

Oracle® Grid Infrastructure

Grid Infrastructure Installation and Upgrade Guide



19c for IBM AIX on POWER Systems (64-Bit)
E96274-11
March 2025

The Oracle logo, consisting of a solid red square with the word "ORACLE" in white, uppercase, sans-serif font centered within it.

ORACLE®

Copyright © 2013, 2025, Oracle and/or its affiliates.

Primary Author: Subhash Chandra

Contributing Authors: Aparna Kamath, Douglas Williams, Mark Bauer, Prakash Jashnani, Janet Stern

Contributors: Markus Michalewicz, Aneesh Khandelwal, Rajesh Dasari, Pallavi Kamath, Donald Graves, James Williams, Ian Cookson, Jonathan Creighton, Angad Gokakkar, Srinivas Poovala, Prasad Bagal, Balaji Pagadala, Neha Avasthy, Apparsamy Perumal, Akshay Shah, Saar Maoz, Barb Glover, Kevin Jernigan, Eric Belden, Mark Scardina, Mark Fuller, Barbara Glover, Saar Maoz, Binoy Sukumaran

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

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

If this is software, software documentation, data (as defined in the Federal Acquisition Regulation), or related documentation that is delivered to the U.S. Government or anyone licensing it on behalf of the U.S. Government, then the following notice is applicable:

U.S. GOVERNMENT END USERS: Oracle programs (including any operating system, integrated software, any programs embedded, installed, or activated on delivered hardware, and modifications of such programs) and Oracle computer documentation or other Oracle data delivered to or accessed by U.S. Government end users are "commercial computer software," "commercial computer software documentation," or "limited rights data" pursuant to the applicable Federal Acquisition Regulation and agency-specific supplemental regulations. As such, the use, reproduction, duplication, release, display, disclosure, modification, preparation of derivative works, and/or adaptation of i) Oracle programs (including any operating system, integrated software, any programs embedded, installed, or activated on delivered hardware, and modifications of such programs), ii) Oracle computer documentation and/or iii) other Oracle data, is subject to the rights and limitations specified in the license contained in the applicable contract. The terms governing the U.S. Government's use of Oracle cloud services are defined by the applicable contract for such services. No other rights are granted to the U.S. Government.

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

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

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

This software or hardware and documentation may provide access to or information about content, products, and services from third parties. Oracle Corporation and its affiliates are not responsible for and expressly disclaim all warranties of any kind with respect to third-party content, products, and services unless otherwise set forth in an applicable agreement between you and Oracle. Oracle Corporation and its affiliates will not be responsible for any loss, costs, or damages incurred due to your access to or use of third-party content, products, or services, except as set forth in an applicable agreement between you and Oracle.

Contents

Preface

Audience	xiii
Documentation Accessibility	xiii
Set Up Java Access Bridge to Implement Java Accessibility	xiv
Related Documentation	xiv
Conventions	xiv

1 Oracle Grid Infrastructure Installation Checklist

Server Hardware Checklist for Oracle Grid Infrastructure	1-1
Operating System Checklist for Oracle Grid Infrastructure on IBM AIX on POWER Systems (64-Bit)	1-2
Server Configuration Checklist for Oracle Grid Infrastructure	1-2
Network Checklist for Oracle Grid Infrastructure	1-3
User Environment Configuration Checklist for Oracle Grid Infrastructure	1-5
Storage Checklist for Oracle Grid Infrastructure	1-7
Cluster Deployment Checklist for Oracle Grid Infrastructure	1-8
Installer Planning Checklist for Oracle Grid Infrastructure	1-9

2 Checking and Configuring Server Hardware for Oracle Grid Infrastructure

Logging In to a Remote System Using X Window System	2-1
Checking Server Hardware and Memory Configuration	2-2

3 Configuring Operating Systems for Oracle Grid Infrastructure on IBM AIX on POWER Systems (64–Bit)

Guidelines for IBM AIX on POWER Systems (64-bit) Operating System Installation	3-2
Reviewing Operating System and Software Upgrade Best Practices	3-2
General Upgrade Best Practices	3-2
New Server Operating System Upgrade Option	3-3
Oracle ASM Upgrade Notifications	3-3
Reviewing Operating System Security Common Practices	3-4
About Installation Fixup Scripts	3-4

About Operating System Requirements	3-5
Operating System Requirements for IBM AIX on POWER Systems (64-bit)	3-5
Supported IBM AIX 7.1 Versions	3-6
Supported IBM AIX 7.2 Versions	3-6
Supported IBM AIX 7.3 Versions	3-7
Additional Drivers and Software Packages for IBM AIX on POWER Systems (64-bit)	3-8
Installing Oracle Messaging Gateway	3-9
Installation Requirements for ODBC and LDAP	3-9
About ODBC Drivers and Oracle Database	3-10
Installing ODBC Drivers for IBM AIX	3-10
About LDAP and Oracle Plug-ins	3-10
Installing the LDAP Package	3-10
Installation Requirements for Programming Environments for IBM AIX	3-10
Installation Requirements for Web Browsers	3-11
Checking the Software Requirements	3-12
Enabling I/O Completion Ports	3-12
Checking Resource Limits for AIX	3-13
Tuning AIX System Environment	3-14
Tuning Virtual Memory Manager (VMM)	3-14
Tuning Virtual Processor Manager (VPM)	3-15
Increasing System Block Size Allocation	3-15
Configuring SSH LoginGraceTime Parameter for AIX	3-15
Configuring User Process Parameters	3-16
Configuring Network Tuning Parameters	3-16
Recommended Values for Network Tuning Parameters	3-18
Using Automatic SSH Configuration During Installation	3-18
Setting Network Time Protocol for Cluster Time Synchronization	3-19
Running the rootpre.sh Script	3-20

4 Configuring Networks for Oracle Grid Infrastructure and Oracle RAC

About Oracle Grid Infrastructure Network Configuration Options	4-2
Understanding Network Addresses	4-2
About the Public IP Address	4-3
About the Private IP Address	4-3
About the Virtual IP Address	4-3
About the Grid Naming Service (GNS) Virtual IP Address	4-4
About the SCAN	4-4
About Shared SCAN	4-6
Network Interface Hardware Minimum Requirements	4-6
Private IP Interface Configuration Requirements	4-7
IPv4 and IPv6 Protocol Requirements	4-9

Oracle Grid Infrastructure IP Name and Address Requirements	4-9
About Oracle Grid Infrastructure Name Resolution Options	4-10
Cluster Name and SCAN Requirements	4-11
IP Name and Address Requirements For Grid Naming Service (GNS)	4-11
IP Name and Address Requirements For Multi-Cluster GNS	4-12
About Multi-Cluster GNS Networks	4-12
Configuring GNS Server Clusters	4-12
Configuring GNS Client Clusters	4-13
Creating and Using a GNS Client Data File	4-13
IP Name and Address Requirements for Manual Configuration of Cluster	4-14
Confirming the DNS Configuration for SCAN	4-15
Broadcast Requirements for Networks Used by Oracle Grid Infrastructure	4-16
Multicast Requirements for Networks Used by Oracle Grid Infrastructure	4-16
Domain Delegation to Grid Naming Service	4-16
Choosing a Subdomain Name for Use with Grid Naming Service	4-17
Configuring DNS for Cluster Domain Delegation to Grid Naming Service	4-17
Configuration Requirements for Oracle Flex Clusters	4-18
Understanding Oracle Flex Clusters	4-18
About Oracle Flex ASM Clusters Networks	4-19
General Requirements for Oracle Flex Cluster Configuration	4-20
Oracle Flex Cluster DHCP-Assigned Virtual IP (VIP) Addresses	4-20
Oracle Flex Cluster Manually-Assigned Addresses	4-21
Grid Naming Service Cluster Configuration Example	4-21
Manual IP Address Configuration Example	4-22
Network Interface Configuration Options	4-23

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

Creating Groups, Users and Paths for Oracle Grid Infrastructure	5-1
Determining If an Oracle Inventory and Oracle Inventory Group Exist	5-2
Creating the Oracle Inventory Group If an Oracle Inventory Does Not Exist	5-3
About Oracle Installation Owner Accounts	5-3
Restrictions for Oracle Software Installation Owners	5-3
Identifying an Oracle Software Owner User Account	5-5
About the Oracle Base Directory for the grid User	5-5
About the Oracle Home Directory for Oracle Grid Infrastructure Software	5-6
About Creating the Oracle Home and Oracle Base Directory	5-7
Oracle Installations with Standard and Job Role Separation Groups and Users	5-8
About Oracle Installations with Job Role Separation	5-8
Standard Oracle Database Groups for Database Administrators	5-9
Extended Oracle Database Groups for Job Role Separation	5-9

Creating an ASMSNMP User	5-10
Oracle Automatic Storage Management Groups for Job Role Separation	5-11
Creating Operating System Privileges Groups	5-11
Creating the OSASM Group	5-12
Creating the OSDBA for ASM Group	5-12
Creating the OSOPER for ASM Group	5-13
Creating the OSDBA Group for Database Installations	5-13
Creating an OSOPER Group for Database Installations	5-13
Creating the OSBACKUPDBA Group for Database Installations	5-14
Creating the OSDGDBA Group for Database Installations	5-14
Creating the OSKMDBA Group for Database Installations	5-14
Creating the OSRACDBA Group for Database Installations	5-15
Creating Operating System Oracle Installation User Accounts	5-15
Creating an Oracle Software Owner User	5-16
Modifying Oracle Owner User Groups	5-17
Identifying Existing User and Group IDs	5-17
Creating Identical Database Users and Groups on Other Cluster Nodes	5-18
Example of Creating Minimal Groups, Users, and Paths	5-19
Example of Creating Role-allocated Groups, Users, and Paths	5-20
Configuring Grid Infrastructure Software Owner User Environments	5-23
Environment Requirements for Oracle Software Owners	5-24
Procedure for Configuring Oracle Software Owner Environments	5-24
Setting Remote Display and X11 Forwarding Configuration	5-27
Preventing Installation Errors Caused by Terminal Output Commands	5-28
Enabling Intelligent Platform Management Interface (IPMI)	5-28
Requirements for Enabling IPMI	5-29
Configuring the IPMI Management Network	5-29
Configuring the BMC	5-30
Example of BMC Configuration Using IPMItool	5-30

