in the cluster.

   Complete the upgrade of all other nodes in the cluster.

```
[root@node2]# ./rootupgrade.sh
```

4. Configure a response file, and provide passwords for the installation.

5. To complete the upgrade, log in as the Grid installation owner, and run `gridSetup.sh`, located in the `Grid_home`, specifying the response file that you created. For example, where the response file is `gridinstall.rsp`:

```
[grid@node1]$ Grid_home/gridSetup.sh -executeConfigTools -responseFile
Grid_home/install/response/gridinstall.rsp
```

> **✎ Note:**
>
> When you reupgrade Oracle Grid Infrastructure, you must use the `-all` flag with the `executeConfigTools` command to run all the configuration tools.

ORACLE®

Continuing Incomplete Installation on First Node

Review this information to continue an incomplete installation of Oracle Grid Infrastructure, if installation fails on the first node.

1. If the root script failure indicated a need to reboot, through the message `CLSRSC-400`, then reboot the first node (the node where the installation was started). Otherwise, manually fix or clear the error condition, as reported in the error output.

2. If necessary, log in as `root` to the first node. Run the `orainstRoot.sh` script on that node again. For example:

    ```
    $ sudo -s
    [root@node1]# cd /u01/app/oraInventory
    [root@node1]# ./orainstRoot.sh
    ```

3. Change directory to the Grid home on the first node, and run the `root` script on that node again. For example:

    ```
    [root@node1]# cd Grid_home
    [root@node1]# ./root.sh
    ```

4. Complete the installation on all other nodes.

5. Configure a response file, and provide passwords for the installation.

6. To complete the installation, log in as the Grid installation owner, and run `gridSetup.sh`, located in the Oracle Grid Infrastructure home, specifying the response file that you created. For example, where the response file is `gridinstall.rsp`:

    ```
    [grid@node1]$ Grid_home/gridSetup.sh -executeConfigTools -
    responseFile /u01/app/21.0.0/grid/install/response/gridinstall.rsp
    ```

Continuing Incomplete Installation on Remote Nodes

Review this information to continue incomplete installation on remote nodes.

1. If the root script failure indicated a need to reboot, through the message `CLSRSC-400`, then reboot the affected node. Otherwise, manually fix or clear the error condition, as reported in the error output.

2. If root automation is being used, click **Retry** on the OUI instance on the first node.

3. If root automation is not being used, follow these steps:

    a. Log into the affected node as `root`, and run the `orainstRoot.sh` script on that node. For example:

        ```
        $ sudo -s
        [root@node6]# cd /u01/app/oraInventory
        [root@node6]# ./orainstRoot.sh
        ```

b. Change directory to the Grid home, and run the `root.sh` script on the affected node. For example:

```
[root@node6]# cd Grid_home
[root@node6]# ./root.sh
```

4. Continue the installation from the OUI instance on the first node.

Converting to Oracle Extended Cluster After Upgrading Oracle Grid Infrastructure

Review this information to convert to an Oracle Extended Cluster after upgrading Oracle Grid Infrastructure. Oracle Extended Cluster enables you to deploy Oracle RAC databases on a cluster, in which some of the nodes are located in different sites.

Ensure that you have upgraded to Oracle Grid Infrastructure 21c as described in this chapter.

1. As the `root` user, log in to the first node, and run the following command:

```
# rootcrs.sh -converttoextended -first -sites list_of_sites -site
node_site
```

list_of_sites is the comma-separated list of sites in the extended cluster, and *node_site* is the node containing the site.

2. Run the following command on all other cluster nodes:

```
# rootcrs.sh -converttoextended -site node_site
```

node_site is the node containing the site.

3. Delete the default site after the associated nodes and storage are migrated.

```
# crsctl delete cluster site site_name
```

11

Removing Oracle Database Software

Use the `deinstall` command that is included in Oracle homes to remove Oracle software. Oracle does not support the removal of individual products or components.

> ⚠ **Caution:**
>
> If you have a single-instance database on a node in a cluster, and if you have multiple databases with the same global database name (GDN), then you cannot use the `deinstall` command to remove one database only.

- **About Oracle Deinstallation Options**
  You can stop and remove Oracle Database software and components in an Oracle Database home with the `deinstall` command.

- **Oracle Deinstallation (Deinstall)**
  You can run the `deinstall` command from an Oracle home directory after installation.

- **Deinstallation Examples for Oracle Database**
  Use these examples to help you understand how to run the `deinstall` command.

- **Deinstallation Response File Example for Oracle Grid Infrastructure for a Cluster**
  You can run the `deinstall` command with the `-paramfile` option to use the values you specify in the response file.

- **Deinstallation Example for Grid Infrastructure Management Repository**
  Use this example to understand how to deinstall Grid Infrastructure Management Repository (GIMR) on an Oracle Standalone Cluster.

- **Migrating Standalone Oracle Grid Infrastructure Servers to a Cluster**
  If you have an Oracle Database installation using Oracle Restart (that is, an Oracle Grid Infrastructure installation for a standalone server), and you want to configure that server as a cluster member node, then complete the following tasks:

- **Relinking Oracle Grid Infrastructure for a Cluster Binaries**
  After installing Oracle Grid Infrastructure for a cluster (Oracle Clusterware and Oracle ASM configured for a cluster), if you need to modify the binaries, then use the following procedure, where `Grid_home` is the Oracle Grid Infrastructure for a cluster home:

- **Changing the Oracle Grid Infrastructure Home Path**
  After installing Oracle Grid Infrastructure for a cluster (Oracle Clusterware and Oracle ASM configured for a cluster), if you need to change the Grid home path, then use the following example as a guide to detach the old Grid home, and to attach a new Grid home:

- **Unconfiguring Oracle Clusterware Without Removing Binaries**
  Running the `rootcrs.sh` command flags `-deconfig -force` enables you to unconfigure Oracle Clusterware on one or more nodes without removing installed binaries.

About Oracle Deinstallation Options

You can stop and remove Oracle Database software and components in an Oracle Database home with the `deinstall` command.

You can remove the following software using `deinstall`:

- Oracle Database
- Oracle Grid Infrastructure, which includes Oracle Clusterware and Oracle Automatic Storage Management (Oracle ASM)
- Oracle Real Application Clusters (Oracle RAC)
- Oracle Database Client

The `deinstall` command is available in Oracle home directories after installation. It is located in the `$ORACLE_HOME/deinstall` directory.

`deinstall` creates a response file by using information in the Oracle home and using the information you provide. You can use a response file that you generated previously by running the `deinstall` command using the `-checkonly` option. You can also edit the response file template.

If you run `deinstall` to remove an Oracle Grid Infrastructure installation, then the deinstaller prompts you to run the `deinstall` command as the `root` user. For Oracle Grid Infrastructure for a cluster, the script is `rootcrs.sh`, and for Oracle Grid Infrastructure for a standalone server (Oracle Restart), the script is `roothas.sh`.

> **✐ Note:**
>
> - You must run the `deinstall` command from the same release to remove Oracle software. Do not run the `deinstall` command from a later release to remove Oracle software from an earlier release. For example, do not run the `deinstall` command from the 21c Oracle home to remove Oracle software from an existing 19c Oracle home.
>
> - Starting with Oracle Database 12c Release 1 (12.1.0.2), the `roothas.sh` script replaces the `roothas.pl` script in the Oracle Grid Infrastructure home for Oracle Restart, and the `rootcrs.sh` script replaces the `rootcrs.pl` script in the Grid home for Oracle Grid Infrastructure for a cluster.

If the software in the Oracle home is not running (for example, after an unsuccessful installation), then `deinstall` cannot determine the configuration, and you must provide all the configuration details either interactively or in a response file.

In addition, before you run `deinstall` for Oracle Grid Infrastructure installations:

- If Grid Naming Service (GNS) is in use, then notify your DNS administrator to delete the subdomain entry from the DNS.

Files Deleted by deinstall

When you run `deinstall`, if the central inventory (`oraInventory`) contains no other registered homes besides the home that you are deconfiguring and removing, then `deinstall` removes the following files and directory contents in the Oracle base directory of the Oracle Database installation owner:

- `admin`

- `cfgtoollogs`

- `checkpoints`

- `diag`

- `oradata`

- `fast_recovery_area`

Oracle strongly recommends that you configure your installations using an Optimal Flexible Architecture (OFA) configuration, and that you reserve Oracle base and Oracle home paths for exclusive use of Oracle software. If you have any user data in these locations in the Oracle base that is owned by the user account that owns the Oracle software, then `deinstall` deletes this data.

> ⚠️ **Caution:**
>
> `deinstall` deletes Oracle Database configuration files, user data, and fast recovery area (FRA) files even if they are located outside of the Oracle base directory path.

Oracle Deinstallation (Deinstall)

You can run the `deinstall` command from an Oracle home directory after installation.

Purpose

`deinstall` stops Oracle software, and removes Oracle software and configuration files on the operating system for a specific Oracle home.