6 Supported Storage Options for Oracle Database and Oracle Grid Infrastructure

Supported Storage Options for Oracle Grid Infrastructure	6-1
Oracle ACFS and Oracle ADVM	6-3
Oracle ACFS and Oracle ADVM Support on IBM AIX	6-3
Restrictions and Guidelines for Oracle ACFS	6-4
Storage Considerations for Oracle Grid Infrastructure and Oracle RAC	6-5
Guidelines for Using Oracle ASM Disk Groups for Storage	6-6
Guidelines for Configuring Oracle ASM Disk Groups on NFS	6-7
Using Logical Volume Managers with Oracle Grid Infrastructure and Oracle RAC	6-8
Using a Cluster File System for Oracle Clusterware Files	6-8

About NFS Storage for Data Files	6-9
About Direct NFS Client Mounts to NFS Storage Devices	6-9

7 Configuring Storage for Oracle Grid Infrastructure

Configuring Storage for Oracle Automatic Storage Management	7-2
Identifying Storage Requirements for Oracle Automatic Storage Management	7-2
Oracle Clusterware Storage Space Requirements	7-6
About the Grid Infrastructure Management Repository	7-9
Using an Existing Oracle ASM Disk Group	7-10
About Upgrading Existing Oracle Automatic Storage Management Instances	7-10
Selecting Disks to use with Oracle ASM Disk Groups	7-11
Specifying the Oracle ASM Disk Discovery String	7-11
Creating Files on a NAS Device for Use with Oracle Automatic Storage Management	7-12
Using Disk Groups with Oracle Database Files on Oracle ASM	7-13
Identifying and Using Existing Oracle Database Disk Groups on Oracle ASM	7-13
Configuring Disk Devices for Oracle ASM on IBM AIX	7-14
Creating Disk Groups for Oracle Database Data Files	7-15
Creating Directories for Oracle Database Files	7-15
Configuring File System Storage for Oracle Database	7-16
Configuring NFS Buffer Size Parameters for Oracle Database	7-17
Checking TCP Network Protocol Buffer for Direct NFS Client	7-17
Creating an orafstab File for Direct NFS Client	7-18
Enabling and Disabling Direct NFS Client Control of NFS	7-20
Enabling HCC on Direct NFS Client	7-21
Configuring HACMP Multinode Disk Heartbeat (MNDHB)	7-21
About Using HACMP with Oracle Clusterware	7-22
Deploying HACMP and MDNDHB for Oracle Clusterware	7-22
Upgrading an Existing Oracle Clusterware and HACMP Installation	7-25
Creating Member Cluster Manifest File for Oracle Member Clusters	7-26
Creating and Using Oracle ASM Credentials File	7-27
Configuring Oracle Automatic Storage Management Cluster File System	7-27

8 Installing Oracle Grid Infrastructure

About Image-Based Oracle Grid Infrastructure Installation	8-2
Understanding Cluster Configuration Options	8-2
About Oracle Standalone Clusters	8-3
About Oracle Cluster Domain and Oracle Domain Services Cluster	8-3
About Oracle Member Clusters	8-4
About Oracle Extended Clusters	8-5
Installing Oracle Grid Infrastructure for a New Cluster	8-6

About Oracle Grid Infrastructure Installation	8-6
Installing Oracle Standalone Cluster	8-7
Installing Oracle Domain Services Cluster	8-12
Installing Oracle Member Clusters	8-17
Installing Oracle Grid Infrastructure Using a Cluster Configuration File	8-22
Installing Only the Oracle Grid Infrastructure Software	8-23
Installing Software Binaries for Oracle Grid Infrastructure for a Cluster	8-24
Configuring Software Binaries for Oracle Grid Infrastructure for a Cluster	8-25
Configuring the Software Binaries Using a Response File	8-25
Setting Ping Targets for Network Checks	8-26
About Deploying Oracle Grid Infrastructure Using Oracle Fleet Patching and Provisioning	8-26
Confirming Oracle Clusterware Function	8-28
Confirming Oracle ASM Function for Oracle Clusterware Files	8-29
Understanding Offline Processes in Oracle Grid Infrastructure	8-29

9 Oracle Grid Infrastructure Postinstallation Tasks

Required Postinstallation Tasks	9-1
Downloading Release Update Patches	9-2
Recommended Postinstallation Tasks	9-2
Configuring IPMI-based Failure Isolation Using Crsctl	9-3
Creating a Backup of the root.sh Script	9-4
About Installing Oracle Autonomous Health Framework	9-4
Creating a Fast Recovery Area	9-4
About the Fast Recovery Area and the Fast Recovery Area Disk Group	9-5
Creating the Fast Recovery Area Disk Group	9-5
Checking the SCAN Configuration	9-6
Setting Resource Limits for Oracle Clusterware and Associated Databases and Applications	9-7
About Changes in Default SGA Permissions for Oracle Database	9-7
Using Earlier Oracle Database Releases with Oracle Grid Infrastructure	9-8
General Restrictions for Using Earlier Oracle Database Releases	9-8
Configuring Earlier Release Oracle Database on Oracle ACFS	9-9
Making Oracle ASM Available to Earlier Oracle Database Releases	9-10
Using ASMCA to Administer Disk Groups for Earlier Database Releases	9-10
Using the Correct LSNRCTL Commands	9-11
Modifying Oracle Clusterware Binaries After Installation	9-11

10 Removing Oracle Database Software

About Oracle Deinstallation Options	10-2
Oracle Deinstallation (Deinstall)	10-3

Unlocking and Deinstalling the Previous Release Grid Home	10-5
Deinstallation Examples for Oracle Database	10-6
Deinstallation Response File Example for Oracle Grid Infrastructure for a Cluster	10-7
Migrating Standalone Oracle Grid Infrastructure Servers to a Cluster	10-10
Relinking Oracle Grid Infrastructure for a Cluster Binaries	10-12
Changing the Oracle Grid Infrastructure Home Path	10-13
Unconfiguring Oracle Clusterware Without Removing Binaries	10-14
Unconfiguring Oracle Member Cluster	10-15

11 Upgrading Oracle Grid Infrastructure

Understanding Out-of-Place Upgrade	11-2
About Oracle Grid Infrastructure Upgrade and Downgrade	11-2
Options for Oracle Grid Infrastructure Upgrades	11-3
Restrictions for Oracle Grid Infrastructure Upgrades	11-4
Preparing to Upgrade an Existing Oracle Clusterware Installation	11-5
Upgrade Checklist for Oracle Grid Infrastructure	11-6
Checks to Complete Before Upgrading Oracle Grid Infrastructure	11-8
Moving Oracle Clusterware Files from NFS to Oracle ASM	11-9
Running the Oracle ORAchk Upgrade Readiness Assessment	11-10
Using CVU to Validate Readiness for Oracle Clusterware Upgrades	11-10
About the CVU Upgrade Validation Command Options	11-10
Example of Verifying System Upgrade Readiness for Grid Infrastructure	11-12
Using Dry-Run Upgrade Mode to Check System Upgrade Readiness	11-12
About Oracle Grid Infrastructure Dry-Run Upgrade Mode	11-12
Performing Dry-Run Upgrade Using Oracle Universal Installer	11-13
Understanding Rolling Upgrades Using Batches	11-14
Performing Rolling Upgrade of Oracle Grid Infrastructure	11-15
Upgrading Oracle Grid Infrastructure from an Earlier Release	11-15
Completing an Oracle Clusterware Upgrade when Nodes Become Unreachable	11-17
Joining Inaccessible Nodes After Forcing an Upgrade	11-18
Changing the First Node for Install and Upgrade	11-18
About Upgrading Oracle Grid Infrastructure Using Oracle Fleet Patching and Provisioning	11-18
Applying Patches to Oracle Grid Infrastructure	11-19
About Individual Oracle Grid Infrastructure Patches	11-20
About Oracle Grid Infrastructure Software Patch Levels	11-20
Patching Oracle Grid Infrastructure	11-20
Applying Patches During an Oracle Grid Infrastructure Installation or Upgrade	11-20
Applying Patches After an Oracle Grid Infrastructure Installation or Upgrade	11-21
Applying Patches when Oracle Clusterware Fails to Start	11-22
Patching and Switching Oracle Grid Infrastructure Homes	11-23
Updating Oracle Enterprise Manager Cloud Control Target Parameters	11-25

Updating the Enterprise Manager Cloud Control Target After Upgrades	11-25
Updating the Enterprise Manager Agent Base Directory After Upgrades	11-26
Registering Resources with Oracle Enterprise Manager After Upgrades	11-26
Unlocking and Deinstalling the Previous Release Grid Home	11-27
Checking Cluster Health Monitor Repository Size After Upgrading	11-28
Downgrading Oracle Clusterware to an Earlier Release	11-28
Options for Oracle Grid Infrastructure Downgrades	11-29
Restrictions for Oracle Grid Infrastructure Downgrades	11-30
Downgrading Oracle Clusterware to 18c	11-30
Downgrading Oracle Standalone Cluster to 18c	11-30
Downgrading Oracle Domain Services Cluster to 18c	11-32
Downgrading Oracle Member Cluster to 18c	11-35
Downgrading Oracle Grid Infrastructure to 18c when Upgrade Fails	11-37
Downgrading Oracle Clusterware to 12c Release 2 (12.2)	11-37
Downgrading Oracle Standalone Cluster to 12c Release 2 (12.2)	11-37
Downgrading Oracle Domain Services Cluster to 12c Release 2 (12.2)	11-39
Downgrading Oracle Member Cluster to 12c Release 2 (12.2)	11-42
Downgrading Oracle Grid Infrastructure to 12c Release 2 (12.2) when Upgrade Fails	11-44
Downgrading to Oracle Grid Infrastructure 12c Release 1 (12.1)	11-44
Downgrading to Oracle Grid Infrastructure 11g Release 2 (11.2)	11-46
Downgrading Oracle Grid Infrastructure Using Online Abort Upgrade	11-47
Completing Failed or Interrupted Installations and Upgrades	11-49
Completing Failed Installations and Upgrades	11-49
Continuing Incomplete Upgrade of First Node	11-50
Continuing Incomplete Upgrades on Remote Nodes	11-50
Continuing Incomplete Installation on First Node	11-51
Continuing Incomplete Installation on Remote Nodes	11-51
Converting to Oracle Extended Cluster After Upgrading Oracle Grid Infrastructure	11-52

A Completing Preinstallation Tasks Manually

Configuring SSH Manually on All Cluster Nodes	A-1
Checking Existing SSH Configuration on the System	A-2
Configuring SSH on Cluster Nodes	A-2
Create SSH Directory and Create SSH Keys On Each Node	A-2
Add All Keys to a Common authorized_keys File	A-3
Enabling SSH User Equivalency on Cluster Nodes	A-5
Configuring Shell Limits and System Configuration Parameters for IBM AIX	A-5
Configuring System Configuration Parameters for IBM AIX	A-6
Configuring Shell Limits on IBM AIX	A-7
Checking Asynchronous Input Output Processes	A-8

B Response Files

How Response Files Work	B-1
Reasons for Using Silent Mode or Response File Mode	B-2
Using Response Files	B-2
Preparing Response Files	B-3
Editing a Response File Template	B-3
Recording Response Files	B-4
Running Oracle Universal Installer Using a Response File	B-5
Running Configuration Assistants Using Response Files	B-7
Running Oracle DBCA Using Response Files	B-7
Running Net Configuration Assistant Using Response Files	B-8
Postinstallation Configuration Using Response File Created During Installation	B-9
Using the Installation Response File for Postinstallation Configuration	B-9
Running Postinstallation Configuration Using Response File	B-11
Postinstallation Configuration Using the ConfigToolAllCommands Script	B-12
About the Postinstallation Configuration File	B-12
Creating a Password Response File	B-13
Running Postinstallation Configuration Using a Password Response File	B-14

C Optimal Flexible Architecture

About the Optimal Flexible Architecture Standard	C-1
About Multiple Oracle Homes Support	C-2
About the Oracle Inventory Directory and Installation	C-2
Oracle Base Directory Naming Convention	C-4
Oracle Home Directory Naming Convention	C-4
Optimal Flexible Architecture File Path Examples	C-5

Index

List of Tables

1-1	Server Hardware Checklist for Oracle Grid Infrastructure	1-1
1-2	Operating System General Checklist for Oracle Grid Infrastructure on IBM AIX on POWER Systems (64-Bit)	1-2
1-3	Server Configuration Checklist for Oracle Grid Infrastructure	1-2
1-4	Network Configuration Tasks for Oracle Grid Infrastructure and Oracle RAC	1-3
1-5	User Environment Configuration for Oracle Grid Infrastructure	1-6
1-6	Oracle Grid Infrastructure Storage Configuration Checks	1-7
1-7	Oracle Grid Infrastructure Cluster Deployment Checklist	1-9
1-8	Oracle Universal Installer Checklist for Oracle Grid Infrastructure Installation	1-9
3-1	IBM AIX 7.1 on POWER Systems (64-Bit) Minimum Operating System Requirements	3-6
3-2	IBM AIX 7.2 on POWER Systems (64-Bit) Minimum Operating System Requirements	3-6
3-3	IBM AIX 7.3 on POWER Systems (64-bit) Minimum Operating System Requirements	3-8
3-4	Requirements for Programming Environments for IBM AIX on POWER Systems (64-bit)	3-11
3-5	Recommended Values for Virtual Memory Manager	3-14
3-6	Recommended Values for Network Tuning Parameters	3-18
4-1	Grid Naming Service Cluster Configuration Example	4-22
4-2	Manual Network Configuration Example	4-23
6-1	Supported Storage Options for Oracle Grid Infrastructure	6-2
7-1	Minimum Available Space Requirements for Oracle Standalone Cluster With GIMR Configuration	7-7
7-2	Minimum Available Space Requirements for Oracle Standalone Cluster Without GIMR Configuration	7-7
7-3	Minimum Available Space Requirements for Oracle Member Cluster with Local ASM	7-8
7-4	Minimum Available Space Requirements for Oracle Domain Services Cluster	7-8
8-1	Oracle ASM Disk Group Redundancy Levels for Oracle Extended Clusters with 2 Data Sites	8-6
11-1	Upgrade Checklist for Oracle Grid Infrastructure Installation	11-6
A-1	Minimum System Configuration Parameter Values	A-6
A-2	Recommended Shell Limit Values	A-7
B-1	Response Files for Oracle Database and Oracle Grid Infrastructure	B-3
C-1	Examples of OFA-Compliant Oracle Base Directory Names	C-4
C-2	Optimal Flexible Architecture Hierarchical File Path Examples	C-5

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](#)
- [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.
- [Related Documentation](#)
- [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 customer access to and use of Oracle support services will be pursuant to the terms and conditions specified in their Oracle order for the applicable services.

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*

Related Documentation

For more information, see the following Oracle resources:

Related Topics

- *Oracle Real Application Clusters Installation Guide for Linux and UNIX*
- *Oracle Database Installation Guide*
- *Oracle Clusterware Administration and Deployment Guide*
- *Oracle Real Application Clusters Administration and Deployment Guide*
- *Oracle Database Concepts*
- *Oracle Database New Features Guide*
- *Oracle Database Licensing Information*
- *Oracle Database Release Notes*
- *Oracle Database Examples Installation Guide*
- *Oracle Database Administrator's Reference for Linux and UNIX-Based Operating Systems*
- *Oracle Automatic Storage Management Administrator's Guide*
- *Oracle Database Upgrade Guide*
- *Oracle Database 2 Day DBA*
- *Oracle Application Express Installation Guide*

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.

Convention	Meaning
<i>italic</i>	Italic type indicates book titles, emphasis, or placeholder variables for which you supply particular values.
<code>monospace</code>	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 Grid Infrastructure on IBM AIX on POWER Systems \(64-Bit\)](#)
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.
- [Cluster Deployment Checklist for Oracle Grid Infrastructure](#)
Review the checklist for planning your cluster deployment 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
Runlevel	2
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 Grid Infrastructure on IBM AIX on POWER Systems (64-Bit)

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

Table 1-2 Operating System General Checklist for Oracle Grid Infrastructure on IBM AIX on POWER Systems (64-Bit)

Item	Task
Operating system general requirements	The following IBM AIX on POWER Systems (64-bit) versions are supported: AIX 7.1 Technology Level 5 Service Pack 1 (7100-05-01-1731) or later, 64-bit kernel AIX 7.2 Technology Level 2 Service Pack 1 (7200-02-01-1731) or later, 64-bit kernel AIX 7.3 Technology Level 0 Service Pack 1 (7300-00-01-2148) or later, 64-bit kernel Review the system requirements section for a list of operating system filesets and APARS.

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 5 GB of space in the temporary disk space (/tmp) directory
Swap space allocation relative to RAM	Between 4 GB and 16 GB: Equal to RAM More than 16 GB: 16 GB

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

Check	Task
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 <i>Oracle Grid Infrastructure Installation and Upgrade Guide</i> 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.

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).

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

Check	Task
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.
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).

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

Check	Task
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 <code>/etc/hosts</code> (for example, <code>node1.example.com</code>, address <code>192.0.2.10</code>). The public node name should be the primary host name of each node, which is the name displayed by the <code>hostname</code> 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, <code>node1-vip.example.com</code>, address <code>192.0.2.11</code>). 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 <code>hostname-vip</code>. For example: <code>myclstr2-vip</code>.

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 <code>/etc/oraInst.loc</code> 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 <i>owner</i> 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 <code>/app/owner</code>, where <i>owner</i> 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.
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: <code>\$ORACLE_HOME</code>; <code>\$ORA_NLS10</code>; <code>\$TNS_ADMIN</code>. If you have set <code>\$ORA_CRS_HOME</code> as an environment variable, then unset it before starting an installation or upgrade. Do not use <code>\$ORA_CRS_HOME</code> 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.

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

Check	Task
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 <code>sudo</code> is a UNIX and Linux utility that allows members of the <code>sudoers</code> list privileges to run individual commands as root. Provide the user name and password of an operating system user who is a member of <code>sudoers</code>, and is authorized to run <code>sudo</code> on each cluster member node. To enable <code>sudo</code>, have a system administrator with the appropriate privileges configure a user who is a member of the <code>sudoers</code> 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 13 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 11 GB for Oracle Database Enterprise Edition • Allocate additional storage space as per your cluster configuration, as described in <i>Oracle Clusterware Storage Space Requirements</i>.

Table 1-6 (Cont.) Oracle Grid Infrastructure Storage Configuration Checks

Check	Task
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. - For Oracle Domain Services Cluster deployment, Oracle ASM storage is shared across all nodes, and is available to Oracle Member Clusters. Oracle Member Cluster for Oracle Databases can either use storage services from the Oracle Domain Services Cluster or local Oracle ASM storage shared across all the nodes. Oracle Member Cluster for Applications always use storage services from the Oracle Domain Services Cluster. Before installing Oracle Member Cluster, create a Member Cluster Manifest file that specifies the storage details. 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, you can specify the same or separate Oracle ASM disk group for the GIMR. For Oracle Domain Services Cluster deployment, the GIMR must be configured on a separate Oracle ASM disk group. Oracle Member Clusters use the remote GIMR of the Oracle Domain Services Cluster. You must specify the GIMR details when you create the Member Cluster Manifest file before installation.

Cluster Deployment Checklist for Oracle Grid Infrastructure

Review the checklist for planning your cluster deployment Oracle Grid Infrastructure installation.

Table 1-7 Oracle Grid Infrastructure Cluster Deployment Checklist

Check	Task
Configure an Oracle Cluster that hosts all Oracle Grid Infrastructure services and Oracle ASM locally and accesses storage directly	Deploy an Oracle Standalone Cluster. Use the Oracle Extended Cluster option to extend an Oracle RAC cluster across two, or more, separate sites, each equipped with its own storage.
Configure an Oracle Cluster Domain to standardize, centralize, and optimize your Oracle Real Application Clusters (Oracle RAC) deployment	Deploy an Oracle Domain Services Cluster. To run Oracle Real Application Clusters (Oracle RAC) or Oracle RAC One Node database instances, deploy Oracle Member Cluster for Oracle Databases. To run highly-available software applications, deploy Oracle Member Cluster for Applications.