Syntax

The `deinstall` command uses the following syntax:

```
(./deinstall [-silent] [-checkonly] [-paramfile complete path of input
response file]
[-params name1=value name2=value . . .]
[-o complete path of directory for saving files]
[-tmpdir complete path of temporary directory to use]
[-logdir complete path of log directory to use] [-local] [-
skipLocalHomeDeletion] [-skipRemoteHomeDeletion] [-help]
```

Parameters

Parameter	Description
`-silent`	Use this flag to run `deinstall` in noninteractive mode. This option requires one of the following: • A working system that it can access to determine the installation and configuration information. The `-silent` flag does not work with failed installations. • A response file that contains the configuration values for the Oracle home that is being deinstalled or deconfigured. You can generate a response file to use or modify by running `deinstall` with the `-checkonly` flag. `deinstall` then discovers information from the Oracle home to deinstall and deconfigure. It generates the response file that you can then use with the `-silent` option. You can also modify the template file `deinstall.rsp.tmpl`, located in the `$ORACLE_HOME/deinstall/response` directory.
`-checkonly`	Use this flag to check the status of the Oracle software home configuration. Running `deinstall` with the `-checkonly` flag does not remove the Oracle configuration. The `-checkonly` flag generates a response file that you can then use with the `deinstall` command and `-silent` option.
`-paramfile` *complete path of input response file*	Use this flag to run `deinstall` with a response file in a location other than the default. When you use this flag, provide the complete path where the response file is located. The default location of the response file is `$ORACLE_HOME/deinstall/response`.
`-params` [*name1=value name2=value name3=value . . .*]	Use this flag with a response file to override one or more values to change in a response file you have created.
`-o` *complete path of directory for saving response files*	Use this flag to provide a path other than the default location where the response file (`deinstall.rsp.tmpl`) is saved. The default location of the response file is `$ORACLE_HOME/deinstall/response`.
`-tmpdir` *complete path of temporary directory to use*	Use this flag to specify a non-default location where `deinstall` writes the temporary files for the deinstallation.

Parameter	Description
`-logdir` *complete path of log directory to use*	Use this flag to specify a non-default location where `deinstall` writes the log files for the deinstallation.
`-local`	Use this flag on a multinode environment to deinstall Oracle software in a cluster.
	When you run `deinstall` with this flag, it deconfigures and deinstalls the Oracle software on the local node (the node where `deinstall` is run). On remote nodes, it deconfigures Oracle software, but does not deinstall the Oracle software.
`-skipLocalHomeDeletion`	Use this flag in Oracle Grid Infrastructure installations on a multinode environment to deconfigure a local Grid home without deleting the Grid home.
`-skipRemoteHomeDeletion`	Use this flag in Oracle Grid Infrastructure installations on a multinode environment to deconfigure a remote Grid home without deleting the Grid home.
`-help`	Use this option to obtain additional information about the command option flags.

Deinstallation Examples for Oracle Database

Use these examples to help you understand how to run the `deinstall` command.

Run `deinstall` from the `$ORACLE_HOME/deinstall` directory. The deinstallation starts without prompting you for the Oracle home path.

```
$ ./deinstall
```

You can generate a deinstallation response file by running `deinstall` with the `-checkonly` flag. Alternatively, you can use the response file template located at `$ORACLE_HOME/deinstall/response/deinstall.rsp.tmpl`. If you have a response file, then use the optional flag `-paramfile` to provide a path to the response file.

In the following example, the `deinstall` command is in the path `/u01/app/oracle/product/21.0.0/dbhome_1/deinstall`. It uses a response file called `my_db_paramfile.tmpl` in the software owner location `/home/usr/oracle`:

```
$ cd /u01/app/oracle/product/21.0.0/dbhome_1/deinstall
$ ./deinstall -paramfile /home/usr/oracle/my_db_paramfile.tmpl
```

To remove the Oracle Grid Infrastructure home, use the `deinstall` command in the Oracle Grid Infrastructure home.

ORACLE®

In this example, the Oracle Grid Infrastructure home is `/u01/app/21.0.0/grid`

```
$ cd  /u01/app/21.0.0/grid/deinstall
$ ./deinstall -paramfile /home/usr/oracle/my_grid_paramfile.tmpl
```

Deinstallation Response File Example for Oracle Grid Infrastructure for a Cluster

You can run the `deinstall` command with the `-paramfile` option to use the values you specify in the response file.

Deinstallation Parameter File Description and Example

The following is an example of a response file for a cluster on nodes `node1` and `node2`, in which the Oracle Grid Infrastructure for a cluster software binary owner is `grid`, the Oracle Grid Infrastructure home (Grid home) is in the path `/u01/app/21.0.0/grid`, the Oracle base (the Oracle base for Oracle Grid Infrastructure, containing Oracle ASM log files, Oracle Clusterware logs, and other administrative files) is `/u01/app/grid/`, and the central Oracle Inventory home (`oraInventory`) is `/u01/app/oraInventory`. The local node (the node where you run the deinstallation session from) is `node1`:

```
# Copyright (c) 2005, 2020, Oracle and/or its affiliates.  All rights
reserved.
ORACLE_HOME=/u01/app/21.0.0/grid
CDATA_AUSIZE=1
BIG_CLUSTER=true
ISROLLING=true
LOCAL_NODE=node1
OCR_VD_DISKGROUPS="+DATA1"
ObaseCleanupPtrLoc=/tmp/deinstall2020-09-23_08-53-16AM/utl/
orabase_cleanup.lst
CDATA_BACKUP_QUORUM_GROUPS=
ASM_CREDENTIALS=
EXTENDED_CLUSTER=false
SINGLE_NETWORK=false
LISTENER_USERNAME=aime
local=false
inventory_loc=/u01/app/oraInventory
ORACLE_HOME=/u01/app/21.0.0/grid
ASM_HOME=/u01/app/21.0.0/grid
ASM_DISK_GROUPS="+DATA1"
HUB_NODE_VIPS=node1-v,node2-v
PING_TARGETS=
ORA_DBA_GROUP=oinstall
ASM_DISCOVERY_STRING=/dev/rdsk/*
CDATA_DISKS=/dev/rdsk/c0t600144F0C4A01A3F000056E6A12A0022d0s3, /dev/
rdsk/c0t600144F0C4A01A3F000056E6A12A0022d1s4
MinimumSupportedVersion=11.2.0.1.0
ORACLE_HOME_VERSION=21.0.0.0.0
ASM_DISKSTRING=/dev/rdsk/*
CDATA_QUORUM_GROUPS=
AHF_INSTALL_LOC=/opt/oracle.ahf
```

```
CRS_HOME=true
ODA_CONFIG=
USER_IGNORED_PREREQ=true
RHP_CONF=false
HAS_NOMULTICAST_ENV=
GNS_TYPE=
CRS_STORAGE_OPTION=1
CDATA_SITES=
GIMR_CONFIG=no
CDATA_BACKUP_SIZE=0
GPNPGCONFIGDIR=$ORACLE_HOME
CDATA_DISK_GROUP=DATA1
LANGUAGE_ID=AMERICAN_AMERICA.AL32UTF8
CDATA_BACKUP_FAILURE_GROUPS=
RHP_PROVISIONED=
CRS_NODEVIPS='node1-v/255.255.252.0/eth0,node2-v/255.255.252.0/eth0'
ORACLE_OWNER=grid
GNS_ALLOW_NET_LIST=
silent=true
INSTALL_NODE=node1.example.com
ORACLE_HOME_VERSION_VALID=true
inst_group=oinstall
LOGDIR=/tmp/deinstall2020-09-23_08-53-16AM/logs
CDATA_EXC_VAULT=
EXTENDED_CLUSTER_SITES=crs-node
CDATA_REDUNDANCY=NORMAL
CDATA_BACKUP_DISK_GROUP=
APPLICATION_VIP=
HUB_NODE_LIST=node1,node2
GNS_DENY_ITF_LIST=
NODE_NAME_LIST=node1,node2
ORA_CRS_HOME=/u01/app/21.0.0/grid
ASM_LOCAL_SID=+ASM1
ORACLE_BASE=/u01/app/grid
GNS_CONF=false
CLUSTER_CLASS=STANDALONE
ORACLE_BINARY_OK=true
ASM_CONFIG=near
CDATA_BACKUP_REDUNDANCY=
CDATA_FAILURE_GROUPS=
OCR_LOCATIONS=
GIMR_CREDENTIALS=
ASM_ORACLE_BASE=/u01/app/grid
GPNPCONFIGDIR=$ORACLE_HOME
ORA_ASM_GROUP=dba
GNS_CREDENTIALS=
CDATA_BACKUP_AUSIZE=1
GNS_DENY_NET_LIST=
OLD_CRS_HOME=
GNS_DOMAIN_LIST=
ASM_UPGRADE=false
NETCA_LISTENERS_REGISTERED_WITH_CRS=LISTENER
CDATA_BACKUP_DISKS=
ASMCA_ARGS=
CLUSTER_NODES=node1,node2
```

```
ASM_DIAGNOSTIC_DEST=/u01/app/grid
AFD_LABELS_NO_DG=
AFD_CONFIGURED=false
MGMT_DB=false
SCAN_PORT=1522
ASM_DROP_DISKGROUPS=true
OPC_NAT_ADDRESS=
CLUSTER_TYPE=DB
NETWORKS="eth0"/IP_Address:public,"eth1"/IP_Address:asm,"eth1"/
IP_Address:cluster_interconnect
OCR_VOTINGDISK_IN_ASM=true
HUB_SIZE=32
CDATA_BACKUP_SITES=
CDATA_SIZE=0
REUSEDG=false
ASM_IN_HOME=true
ORACLE_SVCUSER=
HOME_TYPE=CRS
GNS_ADDR_LIST=
CLUSTER_NAME=node1-cluster
AFD_CONF=false
OPC_CLUSTER_TYPE=
VOTING_DISKS=
SILENT=true
SCAN_CREDENTIALS=
TZ=UTC
SCAN_TYPE=LOCAL
GPNP_PA=
DC_HOME=/tmp/deinstall2020-09-23_08-53-16AM/logs
CSS_LEASEDURATION=400
REMOTE_NODES=node2
ASM_SPFILE=
AHF_REPOS_LOC=/u01/app/grid/data/repository
SCAN_NAME=node1-cluster-scan
RIM_NODE_LIST=
EXTENDED_CLUSTER_SITE_GUIDS=
INVENTORY_LOCATION=/u01/app/oraInventory
FASTDEPLOY_MODE=
```

> **Note:**
>
> Do not use quotation marks with variables except in the following cases:
>
> - Around addresses in *CRS_NODEVIPS*:
>
>   ```
>   CRS_NODEVIPS='n1-vip/255.255.252.0/eth0,n2-vip/255.255.252.0/eth0'
>   ```
>
> - Around interface names in *NETWORKS*:
>
>   ```
>   NETWORKS="eth0"/IP_Address:public,"eth1"/IP_Address:asm,"eth1"/
>   IP_Address:cluster_interconnect
>   ```

Deinstallation Example for Grid Infrastructure Management Repository

Use this example to understand how to deinstall Grid Infrastructure Management Repository (GIMR) on an Oracle Standalone Cluster.

1. As the `grid` user, prepare GIMR for deinstallation.

   ```
   $ /u01/app/21.0.0/gimr/bin/mgmtca updateConfig
   ```

2. Run `deinstall` from the *GIMR_home*/deinstall directory. The deinstallation starts without prompting you for the GIMR home path.

   ```
   $ ./deinstall
   ```

   You can generate a deinstallation response file by running `deinstall` with the `-checkonly` flag. Alternatively, you can use the response file template located at *GIMR_home*/deinstall/response/deinstall.rsp.tmpl. If you have a response file, then use the optional flag `-paramfile` to provide a path to the response file.

Migrating Standalone Oracle Grid Infrastructure Servers to a Cluster

If you have an Oracle Database installation using Oracle Restart (that is, an Oracle Grid Infrastructure installation for a standalone server), and you want to configure that server as a cluster member node, then complete the following tasks:

1. List all the Oracle databases on the server with their version, unique name of the database, and Oracle home information:

   ```
   srvctl config database -home
   ```

2. Inspect the Oracle Restart configuration of each database with `srvctl` using the following syntax, where *db_unique_name* is the unique name of the database, and *lsnrname* is the name of the listener:

   ```
   srvctl config database -db db_unique_name
   srvctl config service -db db_unique_name
   srvctl config listener -listener lsnrname
   srvctl config volume -volume volume_name -diskgroup diskgroup_name
   ```

   Write down the configuration information for the server.

3. Stop all of the databases, services, and listeners that you discovered in step 1.

4. Log in as `root`, and change directory to *Grid home*/crs/install. For example:

   ```
   # cd /u01/app/19.0.0/grid/crs/install
   ```

5. Unconfigure the Oracle Grid Infrastructure installation for a standalone server (Oracle Restart), using the following command:

```
# roothas.sh -deconfig -force
```

6. Open the `/etc/oratab` file and remove the entry corresponding to `oracle_restart_home`, in the following format.

```
+ASM:oracle_restart_home:N
```

7. Prepare the server for Oracle Clusterware configuration, as described in this document. In addition, you can install Oracle Grid Infrastructure for a cluster in the same location as Oracle Restart, or in a different location.

   Installing in the Same Location as Oracle Restart

   a. Unlock the Oracle Grid Infrastructure installation for a standalone server (Oracle Restart) home, using the following command:

   ```
   roothas.sh -unlock -hahome oracle_restart_home
   ```

   b. Proceed to step 7.

   Installing in a Different Location than Oracle Restart

   a. Set up Oracle Grid Infrastructure software in the new Grid home software location as described in *Installing Only the Oracle Grid Infrastructure Software*.

   b. Proceed to step 7.

8. Set the environment variables as follows:

```
export oracle_install_asm_UseExistingDG=true or false
export oracle_install_asm_DiskGroupName=disk_group_name
export oracle_install_asm_DiskDiscoveryString=asm_discovery_string
```

   If `oracle_install_asm_UseExistingDG` is set to false, then you do not need to specify other environment variables.

9. As the Oracle Grid Infrastructure installation owner, create and stage the response file for this installation as described in *Recording Response Files*.

10. Complete the installation in the silent mode using the following command:

```
$ Grid_home/gridSetup.sh -silent -responseFile $ORACLE_HOME/GI.rsp
```

11. Run `root.sh`.

12. Mount the Oracle ASM disk group used by Oracle Restart.

13. Enter the `volenable` command to enable all Oracle Restart disk group volumes.

14. Add back Oracle Clusterware services to the Oracle Clusterware home, using the information you wrote down in step 1. For example:

```
/u01/app/grid/product/19.0.0/grid/bin/srvctl add filesystem -device
/dev/asm/db1  -diskgroup ORestartData -volume db1 -mountpointpath
/u01/app/grid/product/19.0.0/db1 -user grid
```

15. Add the Oracle Database for support by Oracle Grid Infrastructure for a cluster, using the configuration information you recorded in step 1. Use the following command syntax, where *db_unique_name* is the unique name of the database on the node, and nodename is the name of the node:

```
srvctl add database -db db_unique_name -spfile spfile_name -pwfile
pwfile_name -oraclehome $ORACLE_HOME -node nodename
```

  a. For example, first verify that the `ORACLE_HOME` environment variable is set to the location of the database home directory.

  b. Next, to add the database name `mydb`, enter the following command:

```
srvctl add database -db mydb -spfile spfile_name -pwfile pwfile_name -
oraclehome $ORACLE_HOME -node node1
```

  c. Add each service to the database, using the command `srvctl add service`. For example, add `myservice` as follows:

```
srvctl add service -db mydb –service myservice –preferred myinstance
```

16. Add nodes to your cluster, as required, using the Oracle Grid Infrastructure installer.

> **✎ See Also:**
>
> *Oracle Clusterware Administration and Deployment Guide* for information about adding nodes to your cluster.

Relinking Oracle Grid Infrastructure for a Cluster Binaries

After installing Oracle Grid Infrastructure for a cluster (Oracle Clusterware and Oracle ASM configured for a cluster), if you need to modify the binaries, then use the following procedure, where `Grid_home` is the Oracle Grid Infrastructure for a cluster home:

> **⚠ Caution:**
>
> Before relinking executables, you must shut down all executables that run in the Oracle home directory that you are relinking. In addition, shut down applications linked with Oracle shared libraries.

As `root`:

```
# cd Grid_home/crs/install
# rootcrs.sh -unlock
```

As the Oracle Grid Infrastructure for a cluster owner:

```
$ export ORACLE_HOME=Grid_home
$ Grid_home/bin/relink
```

As root again:

```
# cd Grid_home/rdbms/install/
# ./rootadd_rdbms.sh
# cd Grid_home/crs/install
# rootcrs.sh -lock
```

You must relink the Oracle Clusterware and Oracle ASM binaries every time you apply an operating system patch or after you perform an operating system upgrade that does not replace the root file system. For an operating system upgrade that results in a new root file system, you must remove the node from the cluster and add it back into the cluster.

For upgrades from previous releases, if you want to deinstall the prior release Grid home, then you must first unlock the prior release Grid home. Unlock the previous release Grid home by running the command `rootcrs.sh -unlock` from the previous release home. After the script has completed, you can run the `deinstall` command.

Changing the Oracle Grid Infrastructure Home Path

After installing Oracle Grid Infrastructure for a cluster (Oracle Clusterware and Oracle ASM configured for a cluster), if you need to change the Grid home path, then use the following example as a guide to detach the old Grid home, and to attach a new Grid home:

> **Note:**
>
> Before changing the Grid home, you must shut down all executables that run in the Grid home directory that you are relinking. In addition, shut down applications linked with Oracle shared libraries.

1. As the `root` user, copy the Oracle Grid Infrastructure binaries from the old Grid home location to the new Grid home location. For example, where the old Grid home is `/u01/app/21.0.0/grid` and the new Grid home is `/u01/app/21c/grid`.