Installer Planning Checklist for Oracle Grid Infrastructure

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

Table 1-8 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: <i>Oracle Database Licensing Information User Manual</i>
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: <i>Oracle Clusterware Administration and Deployment Guide</i>

Table 1-8 (Cont.) Oracle Universal Installer Checklist for Oracle Grid Infrastructure Installation

Check	Task
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: <pre>./orachk -u -o pre</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 <code>cron</code> jobs do not run during installation	If the installer is running when daily <code>cron</code> jobs start, then you may encounter unexplained installation problems if your <code>cron</code> job is performing cleanup, and temporary files are deleted before the installation is finished. Oracle recommends that you complete installation before daily <code>cron</code> jobs are run, or disable daily <code>cron</code> 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/
Check running Oracle processes, and shut down processes if necessary	- On a node with a standalone 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 standalone 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 clients to have full access to the original server 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 physical RAM size on the server:

```
# /usr/sbin/lsattr -E -l sys0 -a realmem
```

2. To determine the available RAM and swap space:

```
# /usr/sbin/lsp -s
```

Oracle recommends that you take multiple values for the available RAM and swap space before finalizing a value. This is because the available RAM and swap space keep changing depending on the user interactions with the computer.

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

```
# /usr/sbin/lsp -a
```

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

4. Determine the amount of space available in the `/tmp` directory:

```
# df -m /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.
 - Extend the file system that contains the /tmp directory.
5. Determine the amount of free disk space on the system:

```
# df -g
# df -m
```

6. Determine if the system architecture can run the software:

```
# /usr/bin/getconf HARDWARE_BITMODE
```

The expected output of this command is 64. If you do not see the expected output, then you cannot install the software on this system.

To determine if the system is started in 64-bit mode:

```
# /usr/sbin/bootinfo -K
```

The result of this command must be 64, indicating that the 64-bit kernel is enabled.

Verify that the processor architecture matches the Oracle software release to install. If you do not see the expected output, then you cannot install the software on this system.

 Note:

For AIX Based Systems, Oracle Database supports 64-bit kernel and does not provide support for 32-bit kernel applications.

7. To determine the run level:

```
/usr/bin/who -r
```

3

Configuring Operating Systems for Oracle Grid Infrastructure on IBM AIX on POWER Systems (64–Bit)

Complete operating system configuration requirements and checks for IBM AIX on POWER Systems (64–Bit) operating systems before you start installation.

- [Guidelines for IBM AIX on POWER Systems \(64-bit\) Operating System Installation](#)
Operating system guidelines to be aware of 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 IBM AIX on POWER Systems \(64-bit\)](#)
The versions, filesets, and Authorized Problem Analysis Reports (APARs) listed in this section are supported for this release on IBM AIX on POWER Systems (64-bit).
- [Additional Drivers and Software Packages for IBM AIX on POWER Systems \(64-bit\)](#)
Information about optional drivers and software packages.
- [Checking the Software Requirements](#)
Check your software to see if they meet the minimum version, filesets, and Authorized Problem Analysis Reports (APARs) requirements for installation:
- [Enabling I/O Completion Ports](#)
To ensure successful database, client, and grid infrastructure installation, set IOCP to `Available` using this procedure.
- [Checking Resource Limits for AIX](#)
Set control and limit resources for users on the system to prevent denial of service attacks.
- [Tuning AIX System Environment](#)
Perform the following system tuning and configuration all cluster nodes.
- [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.

- [Running the rootpre.sh Script](#)
Run the `rootpre.sh` script only the first time you install Oracle Database on IBM AIX on POWER Systems (64-bit).

Guidelines for IBM AIX on POWER Systems (64-bit) Operating System Installation

Operating system guidelines to be aware of before proceeding with an Oracle installation.

For IBM AIX on POWER Systems (64-bit), Oracle Database supports 64-bit kernel and does not provide support for 32-bit kernel applications. Refer to your IBM AIX on POWER Systems (64-bit) documentation to obtain information about installing IBM AIX on POWER Systems (64-bit) 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. 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.

- You can upgrade Oracle Automatic Storage Management (Oracle ASM) 11g release 2 (11.2) and later 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 standalone database on a cluster that uses Oracle ASM, then you must shut down the standalone database before upgrading.
- The location of the Oracle ASM home changed in Oracle Grid Infrastructure 11g release 2 (11.2) so that 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 11g release 2 (11.2) or Oracle Grid Infrastructure 12c release 1 (12.1) 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 IBM AIX on POWER Systems (64-bit)

The versions, filesets, and Authorized Problem Analysis Reports (APARs) listed in this section are supported for this release on IBM AIX on POWER Systems (64-bit).

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 IBM AIX on POWER Systems (64-bit) version, and ensure that you have a supported kernel, filesets and APARs installed before starting installation.

- **Supported IBM AIX 7.1 Versions**
Use the following information to check the minimum supported IBM AIX 7.1 on POWER Systems (64-bit) versions:
- **Supported IBM AIX 7.2 Versions**
Use the following information to check the minimum supported IBM AIX 7.2 on POWER Systems (64-bit) versions:
- **Supported IBM AIX 7.3 Versions**
Use the following information to check the minimum supported IBM AIX 7.3 on POWER Systems (64-bit) versions:

Supported IBM AIX 7.1 Versions

Use the following information to check the minimum supported IBM AIX 7.1 on POWER Systems (64-bit) versions:

Table 3-1 IBM AIX 7.1 on POWER Systems (64-Bit) Minimum Operating System Requirements

Item	Minimum Requirements
SSH Requirement	OpenSSH_7.5p1 or later
AIX 7.1 Operating System	AIX 7.1 Technology Level 5 Service Pack 1 (7100-05-01-1731) or later, 64-bit kernel
AIX 7.1 Operating System Filesets	The following operating system filesets are required: • bos.adt.base • bos.adt.lib • bos.adt.libm • bos.perf.libperfstat • bos.perf.perfstat • bos.perf.proctools • security.pkcs11.7.1.3.15 or later • xlc.aix61.rte:13.1.0.1 or later • xlc.rte:13.1.0.1 or later • xlf.rte.aix61-15.1.0.9 or later The following filesets are required for Oracle Real Application Clusters (Oracle RAC) Installations only: • rsct.basic.rte • rsct.compat.clients.rte
AIX 7.1 APARs and Other Operating System Fixes	If you are using the minimum operating system TL level for AIX 7.1 listed above, then install the following AIX APAR fix: • AIX7.1 TL05 : IJ04311 – GETSOCKNAME returns invalid path name for AF_UNIX sockets AIX APAR numbers are tied to AIX versions and technology levels. Download and install the APAR that matches your AIX versions and Technology Levels from the IBM fix central website at the following URL: https://www.ibm.com/

Supported IBM AIX 7.2 Versions

Use the following information to check the minimum supported IBM AIX 7.2 on POWER Systems (64-bit) versions:

Table 3-2 IBM AIX 7.2 on POWER Systems (64-Bit) Minimum Operating System Requirements

Item	Minimum Requirements
SSH Requirement	OpenSSH_7.5p1 or later
AIX 7.2 Operating System	AIX 7.2 Technology Level 2 Service Pack 1 (7200-02-01-1731) or later, 64-bit kernel

Table 3-2 (Cont.) IBM AIX 7.2 on POWER Systems (64-Bit) Minimum Operating System Requirements

Item	Minimum Requirements
AIX 7.2 Operating System Filesets	The following operating system filesets are required: • bos.adt.base • bos.adt.lib • bos.adt.libm • bos.perf.libperfstat • bos.perf.perfstat • bos.perf.proctools • xlc.aix61.rte:13.1.0.1 or later • xlc.rte:13.1.0.1 or later • xlf.rte.aix61-15.1.0.9 or later The following filesets are required for Oracle Real Application Clusters (Oracle RAC) Installations only: • rsct.basic.rte • rsct.compat.clients.rte
AIX 7.2 APARs and Other Operating System Fixes	If you are using the minimum operating system TL level for AIX 7.2 listed above, then install the following AIX APAR fix: • AIX7.2 TL02 : IJ06143 – GETSOCKNAME returns invalid path name for AF_UNIX sockets



Note:

- If you are using a later TL level than the minimum level listed for this release, then contact IBM to determine if the required APARs listed in this section are included in the TL level that you have on your system. If they are included, then you do not have to install them. If they are not included, then you must install the equivalent APAR for the appropriate TL level.
- AIX APAR numbers are tied to AIX versions and technology levels. Download and install the APAR that matches your AIX versions and Technology Levels from the IBM fix central website at the following URL:
<https://www.ibm.com/>

Supported IBM AIX 7.3 Versions

Use the following information to check the minimum supported IBM AIX 7.3 on POWER Systems (64-bit) versions:

Table 3-3 IBM AIX 7.3 on POWER Systems (64-bit) Minimum Operating System Requirements

Item	Minimum Requirements
SSH Requirement	OpenSSH_7.5p1 or later
AIX 7.3 Operating System	AIX 7.3 Technology Level 0 Service Pack 1 (7300-00-01-2148) or later, 64-bit kernel
AIX 7.3 Operating System Filesets	The following operating system filesets are required: • bos.adt.base • bos.adt.lib • bos.adt.libm • bos.perf.libperfstat • bos.perf.perfstat • bos.perf.proctools • xlc.aix61.rte.13.1.2.0 or later • xlc.rte.13.1.2.0 or later • xlf rte The following filesets are required for Oracle Real Application Clusters (Oracle RAC) Installations only: • rsct.basic.rte • rsct.compat.clients.rte
AIX 7.3 APARs and Other Operating System Fixes	If you are using the minimum operating system TL level for AIX 7.3 listed above, then install the following AIX APAR fix: • APAR IJ38518

 Note:

- If you are using a later TL level than the minimum level listed for this release, then contact IBM to determine if the required APARs listed in this section are included in the TL level that you have on your system. If they are included, then you do not have to install them. If they are not included, then you must install the equivalent APAR for the appropriate TL level.
- AIX APAR numbers are tied to AIX versions and technology levels. Download and install the APAR that matches your AIX versions and Technology Levels from the IBM fix central website at the following URL:
<https://www.ibm.com/>

Additional Drivers and Software Packages for IBM AIX on POWER Systems (64-bit)

Information about optional drivers and software packages.

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 IBM AIX](#)
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.

For AIX 7.1, check My Oracle Support Certification for current Oracle Messaging Gateway requirements:

<https://support.oracle.com/>

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

<http://www.ibm.com>

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 IBM AIX](#)
If you intend to use ODBC, then install the most recent ODBC Driver Manager for IBM AIX on POWER Systems (64-bit).
- [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 IBM AIX

If you intend to use ODBC, then install the most recent ODBC Driver Manager for IBM AIX on POWER Systems (64-bit).

Download and install the ODBC Driver Manager from the following website:

<http://www.unixodbc.org>

Oracle ODBC driver on IBM AIX is certified with ODBC Driver Manager 2.3.4. To use ODBC with Oracle Database, you must also install the following package:

`gcc 4.2.1 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 IBM AIX

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

Table 3-4 Requirements for Programming Environments for IBM AIX on POWER Systems (64-bit)

Programming Environments	Support Requirements
Java Database Connectivity (JDBC) / Oracle Call Interface (OCI)	JDK 8 (64-bit 8.0.0.0 or later) with the JNDI extension with Oracle Java Database Connectivity and Oracle Call Interface drivers.
 Note: These are not mandatory requirements for database installation.	
Oracle C++	IBM XL C/C++ Runtime for AIX, v13.1.3.3 Fix Pack (June 2017)
Oracle C++ Call Interface	IBM XL C/C++ for AIX Fix Pack 11 (November 2018 PTF) for 13.1
Pro*C/C++	Download this software from the following URLs: http://www-01.ibm.com/support/docview.wss?uid=swg24043829
Oracle XML Developer's Kit (XDK)	
Pro*COBOL	IBM COBOL for AIX version 4.1.1, Fix Pack 12, (October 2014 PTF) Micro Focus Server Express 5.1 Micro Focus Visual COBOL for Eclipse 2.3 - Update 2 Micro Focus Visual COBOL v6.0
Pro*FORTRAN	IBM XL Fortran Runtime for AIX, Version 15.1, September 2014 PTF IBM XL Fortran for AIX, Version 15.1, September 2014 PTF
ADA	OC Systems PowerAda 5.5 For more information about OC Systems and PowerAda, go to: http://www.ocsystems.com/w/index.php/OCS:PowerAda
Network	Communications Server for AIX V6.4.0.2 PTF VIO Server 2.2.1.1 FP25 (28 Nov 2011)
Oracle Clusterware	HACMP 7.1 RSCT 3.1.0.3 or higher GPFS 4.1 cluster.license-7.1
Unzip utility	Unzip 6.0 or later. 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.

<https://support.oracle.com>

Related Topics

- *Oracle Enterprise Manager Cloud Control Basic Installation Guide*

Checking the Software Requirements

Check your software to see if they meet the minimum version, filesets, and Authorized Problem Analysis Reports (APARs) requirements for installation:

1. To determine the distribution and version of IBM AIX installed:

```
# oslevel -s
```

If the operating system version is lower than what is required, then upgrade your operating system accordingly to the currently supported or later version and level.

IBM AIX maintenance packages are available from the following website:

<https://www.ibm.com>

2. To determine if the required filesets are installed and committed enter a command similar to the following:

```
# lsllpp -l bos.adt.base bos.adt.lib bos.adt.libm bos.perf.perfstat \
bos.perf.libperfstat bos.perf.proctools
```

3. To determine the supported kernel mode:

```
# getconf KERNEL_BITMODE
```

- The expected output of this command is 64. If you do not see the expected output, then you cannot install the software on this system.
- Oracle Database supports 64-bit kernel and does not provide support for 32-bit kernel applications.

4. To determine if an APAR is installed:

```
# /usr/sbin/instfix -i -k "IV16716 IV20880 IV21128 IV28319"
```

If an APAR is not installed, then download it from the following website and install it:

<https://www.ibm.com>

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

<https://www.ibm.com>

Enabling I/O Completion Ports

To ensure successful database, client, and grid infrastructure installation, set IOCP to `Available` using this procedure.

On IBM AIX on POWER Systems (64-bit), enable I/O completion ports (IOCP) to ensure successful database and grid infrastructure installation.

To check if the IOCP module is enabled, run the `lsdev` command:

```
$ lsdev | grep iocp
```

The following sample output shows the IOCP status is set to `Defined` and hence not enabled:

```
iocp0          Defined          I/O Completion Ports
```

By default, IOCP is set to `Defined`. To enable IOCP, set IOCP to `Available` using the following procedure:

1. Log in as `root` and run the following command:

```
# smitty iocp
```

2. Select **Change / Show Characteristics of I/O Completion Ports**.
3. Change configured state at system restart from `Defined` to `Available`.
4. Run the `lsdev` command to confirm the IOCP status is set to `Available`:

```
$ lsdev | grep iocp
iocp0          Available          I/O Completion Ports
```

Restart the system to make the changes permanent.

Checking Resource Limits for AIX

Set control and limit resources for users on the system to prevent denial of service attacks.

On AIX platforms, the `/etc/security/login.cfg` file controls and limits resources for users on the system. On login, set control and limit resources for users on the system so that authentication is carried out through the `pam.conf` file. By default, resource limits are set to `std_auth` for AIX operating systems.

1. Edit the file `/etc/security/login.cfg` file to change `auth_type` under the `usw` stanza from `std_auth` to `pam_auth`.
2. To ensure that these resource limits are honored, confirm that the line login session required `/usr/lib/security/pam_aix` is set in `/etc/pam.conf`.

For example:

<code>dtsession</code>	<code>auth</code>	<code>required</code>	<code>/usr/lib/security/pam_aix</code>
<code>dtlogin</code>	<code>session</code>	<code>required</code>	<code>/usr/lib/security/pam_aix</code>
<code>ftp</code>	<code>session</code>	<code>required</code>	<code>/usr/lib/security/pam_aix</code>
<code>imap</code>	<code>session</code>	<code>required</code>	<code>/usr/lib/security/pam_aix</code>
<code>login</code>	<code>session</code>	<code>required</code>	<code>/usr/lib/security/pam_aix</code>
<code>rexec</code>	<code>session</code>	<code>required</code>	<code>/usr/lib/security/pam_aix</code>
<code>rlogin</code>	<code>session</code>	<code>required</code>	<code>/usr/lib/security/pam_aix</code>
<code>rsh</code>	<code>session</code>	<code>required</code>	<code>/usr/lib/security/pam_aix</code>
<code>snapp</code>	<code>session</code>	<code>required</code>	<code>/usr/lib/security/pam_aix</code>
<code>su</code>	<code>session</code>	<code>required</code>	<code>/usr/lib/security/pam_aix</code>
<code>swrole</code>	<code>session</code>	<code>required</code>	<code>/usr/lib/security/pam_aix</code>
<code>telnet</code>	<code>session</code>	<code>required</code>	<code>/usr/lib/security/pam_aix</code>
<code>xdm</code>	<code>session</code>	<code>required</code>	<code>/usr/lib/security/pam_aix</code>

```
OTHER    session required    /usr/lib/security/pam_prohibit
websm_rlogin session required /usr/lib/security/pam_aix
websm_su  session required   /usr/lib/security/pam_aix
wbem      session required   /usr/lib/security/pam_aix
```

Tuning AIX System Environment

Perform the following system tuning and configuration all cluster nodes.

- [Tuning Virtual Memory Manager \(VMM\)](#)
Oracle recommends that you use the `vmo` command to tune virtual memory.
- [Tuning Virtual Processor Manager \(VPM\)](#)
Oracle recommends that you use the `vpm_xvcpus` command to tune virtual processor.
- [Increasing System Block Size Allocation](#)
Perform the following tuning configuration on AIX systems.
- [Configuring SSH LoginGraceTime Parameter for AIX](#)
On AIX systems, configure the OpenSSH parameter to avoid timeout errors.
- [Configuring User Process Parameters](#)
Verify that the maximum number of processes allowed for each user is set to 16384 or greater.
- [Configuring Network Tuning Parameters](#)
View the current value specified for network tuning parameters, and change them if necessary.
- [Recommended Values for Network Tuning Parameters](#)
Verify that the network tuning parameters are set to the values shown or higher values.

Tuning Virtual Memory Manager (VMM)

Oracle recommends that you use the `vmo` command to tune virtual memory.

Note:

- The `lru_file_repage` parameter does not apply to IBM AIX 7.1 and later updates to IBM AIX 7.1.
- The tunable parameters `strict_maxclient` and `strict_maxperm` are deprecated in IBM AIX 7.2 and later updates to IBM AIX 7.2.

Table 3-5 Recommended Values for Virtual Memory Manager

Parameter	Value
<code>minperm%</code>	3
<code>maxperm%</code>	90
<code>maxclient% = 90</code>	90

For example:

```
vmo -p -o minperm%=3
vmo -p -o maxperm%=90
vmo -p -o maxclient%=90
```

You must restart the system for these changes to take effect.

Tuning Virtual Processor Manager (VPM)

Oracle recommends that you use the `vpm_xvcpus` command to tune virtual processor.

When you enable CPU folding with only one CUP, your system may reboot and face performance problems.

1. Check if CPU folding is enabled:

```
schedo -L| grep "vpm_xvcpus" |cut -d " " -f 17
```

The above command returns 0 if CPU folding is enabled.

2. Set value of the `vpm_xvcpus` parameter to at least 2:

```
schedo -o vpm_xvcpus=2
```

Increasing System Block Size Allocation

Perform the following tuning configuration on AIX systems.

- Increase the space allocated for ARG/ENV list to 128. The size is specified by number of 4K blocks.

For example:

```
/usr/sbin/chdev -l sys0 -a ncargs='1024'
```

Configuring SSH LoginGraceTime Parameter for AIX

On AIX systems, configure the OpenSSH parameter to avoid timeout errors.

On AIX systems, the OpenSSH parameter `LoginGraceTime` by default is commented out, and the default behavior of OpenSSH on AIX can sometimes result in timeout errors. To avoid these errors, complete the following procedure:

1. Log in as `root`.
2. Using a text editor, open the OpenSSH configuration file `/etc/ssh/sshd_config`.
3. Locate the comment line `#LoginGraceTime 2m`.
4. Uncomment the line, and change the value to 0 (unlimited).

For example:

```
LoginGraceTime 0
```

5. Save `/etc/ssh/sshd_config`.
6. Restart SSH.

Configuring User Process Parameters

Verify that the maximum number of processes allowed for each user is set to 16384 or greater.

1. Enter the following command:

```
# smit chgsys
```

2. Verify that the value shown for Maximum number of PROCESSES allowed for each user is greater than or equal to 16384. If necessary, edit the existing value.
3. When you have finished making changes, press **Enter**, then **Esc+0 (Exit)** to exit.

For production systems, this value should be at least 128 plus the sum of the `PROCESSES` and `PARALLEL_MAX_SERVERS` initialization parameters for each database running on the system.

Configuring Network Tuning Parameters

View the current value specified for network tuning parameters, and change them if necessary.

1. To check the current values of the network tuning parameters, enter commands similar to the following:

```
# no -a | more
```

2. If you must change the value of any parameter, then enter the following command to determine whether the system is running in compatibility mode:

```
# lsattr -E -l sys0 -a pre520tune
```

If the system is running in compatibility mode, then the output is similar to the following, showing that the value of the `pre520tune` attribute is enabled:

```
pre520tune enable Pre-520 tuning compatibility mode True
```

3. If the system is running in compatibility mode, then follow these steps to change the parameter values:

- a. Enter commands similar to the following to change the value of each parameter:

```
# no -o parameter_name=value
```

For example:

```
# no -o udp_recvspace=655360
```

- b. Add entries similar to the following to the `/etc/rc.net` file for each parameter that you changed in the previous step:

```
if [ -f /usr/sbin/no ] ; then  
    /usr/sbin/no -o udp_sendspace=65536
```

```

/usr/sbin/no -o udp_recvspace=655360
/usr/sbin/no -o tcp_sendspace=65536
/usr/sbin/no -o tcp_recvspace=65536
/usr/sbin/no -o rfc1323=1
/usr/sbin/no -o sb_max=4194304
/usr/sbin/no -o ipqmaxlen=512
fi

```

By adding these lines to the `/etc/rc.net` file, the values persist when the system restarts.

- c. You can also use the `chdev` command to change the characteristics of a device or interface. For example, set the RFC1323 value for the network interface `en5` without restarting the system as follows:

```
chdev -l en5 -a rfc1323=1
```

- 4. If the system is not running in compatibility mode, then enter commands similar to the following to change the parameter values:

`ipqmaxlen` parameter:

```
/usr/sbin/no -r -o ipqmaxlen=512
```

Other parameter:

```
/usr/sbin/no -p -o parameter=value
```

If you modify the `ipqmaxlen` parameter, then you must restart the system. These commands modify the `/etc/tunables/nextboot` file, causing the attribute values to persist when the system restarts.

- 5. If you need to change parameters, and you do not restart your system, then use the `ifconfig` command to check each network parameter after you change the `no` global setting.

For example:

```

# ifconfig en0
en0:
flags=1e080863,2c0<UP,BROADCAST,NOTRAILERS,RUNNING,SIMPLEX,MULTICAST,GROUPR
T,6
4BIT,CHECKSUM_OFFLOAD(ACTIVE),LARGESEND,CHAIN,MONITOR>
inet 192.0.2.1 netmask 0xfffff800 broadcast 192.0.2.0
    inet 192.0.2.2 netmask 0xfffff800 broadcast 192.0.2.0
    inet 192.0.2.3 netmask 0xfffff800 broadcast 192.0.2.0
    inet 192.0.2.4 netmask 0xfffff800 broadcast 192.0.2.0
    tcp_sendspace 131072 tcp_recvspace 65536 rfc1323 0

```

For the ISNO parameter `tcp_sendspace`, use the following command to set it:

```
# ifconfig en0 tcp_sendspace 65536
```

Related Topics

- [Recommended Values for Network Tuning Parameters](#)
Verify that the network tuning parameters are set to the values shown or higher values.

Recommended Values for Network Tuning Parameters

Verify that the network tuning parameters are set to the values shown or higher values.

Table 3-6 Recommended Values for Network Tuning Parameters

Network Tuning Parameter	Recommended Value
ipqmaxlen	512
rfc1323	1
sb_max	4194304
tcp_recvspace	65536
tcp_sendspace	65536
udp_recvspace	655360 Note: The recommended value of this parameter is 10 times the value of the <code>udp_sendspace</code> parameter. The value must be less than the value of the <code>sb_max</code> parameter.
udp_sendspace	65536 Note: This value is suitable for a default database installation. For production databases, the minimum value for this parameter is 4 KB plus the value of the database <code>DB_BLOCK_SIZE</code> initialization parameter multiplied by the value of the <code>DB_MULTIBLOCK_READ_COUNT</code> initialization parameter: $(DB_BLOCK_SIZE * DB_FILE_MULTIBLOCK_READ_COUNT) + 4\text{ KB}$

Related Topics

- [Configuring Network Tuning Parameters](#)
View the current value specified for network tuning parameters, and change them if necessary.

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 `sftp` 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 `chronyd` or `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`:

```
# stopsrc -s xntpd
```

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 stat resource ora.ctssd -t -init
```

Running the rootpre.sh Script

Run the `rootpre.sh` script only the first time you install Oracle Database on IBM AIX on POWER Systems (64-bit).



Note:

You do not have to run the `rootpre.sh` script on an IBM AIX on POWER Systems (64-bit) server that has an Oracle Database software already installed.

Use the following procedure to run the `rootpre.sh` script:

1. Download the Oracle Database installation image files (`db_home.zip`) and extract the files into a new Oracle home directory. For example:

```
$ mkdir -p /u01/app/oracle/product/19.0.0/dbhome_1
$ chgrp oinstall /u01/app/oracle/product/19.0.0/dbhome_1
$ cd /u01/app/oracle/product/19.0.0/dbhome_1
$ unzip -q /tmp/db_home.zip
```

2. Switch the user to `root`:

```
$ su - root
password:
#
```

3. Run the `rootpre.sh` script using the following command:

```
# $ORACLE_HOME/clone/rootpre.sh
```

4. Exit from the `root` account:

```
# exit
```


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`.

Clients configured to use IP addresses for Oracle Database releases prior to Oracle Database 11g release 2 can continue to use their existing connection addresses; using SCAN is not required. When you upgrade to Oracle Clusterware 12c release 1 (12.1) or later releases, the SCAN becomes available, and you should use the SCAN for connections to Oracle Database 11g release 2 or later databases. When an earlier release of Oracle Database is upgraded, it registers with the SCAN listeners, and clients can start using the SCAN to connect to that database. 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 Oracle Database 11g release 2 and later 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.

 Note:

Refer to the My Oracle Support Note 2646289.1 for a list of supported bonding modes.

- **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.

Oracle recommends that you configure Jumbo Frames with an MTU size of 9000 (MTU=9000) on your private network.

 Note:

When configuring Jumbo Frames, ensure that all the components of the private network including host adapters, network drivers, and switches support using Jumbo Frames.

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.

Related Topics

- [My Oracle Support Note 2646289.1](#)

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 and Oracle Member Cluster for Applications, 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

Starting with Oracle Grid Infrastructure 12c Release 1 (12.1), 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.
- [Configuring GNS Server Clusters](#)
Review these requirements to configure GNS server clusters.
- [Configuring GNS Client Clusters](#)
To configure a GNS client cluster, check to ensure all of the following requirements are completed.
- [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.

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.

Configuring 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 and Oracle Member Cluster for Applications, 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](#)
Starting with Oracle Grid Infrastructure 12c Release 1 (12.1), as part of an Oracle Flex Cluster installation, Oracle ASM is configured within Oracle Grid Infrastructure to provide storage services.
- [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

Starting with Oracle Grid Infrastructure 12c Release 1 (12.1), as part of an Oracle Flex Cluster installation, Oracle ASM is configured within Oracle Grid Infrastructure to provide storage services.

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, as all servers configured with the cluster node role are peers.
- 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 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.

About Oracle Flex ASM Cluster with Oracle IOServer (IOS) Configuration

An Oracle IOServer instance provides Oracle ASM file access for Oracle Database instances on nodes of Oracle Member Clusters that do not have connectivity to Oracle ASM managed disks. IOS enables you to configure Oracle Member Clusters on such nodes. On the storage cluster, the IOServer instance on each node opens up network ports to which clients send their I/O. The IOServer instance receives data packets from the client and performs the appropriate I/O to Oracle ASM disks similar to any other database client. On the client side, databases can use direct NFS (dNFS) to communicate with an IOServer instance. However, no client side configuration is required to use IOServer, so you are not required to provide a server IP

address or any additional configuration information. On nodes and clusters that are configured to access Oracle ASM files through IO Server, the discovery of the Oracle IOS instance occurs automatically.

To install an Oracle Member Cluster, the administrator of the Oracle Domain Services Cluster creates an Oracle Member Cluster using a `crsctl` command that creates a Member Cluster Manifest file. During Oracle Grid Infrastructure installation, if you choose to install an Oracle Member Cluster, then the installer prompts you for the Member Cluster Manifest file. An attribute in the Member Cluster Manifest file specifies if the Oracle Member Cluster is expected to access Oracle ASM files through an IO Server instance.

Related Topics

- *Oracle Automatic Storage Management Administrator's Guide*

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.



Note:

Do not use DHCP-assigned VIPs for Oracle Grid Infrastructure 19c. If you get VIP subnet configuration error during Oracle Real Application Cluster (Oracle RAC) installation, then kill the `oraagent` process and restart the cluster.

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
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
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 software owner for Oracle Grid Infrastructure, and directories for Oracle home.
- [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.

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 software owner for Oracle Grid Infrastructure, and directories for Oracle home.

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. If you have system administration privileges, then review the topics in this section and configure operating system groups and users as needed.