   ```
   # mkdir /u01/app/21c/grid
   # cp -pR /u01/app/21.0.0/grid /u01/app/21c/grid
   ```

2. Change the owner and group for the directory and files in the new Grid home.

   ```
   # chown -R grid:oinstall /u01/app/21c/grid
   ```

3. As the `grid` user, run the `gridSetup.sh` command from the new Grid home directory and in the Select Configuration Option screen, select **Setup Software Only**.

   ```
   $ /u01/app/21c/grid/gridSetup.sh
   ```

   Complete the installation by selecting the installation options as prompted. This step relinks the Oracle Grid Infrastructure binaries and updates the Oracle Inventory (`oraInventory`) with the new Oracle Grid Infrastructure home.

> **Note:**
>
> Do not change the Oracle base path during the installation.

4. As the `root` user, run the `root.sh` script from your new Grid home.

   ```
   # /u01/app/21c/grid/root.sh
   ```

5. As the `grid` user, enable Oracle RAC, dNFS, and `ioracle` from the new Grid home.

   ```
   $ export ORACLE_HOME=/u01/app/21c/grid
   $ cd $ORACLE_HOME/rdbms/lib
   $ make -f ins_rdbms.mk dnfs_on rac_on ioracle ORACLE_HOME=$ORACLE_HOME
   ```

6. As the `root` user, enter the following command to start up in the new home location.

   ```
   # cd Grid_home/crs/install
   # ./rootcrs.sh -move -dstcrshome Grid_home
   ```

7. As the `grid` user, disable CRS from the old Grid home.

   ```
   $ Grid_home/oui/bin/runInstaller -updateNodeList -noClusterEnabled
   ORACLE_HOME=$ORACLE_HOME
   "CLUSTER_NODES={node1,node2}" CRS=false -doNotUpdateNodeList
   LOCAL_NODE=list_of_nodes
   ```

8. Enable CRS from the new Grid home.

   ```
   $ Grid_home/oui/bin/runInstaller -updateNodeList -noClusterEnabled
   ORACLE_HOME=$ORACLE_HOME
   "CLUSTER_NODES={node1,node2}" CRS=true -doNotUpdateNodeList
   LOCAL_NODE=list_of_nodes
   ```

9. Repeat steps 1 through 8 on each cluster member node.

Unconfiguring Oracle Clusterware Without Removing Binaries

Running the `rootcrs.sh` command flags `-deconfig -force` enables you to unconfigure Oracle Clusterware on one or more nodes without removing installed binaries.

This feature is useful if you encounter an error on one or more cluster nodes during installation when running the `root.sh` command, such as a missing operating system package on one node. By running `rootcrs.sh -deconfig -force` on nodes where you encounter an installation error, you can unconfigure Oracle Clusterware on those nodes, correct the cause of the error, and then run `root.sh` again.

> **Note:**
>
> Stop any databases, services, and listeners that may be installed and running before deconfiguring Oracle Clusterware.

> **⚠ Caution:**
>
> Commands used in this section remove the Oracle Grid infrastructure installation for the entire cluster. If you want to remove the installation from an individual node, then see *Oracle Clusterware Administration and Deployment Guide*.

To unconfigure Oracle Clusterware:

1. Log in as the `root` user on a node where you encountered an error.

2. Change directory to *Grid_home*/crs/install. For example:

   ```
   # cd /u01/app/19.0.0/grid/crs/install
   ```

3. Run `rootcrs.sh` with the `-deconfig` and `-force` flags. For example:

   ```
   # ./rootcrs.sh -deconfig -force
   ```

   Repeat on other nodes as required.

4. If you are deconfiguring Oracle Clusterware on all nodes in the cluster, then on the last node, enter the following command:

   ```
   # ./rootcrs.sh -deconfig -force -lastnode
   ```

   The `-lastnode` flag completes deconfiguration of the cluster, including the OCR and voting files.

A

Completing Preinstallation Tasks Manually

You can complete the preinstallation configuration tasks manually.

Oracle recommends that you use Oracle Universal Installer and Cluster Verification Utility fixup scripts to complete minimal configuration settings. If you cannot use fixup scripts, then complete minimum system settings manually.

- Configuring SSH Manually on All Cluster Nodes
  SSH is used during installation to configure cluster member nodes, and SSH is used after installation by configuration assistants, Oracle Enterprise Manager, Opatch, and other features.

- Configuring Kernel Parameters on HP-UX Itanium Systems
  These topics explain how to configure kernel parameters manually for HP-UX Itanium systems if you cannot complete them using the fixup scripts.

Configuring SSH Manually on All Cluster Nodes

SSH is used during installation to configure cluster member nodes, and SSH is used after installation by configuration assistants, Oracle Enterprise Manager, Opatch, and other features.

- About Passwordless SSH Configuration
  Passwordless SSH configuration is a mandatory installation requirement for Oracle Grid Infrastructure installation to configure cluster member nodes.

- Checking Existing SSH Configuration on the System
  Use this procedure to check if SSH is already configured and running on your server.

- Configuring SSH on Cluster Nodes
  You must configure SSH separately for each Oracle software installation owner that you intend to use for installation.

- Enabling SSH User Equivalency on Cluster Nodes
  After you have copied the authorized_keys file that contains all keys to each node in the cluster, complete the following procedure.

About Passwordless SSH Configuration

Passwordless SSH configuration is a mandatory installation requirement for Oracle Grid Infrastructure installation to configure cluster member nodes.

Automatic Passwordless SSH configuration using OUI creates RSA encryption keys on all nodes of the cluster. If you have system restrictions that require you to set up SSH manually, such as using DSA keys, then use this procedure as a guide to set up passwordless SSH. If SSH is not available, then Oracle Universal Installer (OUI) attempts to use `rcp` instead. However, these services are disabled by default on most Linux systems.

To configure SSH, you must first create RSA or DSA keys on each cluster node, and then copy all the keys generated on all cluster node members into an authorized keys file that is identical on each node. Note that the SSH files must be readable only by `root` and by the

software installation user (`oracle`, `grid`), as SSH ignores a private key file if it is accessible by others. In the examples that follow, the DSA key is used.

> **Note:**
>
> The supported version of SSH for Linux distributions is OpenSSH.

Checking Existing SSH Configuration on the System

Use this procedure to check if SSH is already configured and running on your server.

You need either an RSA or a DSA key for the SSH protocol. RSA is used with the SSH 1.5 protocol, while DSA is the default for the SSH 2.0 protocol. With OpenSSH, you can use either RSA or DSA. The instructions that follow are for SSH1. If you have an SSH2 installation, and you cannot use SSH1, then refer to your SSH distribution documentation to configure SSH1 compatibility or to configure SSH2 with DSA.

1. To determine if SSH is running, enter the following command:

   ```
   $ pgrep sshd
   ```

   If SSH is running, then the response to this command is one or more process ID numbers.

2. In the home directory of the installation software owner (`grid` or `oracle`), use the command `ls -al` to ensure that the `.ssh` directory is owned and writable only by the user.

Configuring SSH on Cluster Nodes

You must configure SSH separately for each Oracle software installation owner that you intend to use for installation.

- Create SSH Directory and Create SSH Keys On Each Node
  To configure SSH, you must first create RSA or DSA keys on each cluster node.

- Add All Keys to a Common authorized_keys File
  To configure SSH, copy all the generated keys on all cluster node members into an authorized keys file that is identical on each node.

Create SSH Directory and Create SSH Keys On Each Node

To configure SSH, you must first create RSA or DSA keys on each cluster node.

Complete the following steps on each node:

1. Log in as the software owner (in this example, the `grid` user).

2. To ensure that you are logged in as the `grid` user, and to verify that the user ID matches the expected user ID you have assigned to the `grid` user, enter the commands:

   ```
   $ id
   $ id grid
   ```

Ensure that Oracle user group and user and the user terminal window process you are using have group and user IDs are identical.

For example:

```
uid=54322(grid) gid=54321(oinstall)
groups=54321(oinstall),54322(grid,asmadmin,asmdba)
$ id grid uid=54322(grid) gid=54321(oinstall)
groups=54321(oinstall),54322(grid,asmadmin,asmdba)
```

3. If necessary, create the `.ssh` directory in the `grid` user's home directory, and set permissions on it to ensure that only the oracle user has read and write permissions:

```
$ mkdir ~/.ssh
$ chmod 700 ~/.ssh
```

Note that the SSH configuration fails if the permissions are not set to 700.

4. Enter the following command:

```
$ /usr/bin/ssh-keygen -t dsa
```

> **Note:**
>
> If you have OpenSSH version 7.8 or higher installed on your system, then enter the following command to create SSH keys on each node:
>
> ```
> $ /usr/bin/ssh-keygen -t dsa -m PEM
> ```

At the prompts, accept the default location for the key file (press **Enter**).

Never distribute the private key to anyone not authorized to perform Oracle software installations.

This command writes the DSA public key to the `~/.ssh/id_dsa.pub` file and the private key to the `~/.ssh/id_dsa` file.

5. Repeat steps 1 through 4 on each node that you intend to make a member of the cluster, using the DSA key.

Add All Keys to a Common authorized_keys File

To configure SSH, copy all the generated keys on all cluster node members into an authorized keys file that is identical on each node.

Complete the following steps:

1. On the local node, change directories to the `.ssh` directory in the Oracle Grid Infrastructure owner's home directory (typically, either `grid` or `oracle`). Then, add the DSA key to the `authorized_keys` file using the following commands:

```
$ cat id_dsa.pub >> authorized_keys
$ ls
```

In the `.ssh` directory, you should see the `id_dsa.pub` keys that you have created, and the file `authorized_keys`.

2. On the local node, use SCP (Secure Copy) or SFTP (Secure FTP) to copy the `authorized_keys` file to the `oracle` user `.ssh` directory on a remote node. The following example is with SCP, on a node called node2, with the Oracle Grid Infrastructure owner `grid`, where the `grid` user path is `/home/grid`:

```
[grid@node1 .ssh]$ scp authorized_keys node2:/home/grid/.ssh/
```

   a. You are prompted to accept a DSA key. Enter Yes, and you see that the node you are copying to is added to the `known_hosts` file.

   b. When prompted, provide the password for the grid user, which should be the same on all nodes in the cluster. The `authorized_keys` file is copied to the remote node.

   Your output should be similar to the following, where `xxx` represents parts of a valid IP address:

```
[grid@node1 .ssh]$ scp authorized_keys node2:/home/grid/.ssh/
The authenticity of host 'node2 (xxx.xxx.173.152) can't be
established.
DSA key fingerprint is
7e:60:60:ae:40:40:d1:a6:f7:4e:zz:me:a7:48:ae:f6:7e.
Are you sure you want to continue connecting (yes/no)? yes
Warning: Permanently added 'node1,xxx.xxx.173.152' (dsa) to the list
of known hosts
grid@node2's password:
authorized_keys          100%       828       7.5MB/s       00:00
```

3. Using SSH, log in to the node where you copied the `authorized_keys` file. Then change to the `.ssh` directory, and using the `cat` command, add the DSA keys for the second node to the `authorized_keys` file, clicking **Enter** when you are prompted for a password, so that passwordless SSH is set up:

```
[grid@node1 .ssh]$ ssh node2 [grid@node2 grid]$ cd .ssh [grid@node2
ssh]
$ cat id_dsa.pub  >> authorized_keys
```

4. Repeat steps 2 and 3 from each node to each other member node in the cluster.

5. When you have added keys from each cluster node member to the `authorized_keys` file on the last node you want to have as a cluster node member, then use `scp` to copy the `authorized_keys` file with the keys from all nodes back to each cluster node member, overwriting the existing version on the other nodes. To confirm that you have all nodes in the `authorized_keys` file, enter the command `more authorized_keys`, and determine if there is a DSA key for each member

node. The file lists the type of key (`ssh-dsa`), followed by the key, and then followed by the user and server. For example:

```
ssh-dsa AAAABBBB . . . = grid@node1
```

The `grid` user's `/.ssh/authorized_keys` file on every node must contain the contents from all of the `/.ssh/id_dsa.pub` files that you generated on all cluster nodes.

Enabling SSH User Equivalency on Cluster Nodes

After you have copied the authorized_keys file that contains all keys to each node in the cluster, complete the following procedure.

In this example, the Oracle Grid Infrastructure software owner is named `grid`.

Do the following:

1. On the system where you want to run OUI, log in as the `grid` user.

2. Use the following command syntax, where *hostname1*, *hostname2*, and so on, are the public host names (alias and fully qualified domain name) of nodes in the cluster to run SSH from the local node to each node, including from the local node to itself, and from each node to each other node:

   ```
   [grid@nodename]$ ssh hostname1 date [grid@nodename]$ ssh hostname2
   date      .    .    .
   ```

   At the end of this process, the public host name for each member node should be registered in the `known_hosts` file for all other cluster nodes. If you are using a remote client to connect to the local node, and you see a message similar to "`Warning: No xauth data; using fake authentication data for X11 forwarding,`" then this means that your authorized keys file is configured correctly, but your SSH configuration has X11 forwarding enabled. To correct this issue, see Setting Remote Display and X11 Forwarding Configuration.

3. Repeat step 2 on each cluster node member.

If you have configured SSH correctly, then you can now use the `ssh` or `scp` commands without being prompted for a password. For example:

```
[grid@node1 ~]$ ssh node2 date
Mon Feb 26 23:34:42 UTC 2009
[grid@node1 ~]$ ssh node1 date
Mon Feb 26 23:34:48 UTC 2009
```

If any node prompts for a password, then verify that the `~/.ssh/authorized_keys` file on that node contains the correct public keys, and that you have created an Oracle software owner with identical group membership and IDs.

Configuring Kernel Parameters on HP-UX Itanium Systems

These topics explain how to configure kernel parameters manually for HP-UX Itanium systems if you cannot complete them using the fixup scripts.

- **Minimum Parameter Settings for Installation**
  Use this table to set parameters manually if you cannot use the fixup scripts. Verify that the kernel parameters are set to values greater than or equal to the minimum values shown in the following table.

- **Checking Kernel Parameter Values**
  Use the `kcweb` application to check and modify the formula or value kernel parameters on an HP-UX Itanium system.

- **Setting UDP and TCP Kernel Parameters Manually**
  If you do not use a Fixup script or CVU to set ephemeral ports, then set TCP/IP ephemeral port range parameters to provide enough ephemeral ports for the anticipated server workload.

Minimum Parameter Settings for Installation

Use this table to set parameters manually if you cannot use the fixup scripts. Verify that the kernel parameters are set to values greater than or equal to the minimum values shown in the following table.

Table A-1   Minimum HP-UX Itanium Kernel Parameter Settings

Parameter	Minimum Value
executable_stack	0
ksi_alloc_max	32768
max_thread_proc	1024
maxdsiz	1073741824 (1 GB)
maxdsiz_64bit	2147483648 (2 GB)
maxfiles	1024
maxfiles_lim	63488
maxssiz	134217728 (128 MB)
maxssiz_64bit	1073741824 (1 GB)
maxuprc	3686
msgmni	4096
msgtql	4096
ncsize	35840
nflocks	4096
ninode	34816
nkthread	7184
nproc	4096
semmni	4096
semmns	8192
semmnu	4092
semvmx	32767
shmmax	1073741824
shmmni	4096
shmseg	512

Table A-1 (Cont.) Minimum HP-UX Itanium Kernel Parameter Settings

Parameter	Minimum Value
tcp_largest_anon_port	65500
udp_largest_anon_port	65500

Guidelines for Setting Kernel Parameter Values

- If the current value for any parameter is higher than the value listed in this table, then do not change the value of that parameter.

- Do not specify values for the following parameters as they are obsolete on HP-UX 11.31:

    – msgmap

    – msgseg

- If you do not use HFS, then retain the default ninode value.

Checking Kernel Parameter Values

Use the kcweb application to check and modify the formula or value kernel parameters on an HP-UX Itanium system.

1. Start the kcweb application:

   ```
   # /usr/sbin/kcweb -F
   ```

2. Check the value or formula specified for each of these parameters and, if necessary, modify that value or formula.

3. Refer to the kcweb online help for more information about completing this step.

> **✐ Note:**
>
> If you modify the value of a parameter that is not dynamic, then restart the system.

Setting UDP and TCP Kernel Parameters Manually

If you do not use a Fixup script or CVU to set ephemeral ports, then set TCP/IP ephemeral port range parameters to provide enough ephemeral ports for the anticipated server workload.

Check your current range for ephemeral ports:

```
# /usr/bin/ndd /dev/tcp tcp_largest_anon_port

65535
```

In the preceding example, tcp_largest_anon_port is set to the default value.

```
/etc/rc.config.d/nddconf
```

```
TRANSPORT_NAME[0]=tcp
NDD_NAME[0]=tcp_largest_anon_port
NDD_VALUE[0]=65500
```

```
TRANSPORT_NAME[1]=udp
NDD_NAME[1]=udp_largest_anon_port
NDD_VALUE[1]=65500
```

Ensure that the entries are numbered in proper order. For example, if there are two entries present for the TCP and UDP ports in `nddconf`, then they are numbered 0 through 1: `TRANSPORT_NAME[0]=tcp` and `TRANSPORT_NAME[1]=udp`.

B

Installing and Configuring Oracle Database Using Response Files

Review the following topics to install and configure Oracle products using response files.

- **Reasons for Using Silent Mode or Response File Mode**
  Review this section for use cases for running the installer in silent mode or response file mode.

- **Using Response Files**
  Use response files to install and configure Oracle products using the installer in silent or response file mode.

- **Preparing Response Files**
  Review this information to prepare response files for use during silent mode or response file mode installations.

- **Running Oracle Universal Installer Using a Response File**
  After creating the response file, run Oracle Universal Installer at the command line, specifying the response file you created, to perform the installation.

- **Running Net Configuration Assistant Using Response Files**
  You can run NETCA in silent mode to configure and start an Oracle Net Listener on the system, and configure naming methods and Oracle Net service names.

- **Postinstallation Configuration Using Response File Created During Installation**
  Use response files to configure Oracle software after installation. You can use the same response file created during installation to also complete postinstallation configuration.