- [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](#)
Create an Oracle Inventory group manually as part of a planned installation, particularly where more than one Oracle software product is installed on servers.
- [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 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 /etc/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

Create an Oracle Inventory group manually as part of a planned installation, particularly where more than one Oracle software product is installed on servers.

By default, if an `oraInventory` group does not exist, then the installer uses the primary group of the installation owner for the Oracle software being installed as the `oraInventory` group. Ensure that this group is available as a primary group for all planned Oracle software installation owners. If the `oraInst.loc` file does not exist, then create the Oracle Inventory group using the following procedure:

1. Enter the following command:

```
# smit security
```

2. Choose the appropriate menu items to create the Oracle Inventory (`oinstall`) group.
3. Press **F10** to exit.

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/19.0.0/grid` is the Oracle Grid Infrastructure home:

```
$ export ORACLE_HOME=/u01/app/19.0.0/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(rac
dba)
```

```
$ 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 `/p/m/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/19.0.0/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/19.0.0`, `/u01/app/19.0.0/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:

```
/u01/19.0.0/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 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 home or Oracle base cannot be `symlinks`, nor can any of their parent directories, all the way to up to the `root` directory.

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 `SYSDG` 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 installation. Members of this group are granted the SYSRAC privileges to perform day to day administration of Oracle databases on a RAC cluster.

Creating the OSASM Group

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

Create an OSASM group using the group name `asmadmin` unless a group with that name already exists:

1. Enter the following command:

```
# smit security
```

2. Choose the appropriate menu items to create the `asmadmin` group.
3. Press **F10** to exit.

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:

1. Enter the following command:

```
# smit security
```

2. Choose the appropriate menu items to create the `asmadmin` group.
3. Press **F10** to exit.

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:

1. Enter the following command:

```
# smit security
```

2. Choose the appropriate menu items to create the `asmoper` group.
3. Press **F10** to exit.

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:

1. Enter the following command:

```
# smit security
```

2. Choose the appropriate menu items to create the `dba` group.
3. Press **F10** to exit.

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. To use an OSOPER group, then you must 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:

1. Enter the following command:

```
# smit security
```

2. Choose the appropriate menu items to create the `oper` group.
3. Press **F10** to exit.

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:

1. Enter the following command:

```
# smit security
```

2. Choose the appropriate menu items to create the `backupdba` group.
3. Press **F10** to exit.



Note:

On IBM AIX on POWER Systems (64-bit), the user and group name length limit parameter default value is 9 characters. Refer your operating system documentation, to increase the user and group name length limit.

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:

1. Enter the following command:

```
# smit security
```

2. Choose the appropriate menu items to create the `dgdba` group.
3. Press **F10** to exit.

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.

Create the OSKMDBA group using the groups name `kmdba` unless a group with that name already exists:

1. Enter the following command:

```
# smit security
```

2. Choose the appropriate menu items to create the `kmdba` group.
3. Press **F10** to exit.

Creating the OSRACDBA Group for Database Installations

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

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

1. Enter the following command:

```
# smit security
```

2. Choose the appropriate menu items to create the `racdba` group.
3. Press **F10** to exit.

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.

1. Enter the following command:

```
# smit security
```

2. Choose the appropriate menu items to create the `oracle` user, specifying the following information:

- In the **Primary GROUP** field, specify the Oracle Inventory group, for example `oinstall`.
- In the **Group SET** field, specify the OSDBA group and if required, the OSOPER group. For example `dba`, `asmdba`, or `oper`.

 Note:

The UID for the `oracle` user must be less than 65536.

3. Press **F10** to exit.
4. Set the password of the `oracle` user:

```
# passwd oracle
```

5. Ensure that the Oracle software owner user (`oracle` or `grid`) has the capabilities `CAP_NUMA_ATTACH`, `CAP_BYPASS_RAC_VMM`, and `CAP_PROPAGATE`.
 - a. To check existing capabilities, enter the following command as `root`; in this example, the Oracle Database installation user account is `oracle`:

```
# /usr/bin/luser -a capabilities oracle
```

- b. To add capabilities, enter a command similar to the following:

```
# /usr/bin/chuser  
capabilities=CAP_NUMA_ATTACH,CAP_BYPASS_RAC_VMM,CAP_PROPAGATE oracle
```

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

 Note:

On IBM AIX on POWER Systems (64-bit), Oracle Database software installations using an Active Directory (AD) account is not supported.

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:

1. Enter the following command:

```
# smit security
```

2. Choose the appropriate menu items to modify the Oracle installation owner user group.
3. In the **Primary GROUP** field, specify the Oracle Inventory group. For example: oinstall.
4. In the **Group SET** field, specify the required secondary groups. For example: dba, asmdba, backupdba, dgdba, kmdba, and racdba.
5. Press **F10** to exit.

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 Preinstallation RPM or prior installations, then the `oinstall` and `dba` groups.

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

```
# mkgroup -A id=54421 oinstall
# mkgroup -A id=54322 dba
# mkgroup -A id=54323 oper
# mkgroup -A id=54324 backupdba
# mkgroup -A id=54325 dgdba
# mkgroup -A id=54326 kmdba
# mkgroup -A id=54327 asmdba
# mkgroup -A id=54328 asmoper
# mkgroup -A id=54329 asmadmin
# mkgroup -A id=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:

```
# mkuser id=54322 pgrp=oinstall groups=asmadmin,asmdba,racdba grid
```

- The `id` option specifies the user ID, which must be the user ID that you identified in the previous subsection.
- The `pgrp` 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 `groups` 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.

4. Set the password of the user.

For example:

```
# passwd grid
```

5. Ensure that the Oracle software owner user has the capabilities `CAP_NUMA_ATTACH`, `CAP_BYPASS_RAC_VMM`, and `CAP_PROPAGATE`.

- a. To check existing capabilities, enter the following command as `root`; in this example, the Oracle Database installation user account is `grid`:

```
# /usr/bin/lsuser -a capabilities grid
```

- b. To add capabilities, enter a command similar to the following:

```
# /usr/bin/chuser
capabilities=CAP_NUMA_ATTACH,CAP_BYPASS_RAC_VMM,CAP_PROPAGATE grid
```

6. 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:

```
# mkgroup -A id='54421' adms='root' oinstall
# mkgroup -A id='54422' adms='root' dba
# mkuser id='54321' pgrp='oinstall' groups='dba' home='/home/oracle' oracle
# mkuser id='54322' pgrp='oinstall' groups='dba' home='/home/grid' grid
# 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/
```

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/19.0.0/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.

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:

```
# mkgroup -'A' id='54321' adms='root' oinstall
# mkgroup -'A' id='54322' adms='root' dba1
# mkgroup -'A' id='54332' adms='root' dba2
# mkgroup -'A' id='54323' adms='root' oper1
# mkgroup -'A' id='54333' adms='root' oper2
# mkgroup -'A' id='54324' adms='root' backupdba1
# mkgroup -'A' id='54334' adms='root' backupdba2
# mkgroup -'A' id='54325' adms='root' dgdba1
# mkgroup -'A' id='54335' adms='root' dgdba2
# mkgroup -'A' id='54326' adms='root' kmdba1
# mkgroup -'A' id='54336' adms='root' kmdba2
# mkgroup -'A' id='54327' adms='root' asmdba
# mkgroup -'A' id='54328' adms='root' asmoper
# mkgroup -'A' id='54329' adms='root' asmadmin
# mkgroup -'A' id='54330' adms='root' racdba1
# mkgroup -'A' id='54340' adms='root' racdba2
# mkuser id='53322' pgrp='oinstall' groups='asmadmin,asmdba,racdba1,racdba2'
home='/home/grid' grid
# mkuser id='53321' pgrp='oinstall'
groups='dba1,backupdba1,dgdba1,kmdba1,asmdba,racdba1,asmoper' home='/home/
oracle' oracle1
# mkuser id='53323' pgrp='oinstall'
groups='dba2,backupdba2,dgdba2,kmdba2,asmdba,racdba2' home='/home/oracle'
oracle1
# mkdir -p /u01/app/19.0.0/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/19.0.0/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 `oracle1: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](#)
You must make the following changes to configure Oracle software owner environments:
- [Procedure for Configuring Oracle Software Owner Environments](#)
Configure each Oracle installation owner user account environment.
- [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

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.
- Unset `EXTSHM`.

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. 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.

12. If the `/tmp` directory has less than 5 GB of free space, then identify a file system with at least 5 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
```

13. 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.

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](#)
Configure BMC on each node for remote control using LAN for IPMI-based node fencing to function properly.

- [Example of BMC Configuration Using IPMITool](#)
The following is an example of configuring BMC using `ipmitool` (version 1.8.6).

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

Configure BMC on each node for remote control using LAN for IPMI-based node fencing to function properly.

You can configure BMC from the BIOS prompt, using a distribution-specific management utility, or you can configure BMC using publicly available utilities, such as the following: IPMITool, which is available for Linux:

<http://ipmitool.sourceforge.net>

IPMIutil, which is available for Linux:

<http://ipmiutil.sourceforge.net>

Refer to the documentation for the configuration tool you select for details about using the tool to configure the BMC.

When you configure the BMC on each node, you must complete the following:

1. Enable IPMI over LAN, so that the BMC can be controlled over the management network.
2. Enable dynamic IP addressing using DHCP or GNS, or configure a static IP address for the BMC.
3. Establish an administrator user account and password for the BMC.
4. Configure the BMC for VLAN tags, if you will use the BMC on a tagged VLAN.

The configuration tool you use does not matter, but these conditions must be met for the BMC to function properly.

Example of BMC Configuration Using IPMITool

The following is an example of configuring BMC using `ipmitool` (version 1.8.6).

1. Log in as `root`.
2. Verify that `ipmitool` can communicate with the BMC using the IPMI driver by using the command `bmc info`, and looking for a device ID in the output. For example:

```
# ipmitool bmc info
Device ID           : 32
.
.
.
```

If `ipmitool` is not communicating with the BMC, then review the section *Configuring the BMC* and ensure that the IPMI driver is running.

3. Enable IPMI over LAN using the following procedure

- a. Determine the channel number for the channel used for IPMI over LAN. Beginning with channel 1, run the following command until you find the channel that displays LAN attributes (for example, the IP address):

```
# ipmitool lan print 1

. . .
IP Address Source      : 0x01
IP Address              : 140.87.155.89
. . .
```

- b. Turn on LAN access for the channel found. For example, where the channel is 1:

```
# ipmitool -I bmc lan set 1 access on
```

4. Configure IP address settings for IPMI using the static IP addressing procedure:

- **Using static IP Addressing**

If the BMC shares a network connection with ILOM, then the IP address must be on the same subnet. You must set not only the IP address, but also the proper values for netmask, and the default gateway. For example, assuming the channel is 1:

```
# ipmitool -I bmc lan set 1 ipaddr 192.168.0.55
# ipmitool -I bmc lan set 1 netmask 255.255.255.0
# ipmitool -I bmc lan set 1 defgw ipaddr 192.168.0.1
```

Note that the specified address (192.168.0.55) is associated only with the BMC, and does not respond to normal pings.

5. Establish an administration account with a username and password, using the following procedure (assuming the channel is 1):

- a. Set BMC to require password authentication for ADMIN access over LAN. For example:

```
# ipmitool -I bmc lan set 1 auth ADMIN MD5,PASSWORD
```

- b. List the account slots on the BMC, and identify an unused slot less than the maximum ID and not listed, for example, ID 4 in the following example. Note that some slots may be reserved and not available for reuse on some hardware.