Reasons for Using Silent Mode or Response File Mode

Review this section for use cases for running the installer in silent mode or response file mode.

Mode	Uses
Silent	Use silent mode for the following installations:
	• Complete an unattended installation, which you schedule using operating system utilities such as `at`.
	• Complete several similar installations on multiple systems without user interaction.
	• Install the software on a system that does not have X Window System software installed on it.
	The installer displays progress information on the terminal that you used to start it, but it does not display any of the installer screens.
Response file	Use response file mode to complete similar Oracle software installations on more than one system, providing default answers to some, but not all of the installer prompts.

Using Response Files

Use response files to install and configure Oracle products using the installer in silent or response file mode.

> **Note:**
>
> You must complete all required preinstallation tasks on a system before running the installer in silent or response file mode.

1. Prepare a response file.
2. Run the installer in silent or response file mode.
3. Run the root scripts as prompted by Oracle Universal Installer.
4. If you completed a software-only installation, then run Net Configuration Assistant and Oracle DBCA in silent or response file mode to create the database listener and an Oracle Database instance respectively.

Preparing Response Files

Review this information to prepare response files for use during silent mode or response file mode installations.

- Editing a Response File Template
  Oracle provides response file templates for each product and each configuration tool.

- Recording Response Files
  You can use OUI in interactive mode to record response files, which you can then edit and use to complete silent mode or response file mode installations.

Editing a Response File Template

Oracle provides response file templates for each product and each configuration tool.

About Response File Templates

For Oracle Database, the response file templates are located in the `$ORACLE_HOME/install/response` directory. For Oracle Grid Infrastructure, the response file templates are located in the `Grid_home/install/response` directory.

Where, *Grid_home* is the Oracle Grid Infrastructure home directory path.

> **Note:**
>
> If you copied the software to a hard disk, then the response files are located in the `$ORACLE_HOME/install/response` directory.

All response file templates contain comment entries, sample formats, examples, and other useful instructions. Read the response file instructions to understand how to specify values for the response file variables, so that you can customize your installation.

The following table lists the response files provided with this software:

Table B-1   Response Files for Oracle Database and Oracle Grid Infrastructure

Response File	Description
db_install.rsp	Silent installation of Oracle Database.
dbca.rsp	Silent creation and configuration of Oracle Database using Oracle DBCA.
netca.rsp	Silent configuration of Oracle Net using Oracle NETCA.
gridsetup.rsp	Silent configuration of Oracle Grid Infrastructure installations.

> ⚠️ **Caution:**
>
> When you modify a response file template and save a file for use, the response file may contain plain text passwords. Ownership of the response file should be given to the Oracle software installation owner only, and permissions on the response file should be changed to 600. Oracle strongly recommends that database administrators or other administrators delete or secure response files when they are not in use.

To copy and modify a response file:

1. Copy the response file from the response file directory to a directory on your system:

   For example, for Oracle Database:

   ```
   $ cp $ORACLE_HOME/install/response/db_install.rsp local_directory
   ```

2. Open the response file in a text editor:

   ```
   $ vi /local_directory/db_install.rsp
   ```

3. Follow the instructions in the file to edit it.

   > ✎ **Note:**
   >
   > The installer or configuration assistant fails if you do not correctly configure the response file. Also, ensure that your response file name has the `.rsp` suffix.

4. Secure the response file by changing the permissions on the file to 600:

   ```
   $ chmod 600 /local_dir/db_install.rsp
   ```

   Ensure that only the Oracle software owner user can view or modify response files or consider deleting them after the installation succeeds.

> **Note:**
>
> A fully-specified response file for an Oracle Database installation contains the passwords for database administrative accounts and for a user who is a member of the OSDBA group (required for automated backups).

Recording Response Files

You can use OUI in interactive mode to record response files, which you can then edit and use to complete silent mode or response file mode installations.

Recording response files is useful for Advanced or software-only installations. You can save all the installation steps into a response file during installation by clicking **Save Response File** on the **Summary** page. You can use the generated response file for a silent installation later. When you record the response file, you can either complete the installation, or you can exit from the installer on the Summary page, before OUI starts to set up the software to the system.

If you use record mode during a response file mode installation, then the installer records the variable values that were specified in the original source response file into the new response file.

> **Note:**
>
> OUI does not save passwords while recording the response file.

1. Complete preinstallation tasks as for a standard installation.

   When you run the installer to record a response file, it checks the system to verify that it meets the requirements to install the software. For this reason, Oracle recommends that you complete all of the required preinstallation tasks and record the response file while completing an installation.

2. Ensure that the Oracle software owner user (typically `oracle`) has permissions to create or write to the Oracle home path that you specify when you run the installer.

3. On each installation screen, specify the required information.

4. When the installer displays the **Summary** screen, perform the following steps:

   a. Click **Save Response File**. In the window, specify a file name and location for the new response file. Click **Save** to write the responses you entered to the response file.

   b. Click **Finish** to continue with the installation. Click **Cancel** if you do not want to continue with the installation. The installation stops, but the recorded response file is retained.

   > **Note:**
   >
   > Ensure that your response file name has the `.rsp` suffix.

5. Before you use the saved response file on another system, edit the file and make any required changes. Use the instructions in the file as a guide when editing it.

Running Oracle Universal Installer Using a Response File

After creating the response file, run Oracle Universal Installer at the command line, specifying the response file you created, to perform the installation.

Run Oracle Universal Installer at the command line, specifying the response file you created. The Oracle Universal Installer executables, `runInstaller` and `gridSetup.sh`, provide several options. For help information on the full set of these options, run the `gridSetup.sh` or `runInstaller` command with the `-help` option. For example:

- For Oracle Database:

```
$ $ORACLE_HOME/runInstaller -help
```

- For Oracle Grid Infrastructure:

```
$ /u01/app/21.0.0/grid/gridSetup.sh -help
```

The help information appears in a window after some time.

Perform the following steps to run the installer using a response file:

1. Complete the preinstallation tasks for a normal installation.

2. Log in as the software installation owner user.

3. If you are completing a response file mode installation, then set the operating system `DISPLAY` environment variable for the user running the installation.

> **Note:**
>
> You do not have to set the `DISPLAY` environment variable if you are completing a silent mode installation.

4. To start the installer in silent or response file mode, enter a command similar to the following:

- For Oracle Database:

```
$ $ORACLE_HOME/runInstaller [-silent] -responseFile responsefilename
```

- For Oracle Grid Infrastructure:

```
$ /u01/app/21.0.0/grid/gridSetup.sh [-silent] -responseFile
responsefilename
```

> **Note:**
>
> Do not specify a relative path to the response file. If you specify a relative path, then the installer fails.

In this example:

- `-silent` runs the installer in silent mode.

- *responsefilename* is the full path and file name of the installation response file that you configured.

5. If this is the first time you are installing Oracle software on your system, then Oracle Universal Installer prompts you to run the `orainstRoot.sh` script. Log in as the `root` user and run the `orainstRoot.sh` script:

```
$ su root
password:
# /u01/app/oraInventory/orainstRoot.sh
```

> **Note:**
>
> You do not have to manually create the `oraInst.loc` file. Running the `orainstRoot.sh` script is sufficient as it specifies the location of the Oracle Inventory directory.

6. When the installation completes, log in as the `root` user and run the `root.sh` script. For example:

```
$ su root
password:
# $ORACLE_HOME/root.sh
```

Running Net Configuration Assistant Using Response Files

You can run NETCA in silent mode to configure and start an Oracle Net Listener on the system, and configure naming methods and Oracle Net service names.

To run Oracle Net Configuration Assistant (NETCA) in silent mode, you must copy and edit a response file template. Oracle provides a response file template named `netca.rsp` in the `$ORACLE_HOME/assistants/netca` directory.

To run Net Configuration Assistant using a response file:

1. Copy the `netca.rsp` response file template from the response file directory to a directory on your system:

```
$ cp $ORACLE_HOME/assistants/netca/netca.rsp local_directory
```

2. Open the response file in a text editor:

```
$ vi /local_directory/netca.rsp
```

3. Follow the instructions in the file to edit it.

> **Note:**
>
> Net Configuration Assistant fails if you do not correctly configure the response file.

4. Log in as the Oracle software owner user, and set the `ORACLE_HOME` environment variable to specify the correct Oracle home directory.

5. Enter a command similar to the following to run Net Configuration Assistant in silent mode:

```
$ $ORACLE_HOME/bin/netca /silent /responsefile /local_directory/netca.rsp
```

In this command:

- The `/silent` option indicates to run Net Configuration Assistant in silent mode.
- *local_directory* is the full path of the directory where you copied the `netca.rsp` response file template.

Postinstallation Configuration Using Response File Created During Installation

Use response files to configure Oracle software after installation. You can use the same response file created during installation to also complete postinstallation configuration.

- Using the Installation Response File for Postinstallation Configuration
  You can use the response file created during installation to also complete postinstallation configuration.
- Running Postinstallation Configuration Using Response File
  You can use a response file to complete postinstallation tasks on one or more servers simultaneously.

Using the Installation Response File for Postinstallation Configuration

You can use the response file created during installation to also complete postinstallation configuration.

Run the installer with the `-executeConfigTools` option to configure configuration assistants after installing Oracle Grid Infrastructure or Oracle Database. You can use the response file located at `$ORACLE_HOME/install/response/`*product_timestamp.rsp* to obtain the passwords required to run the configuration tools. You must update the response file with the required passwords before running the `-executeConfigTools` command.

Oracle strongly recommends that you maintain security with a password response file:

- Permissions on the response file should be set to 600.
- The owner of the response file should be the installation owner user, with the group set to the central inventory (`oraInventory`) group.

Example B-1   Response File Passwords for Oracle Grid Infrastructure (`grid` user)

```
grid.install.crs.config.ipmi.bmcPassword=password
grid.install.asm.SYSASMPassword=password
grid.install.asm.monitorPassword=password
grid.install.config.emAdminPassword=password
```

If you do not have a BMC card, or you do not want to enable IPMI, then leave the `ipmi.bmcPassword` input field blank.

If you do not want to enable Oracle Enterprise Manager for management, then leave the `emAdminPassword` password field blank.

Example B-2   Response File Passwords for Oracle Grid Infrastructure for a Standalone Server (`oracle` user)

```
oracle.install.asm.SYSASMPassword=password
oracle.install.asm.monitorPassword=password
oracle.install.config.emAdminPassword=password
```

If you do not want to enable Oracle Enterprise Manager for management, then leave the `emAdminPassword` password field blank.

Example B-3   Response File Passwords for Oracle Database (`oracle` user)

This example illustrates the passwords to specify for use with the database configuration assistants.

```
oracle.install.db.config.starterdb.password.SYS=password
oracle.install.db.config.starterdb.password.SYSTEM=password
oracle.install.db.config.starterdb.password.DBSNMP=password
oracle.install.db.config.starterdb.password.PDBADMIN=password
oracle.install.db.config.starterdb.emAdminPassword=password
oracle.install.db.config.asm.ASMSNMPPassword=password
```

You can also specify
`oracle.install.db.config.starterdb.password.ALL=`*password* to use the same password for all database users.

The database configuration assistants require the SYS, SYSTEM, PDBADMIN, and DBSNMP passwords for use with Oracle DBCA. You must specify the following passwords, depending on your system configuration:

- If the database uses Oracle Automatic Storage Management (Oracle ASM) for storage, then you must specify a password for the `ASMSNMPPassword` variable. If you are not using Oracle ASM, then leave the value for this password variable blank.

Running Postinstallation Configuration Using Response File

You can use a response file to complete postinstallation tasks on one or more servers simultaneously.

Complete this procedure to run configuration assistants with the `executeConfigTools` command and a response file.

1. Edit the response file and specify the required passwords for your configuration. You can use the response file created during installation, located at `$ORACLE_HOME/install/response/`*product_timestamp.rsp*.

2. Change directory to the Oracle home containing the installation software. For example:

   For Oracle Grid Infrastructure:

   ```
   cd Grid_home
   ```

   Where, *Grid_home* is the path to the Oracle Grid Infrastructure home directory `/u01/app/21.0.0/grid`

   For Oracle Database:

   ```
   cd $ORACLE_HOME
   ```

3. Run the configuration script using the following syntax:

   For Oracle Grid Infrastructure:

   ```
   $ ./gridSetup.sh -executeConfigTools -responseFile Grid_home/install/
   response/product_timestamp.rsp
   ```

   For Oracle Database:

   ```
   $ ./runInstaller -executeConfigTools -responseFile $ORACLE_HOME/install/
   response/product_timestamp.rsp
   ```

   For Oracle Database, you can also run the response file located in the directory `$ORACLE_HOME/install/response/`:

   ```
   $ ./runInstaller -executeConfigTools -responseFile $ORACLE_HOME/install/
   response/db_install.rsp
   ```

   The postinstallation configuration tool runs the installer in the graphical user interface mode, displaying the progress of the postinstallation configuration. Specify the `[-silent]` option to run the postinstallation configuration in the silent mode.

   For example, for Oracle Grid Infrastructure:

   ```
   $ ./gridSetup.sh -executeConfigTools -responseFile /u01/app/21.0.0/grid/
   install/response/grid_2019-01-09_01-03-36PM.rsp [-silent]
   ```

For Oracle Database:

```
$ ./runInstaller -executeConfigTools -responseFile /u01/app/oracle/
product/21.0.0/dbhome_1/response/db_2019-01-09_01-03-36PM.rsp [-
silent]
```

C

Optimal Flexible Architecture

Oracle Optimal Flexible Architecture (OFA) rules are a set of configuration guidelines created to ensure well-organized Oracle installations, which simplifies administration, support and maintenance.

- About the Optimal Flexible Architecture Standard
  Oracle Optimal Flexible Architecture (OFA) rules help you to organize database software and configure databases to allow multiple databases, of different versions, owned by different users to coexist.

- About Multiple Oracle Homes Support
  Oracle Database supports multiple Oracle homes. You can install this release or earlier releases of the software more than once on the same system, in different Oracle home directories.

- About the Oracle Inventory Directory and Installation
  The directory that you designate as the Oracle Inventory directory (oraInventory) stores an inventory of all software installed on the system.

- Oracle Base Directory Naming Convention
  The Oracle Base directory is the database home directory for Oracle Database installation owners, and the log file location for Oracle Grid Infrastructure owners.

- Oracle Home Directory Naming Convention
  By default, Oracle Universal Installer configures Oracle home directories using these Oracle Optimal Flexible Architecture conventions.

- Optimal Flexible Architecture File Path Examples
  Review examples of hierarchical file mappings of an Optimal Flexible Architecture-compliant installation.

About the Optimal Flexible Architecture Standard

Oracle Optimal Flexible Architecture (OFA) rules help you to organize database software and configure databases to allow multiple databases, of different versions, owned by different users to coexist.

In earlier Oracle Database releases, the OFA rules provided optimal system performance by isolating fragmentation and minimizing contention. In current releases, OFA rules provide consistency in database management and support, and simplifies expanding or adding databases, or adding additional hardware.

By default, Oracle Universal Installer places Oracle Database components in directory locations and with permissions in compliance with OFA rules. Oracle recommends that you configure all Oracle components in accordance with OFA guidelines.

Oracle recommends that you accept the OFA default. Following OFA rules is especially of value if the database is large, or if you plan to have multiple databases.

> **Note:**
>
> OFA assists in identification of an ORACLE_BASE with its Automatic Diagnostic Repository (ADR) diagnostic data to properly collect incidents.

About Multiple Oracle Homes Support

Oracle Database supports multiple Oracle homes. You can install this release or earlier releases of the software more than once on the same system, in different Oracle home directories.

Careful selection of mount point names can make Oracle software easier to administer. Configuring multiple Oracle homes in compliance with Optimal Flexible Architecture (OFA) rules provides the following advantages:

- You can install this release, or earlier releases of the software, more than once on the same system, in different Oracle home directories. However, you cannot install products from one release of Oracle Database into an Oracle home directory of a different release.

- Multiple databases, of different versions, owned by different users can coexist concurrently.

- To install Oracle Database software in multiple Oracle homes, you must extract the image file in each Oracle home, and then run the setup wizard from the respective Oracle home.

- You must install a new Oracle Database release in a new Oracle home that is separate from earlier releases of Oracle Database.

  You cannot install multiple releases in one Oracle home. Oracle recommends that you create a separate Oracle Database Oracle home for each release, in accordance with the Optimal Flexible Architecture (OFA) guidelines.

- In production, the Oracle Database server software release is the release number in the format of major and RU release number. For example, with the release number 19.3.0.0.0, the major release is 19 and the RU release number is 3.

- Later Oracle Database releases can access earlier Oracle Database releases. However, this access is only for upgrades. For example, Oracle Database 19c can access an Oracle Database 18c if the 18c database is started up in upgrade mode.

- Oracle Database Client can be installed in the same Oracle Database home if both products are at the same release level. For example, you can install Oracle Database Client 19c into an existing Oracle Database 19c home but you cannot install Oracle Database Client 19c into an existing Oracle Database 18c home. If you apply a patch set or release update before installing the client, then you must apply the patch set or release update again.

- Structured organization of directories and files, and consistent naming for database files simplify database administration.

- Login home directories are not at risk when database administrators add, move, or delete Oracle home directories.

- You can test software upgrades in an Oracle home in a separate directory from the Oracle home where your production database is located.

- For information about release support timelines, refer to My Oracle Support Doc ID 742060.1

Related Topics

- My Oracle Support Note 742060.1

About the Oracle Inventory Directory and Installation

The directory that you designate as the Oracle Inventory directory (oraInventory) stores an inventory of all software installed on the system.

All Oracle software installation owners on a server are granted the OINSTALL privileges to read and write to this directory. If you have previous Oracle software installations on a server, then additional Oracle software installations detect this directory from the `/var/opt/oracle/oraInst.loc` file, and continue to use that Oracle Inventory. Ensure that the group designated as the OINSTALL group is available as a primary group for all planned Oracle software installation owners.

If you are installing Oracle software for the first time, then OUI creates an Oracle base and central inventory, and creates an Oracle inventory using information in the following priority:

- In the path indicated in the `ORACLE_BASE` environment variable set for the installation owner user account

- In an Optimal Flexible Architecture (OFA) path (u[01–99]`/app/`*owner* where *owner* is the name of the user account running the installation), and that user account has permissions to write to that path

- In the user home directory, in the path `/app/`*owner*, where *owner* is the name of the user account running the installation

For example:

If you are performing an Oracle Database installation, and you set `ORACLE_BASE` for user `oracle` to the path `/u01/app/oracle` before installation, and grant 755 permissions to `oracle` for that path, then Oracle Universal Installer creates the Oracle Inventory directory one level above the ORACLE_BASE in the path `ORACLE_BASE/../oraInventory`, so the Oracle Inventory path is `/u01/app/oraInventory`. Oracle Universal Installer installs the software in the `ORACLE_BASE` path. If you are performing an Oracle Grid Infrastructure for a Cluster installation, then the Grid installation path is changed to `root` ownership after installation, and the Grid home software location should be in a different path from the Grid user Oracle base.

If you create the OFA path `/u01`, and grant `oracle` 755 permissions to write to that path, then the Oracle Inventory directory is created in the path `/u01/app/oraInventory`, and Oracle Universal Installer creates the path `/u01/app/oracle`, and configures the `ORACLE_BASE` environment variable for the Oracle user to that path. If you are performing an Oracle Database installation, then the Oracle home is installed under the Oracle base. However, if you are installing Oracle Grid Infrastructure for a cluster, then be aware that ownership of the path for the Grid home is changed to `root` after installation and the Grid base and Grid home should be in different locations, such as `/u01/app/21.0.0/grid` for the Grid home path, and `/u01/app/grid` for the Grid base. For example:

`/u01/app/oraInventory`, owned by `grid:oinstall`
`/u01/app/oracle`, owned by `oracle:oinstall`
`/u01/app/oracle/product/21.0.0/dbhome_1/`, owned by `oracle:oinistall`

`/u01/app/grid`, owned by `grid:oinstall`
`/u01/app/21.0.0/grid`, owned by `root`

If you have neither set ORACLE_BASE, nor created an OFA-compliant path, then the Oracle Inventory directory is placed in the home directory of the user that is performing the installation, and the Oracle software is installed in the path `/app`/*owner*, where *owner* is the Oracle software installation owner. For example:

```
/home/oracle/oraInventory
/home/oracle/app/oracle/product/21.0.0/dbhome_1
```

Oracle Base Directory Naming Convention

The Oracle Base directory is the database home directory for Oracle Database installation owners, and the log file location for Oracle Grid Infrastructure owners.

Name Oracle base directories using the syntax */pm/h/u*, where *pm* is a string mount point name, *h* is selected from a small set of standard directory names, and *u* is the name of the owner of the directory.

You can use the same Oracle base directory for multiple installations. If different operating system users install Oracle software on the same system, then you must create a separate Oracle base directory for each installation owner. For ease of administration, Oracle recommends that you create a unique owner for each Oracle software installation owner, to separate log files.

Because all Oracle installation owners write to the central Oracle inventory file, and that file mountpoint is in the same mount point path as the initial Oracle installation, Oracle recommends that you use the same */pm/h* path for all Oracle installation owners.

Table C-1   Examples of OFA-Compliant Oracle Base Directory Names

Example	Description
`/u01/app/` `oracle`	Oracle Database Oracle base, where the Oracle Database software installation owner name is `oracle`. The Oracle Database binary home is located underneath the Oracle base path.
`/u01/app/` `grid`	Oracle Grid Infrastructure Oracle base, where the Oracle Grid Infrastructure software installation owner name is `grid`.

> ⚠️ **Caution:**
>
> The Oracle Grid Infrastructure Oracle base should not contain the Oracle Grid Infrastructure binaries for an Oracle Grid Infrastructure for a cluster installation. Permissions for the file path to the Oracle Grid Infrastructure binary home are changed to `root` during installation.

> **✎ Note:**
>
> Oracle home or Oracle base cannot be `symlinks`, nor can any of their parent directories, all the way to up to the `root` directory.

Oracle Home Directory Naming Convention

By default, Oracle Universal Installer configures Oracle home directories using these Oracle Optimal Flexible Architecture conventions.

The directory pattern syntax for Oracle homes is /*pm*/*s*/*u*/product/*v*/*type*_[*n*]. The following table describes the variables used in this syntax:

Variable	Description
pm	A mount point name.
s	A standard directory name.
u	The name of the owner of the directory.
v	The version of the software.
type	The type of installation. For example: Database (`dbhome`), Client (`client`), or Oracle Grid Infrastructure (`grid`)
n	An optional counter, which enables you to install the same product more than once in the same Oracle base directory. For example: Database 1 and Database 2 (`dbhome_1`, `dbhome_2`)

For example, the following path is typical for the first installation of Oracle Database on this system:

```
/u01/app/oracle/product/21.0.0/dbhome_1
```

> **✎ Note:**
>
> Oracle home or Oracle base cannot be `symlinks`, nor can any of their parent directories, all the way to up to the `root` directory.

Optimal Flexible Architecture File Path Examples

Review examples of hierarchical file mappings of an Optimal Flexible Architecture-compliant installation.

```
/u02/u03 /u04
```

> **Note:**
>
> - The Grid homes are examples of Grid homes used for an Oracle Grid Infrastructure for a standalone server deployment (Oracle Restart), or a Grid home used for an Oracle Grid Infrastructure for a cluster deployment (Oracle Clusterware). You can have either an Oracle Restart deployment, or an Oracle Clusterware deployment. You cannot have both options deployed at the same time.
>
> - Oracle Automatic Storage Management (Oracle ASM) is included as part of an Oracle Grid Infrastructure installation. Oracle recommends that you use Oracle ASM to provide greater redundancy and throughput.

Table C-2   Optimal Flexible Architecture Hierarchical File Path Examples

Directory	Description
`/`	Root directory
`/u01/`	User data mount point 1
`/u01/app/`	Subtree for application software
`/u01/app/ oraInventory`	Central OraInventory directory, which maintains information about Oracle installations on a server. Members of the group designated as the OINSTALL group have permissions to write to the central inventory. All Oracle software installation owners must have the OINSTALL group as their primary group, and be able to write to this group.
`/u01/app/oracle/`	Oracle base directory for user `oracle`. There can be many Oracle Database installations on a server, and many Oracle Database software installation owners. Oracle software homes that an Oracle installation owner owns should be located in the Oracle base directory for the Oracle software installation owner, unless that Oracle software is Oracle Grid Infrastructure deployed for a cluster.
`/u01/app/grid`	Oracle base directory for user `grid`. The Oracle home (Grid home) for Oracle Grid Infrastructure for a cluster installation is located outside of the Grid user. There can be only one Grid home on a server, and only one Grid software installation owner. The Grid home contains log files and other administrative files.
`/u01/app/oracle/ admin/`	Subtree for database administration files

Table C-2   (Cont.) Optimal Flexible Architecture Hierarchical File Path Examples

Directory	Description
`/u01/app/oracle/admin/TAR`	Subtree for support log files
`/u01/app/oracle/admin/db_sales/`	Admin subtree for database named "sales"
`/u01/app/oracle/admin/db_dwh/`	Admin subtree for database named "dwh"
`/u01/app/oracle/fast_recovery_area/`	Subtree for recovery files
`/u01/app/oracle/fast_recovery_area/db_sales`	Recovery files for database named "sales"
`/u01/app/oracle/fast_recovery_area/db_dwh`	Recovery files for database named "dwh"

`/u02/app/oracle/oradata`
`/u03/app/oracle/oradata`
`/u04/app/oracle/oradata`
	Oracle data file directories

`/u01/app/oracle/product/`	Common path for Oracle software products other than Oracle Grid Infrastructure for a cluster
`/u01/app/oracle/product/21.0.0/dbhome_1`	Oracle home directory for Oracle Database 1, owned by Oracle Database installation owner account `oracle`

Table C-2   (Cont.) Optimal Flexible Architecture Hierarchical File Path Examples

Directory	Description
`/u01/app/oracle/product/21.0.0/dbhome_2`	Oracle home directory for Oracle Database 2, owned by Oracle Database installation owner account `oracle`
`/u01/app/oracle2/product/21.0.0/dbhome_2`	Oracle home directory for Oracle Database 2, owned by Oracle Database installation owner account `oracle2`
`/u01/app/oracle/product/21.0.0/grid`	Oracle home directory for Oracle Grid Infrastructure for a standalone server, owned by Oracle Database and Oracle Grid Infrastructure installation owner `oracle`.
`/u01/app/21.0.0/grid`	Oracle home directory for Oracle Grid Infrastructure for a cluster (Grid home), owned by user `grid` before installation, and owned by `root` after installation.
`/u01/app/oracle/product/21.0.0/client_1`	Oracle home directory for Oracle Database Client 1, owned by Oracle Database installation owner account `oracle`