```
# ipmitool user summary 1
Maximum IDs      : 20
Enabled User Count : 3
Fixed Name Count  : 2
# ipmitool user list 1
```

ID	Name	Enabled	Callin	Link Auth	IPMI Msg	Channel Priv Lim
1		true	false	false	true	USER
2	root	true	false	false	true	ADMINISTRATOR
3	sysoper	true	true	false	true	OPERATOR
12	default	true	true	false	true	NO ACCESS
13		true	false	true	false	CALLBACK

In the example above, there are 20 possible slots, and the first unused slot is number 4.

- c. Assign the desired administrator user name and password and enable messaging for the identified slot. (Note that for IPMI v1.5 the user name and password can be at most 16 characters). Also, set the privilege level for that slot when accessed over LAN (channel 1) to ADMIN (level 4). For example, where *username* is the administrative user name, and *password* is the password:

```
# ipmitool user set name 4 username
# ipmitool user set password 4 password
# ipmitool user enable 4
# ipmitool channel setaccess 1 4 privilege=4
# ipmitool channel setaccess 1 4 link=on
# ipmitool channel setaccess 1 4 ipmi=on
```

- d. Verify the setup using the command `lan print 1`. The output should appear similar to the following. Note that the items in **bold text** are the settings made in the preceding configuration steps, and comments or alternative options are indicated within brackets []:

```
# ipmitool lan print 1
Set in Progress      : Set Complete
Auth Type Support    : NONE MD2 MD5 PASSWORD
Auth Type Enable     : Callback : MD2 MD5
                    : User      : MD2 MD5
                    : Operator  : MD2 MD5
                    : Admin    : MD5 PASSWORD
                    : OEM       : MD2 MD5

IP Address Source   : DHCP Address [or Static Address]
IP Address         : 192.168.0.55
Subnet Mask        : 255.255.255.0
MAC Address          : 00:14:22:23:fa:f9
SNMP Community String : public
IP Header             : TTL=0x40 Flags=0x40 Precedence=...
Default Gateway IP : 192.168.0.1
Default Gateway MAC   : 00:00:00:00:00:00
.
.
.
# ipmitool channel getaccess 1 4
Maximum User IDs      : 10
Enabled User IDs      : 2

User ID               : 4
User Name          : username [This is the administration user]
Fixed Name            : No
Access Available      : call-in / callback
Link Authentication : enabled
IPMI Messaging      : enabled
Privilege Level     : ADMINISTRATOR
```

6. Verify that the BMC is accessible and controllable from a remote node in your cluster using the `bmc info` command. For example, if `node2-ipmi` is the network host name assigned the

IP address of `node2`'s BMC, then to verify the BMC on node `node2` from `node1`, with the administrator account `username`, enter the following command on `node1`:

```
$ ipmitool -H node2-ipmi -U username lan print 1
```

You are prompted for a password. Provide the IPMI password.

If the BMC is correctly configured, then you should see information about the BMC on the remote node. If you see an error message, such as `Error: Unable to establish LAN session`, then you must check the BMC configuration on the remote node.

7. Repeat this process for each cluster member node. If the IPMI administrator account credentials on each cluster member node are not identical, then IPMI will fail during configuration.

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:
- [Oracle ACFS and Oracle ADVM](#)
Oracle Automatic Storage Management Cluster File System (Oracle ACFS) extends Oracle ASM technology to support of all of your application data in both single instance and cluster configurations.
- [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:



Note:

The ACFS storage option on IBM AIX is deprecated in Oracle Database 19c.

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
Oracle Automatic Storage Management Cluster File System (Oracle ACFS)	No	No	Yes	Yes for Oracle Database 12c Release 1 (12.1) and later	Yes for Oracle Database 12c Release 1 (12.1) and later
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.

Related Topics

- [Oracle Database Upgrade Guide](#)

Oracle ACFS and Oracle ADVM

Oracle Automatic Storage Management Cluster File System (Oracle ACFS) extends Oracle ASM technology to support of all of your application data in both single instance and cluster configurations.

 Note:

The ACFS storage option on IBM AIX is deprecated in Oracle Database 19c.

To align the support for the IBM AIX platform, Oracle Automatic Storage Management Cluster File System (ACFS) on IBM AIX (called Oracle Advanced Cluster File System starting with Oracle Database 21c) is deprecated in Oracle Database 19c. For data or files currently stored on an ACFS file system that cannot be stored in the Oracle database, the IBM General Parallel File System (GPFS) can be used.

Oracle Automatic Storage Management Dynamic Volume Manager (Oracle ADVM) provides volume management services and a standard disk device driver interface to clients. Oracle ACFS communicates with Oracle ASM through the Oracle ADVM interface.

- [Oracle ACFS and Oracle ADVM Support on IBM AIX](#)
Oracle ACFS and Oracle ADVM are supported on IBM AIX on POWER Systems (64-Bit).
- [Restrictions and Guidelines for Oracle ACFS](#)
Review these topics as part of your storage plan for using Oracle ACFS for single instance and cluster configurations.

Related Topics

- [Oracle Automatic Storage Management Administrator's Guide](#)

Oracle ACFS and Oracle ADVM Support on IBM AIX

Oracle ACFS and Oracle ADVM are supported on IBM AIX on POWER Systems (64-Bit).

 Note:

The ACFS storage option on IBM AIX is deprecated in Oracle Database 19c.

Oracle ACFS and Oracle ADVM are supported on:

- IBM AIX 7.1 and later updates to IBM AIX 7.1
- IBM AIX 7.2 and later updates to IBM AIX 7.2

Oracle ACFS and Oracle ADVM are not supported on IBM AIX Workload Partitions (WPARs).

 See Also:

- My Oracle Support Note 1369107.1 for more information about platforms and releases that support Oracle ACFS and Oracle ADVM:
<https://support.oracle.com/CSP/main/article?cmd=show&type=NOT&id=1369107.1>
- Patch Set Updates for Oracle Products (My Oracle Support Note 854428.1) for current release and support information:
<https://support.oracle.com/CSP/main/article?cmd=show&type=NOT&id=854428.1>

Restrictions and Guidelines for Oracle ACFS

Review these topics as part of your storage plan for using Oracle ACFS for single instance and cluster configurations.

 Note:

The ACFS storage option on IBM AIX is deprecated in Oracle Database 19c.

- Oracle Automatic Storage Management Cluster File System (Oracle ACFS) provides a general purpose file system.
- You can only use Oracle ACFS when Oracle ASM is configured.
- Note the following general guidelines and restrictions for placing Oracle Database and Oracle Grid Infrastructure files on Oracle ACFS:
 - You can place Oracle Database binaries, data files, and administrative files (for example, trace files) on Oracle ACFS.
 - Oracle ACFS does not support encryption or replication with Oracle Database data files, tablespace files, control files, redo logs, archive logs, RMAN backups, Data Pump dumpsets, and flashback files.
 - You can place Oracle Database homes on Oracle ACFS only if the database release is Oracle Database 11g Release 2, or more recent releases. You cannot install earlier releases of Oracle Database on Oracle ACFS.
 - For installations on Oracle Clusterware, you cannot place Oracle Clusterware files on Oracle ACFS.
- Oracle Restart does not support root-based Oracle Clusterware resources. For this reason, the following restrictions apply if you run Oracle ACFS on an Oracle Restart Configuration:
 - Starting with Oracle Database 18c, configuration assistants do not allow the creation of Oracle Database homes on Oracle ACFS in an Oracle Restart configuration.
 - Oracle Restart does not support Oracle ACFS resources on all platforms.
 - Starting with Oracle Database 12c, Oracle Restart configurations do not support the Oracle ACFS registry.
 - On Linux, Oracle ACFS provides an automated mechanism to load and unload drivers and mount and unmount Oracle ACFS file systems on system restart and shutdown.

However, Oracle ACFS does not provide automated recovery of mounted file systems when the system is running. Other than Linux, Oracle ACFS does not provide this automated mechanism on other operating systems.

- Creating Oracle data files on an Oracle ACFS file system is not supported in Oracle Restart configurations. Creating Oracle data files on an Oracle ACFS file system is supported on Oracle Grid Infrastructure for a cluster configurations.
- Oracle ACFS and Oracle ADVM are not supported on IBM AIX Workload Partitions (WPARs).

Related Topics

- *Oracle Automatic Storage Management Administrator's 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 on 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.
- 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 memory maps.
- 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 19c 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.

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.
- For Oracle Domain Services Clusters, you must configure two separate Oracle ASM disk groups, one for OCR and voting files and the other for the GIMR.

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.

With Oracle Database 11g Release 2 (11.2) and later releases, Oracle Database Configuration Assistant (DBCA) does not have the functionality to create disk groups for Oracle ASM.

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.



Note:

The ACFS storage option on IBM AIX is deprecated in Oracle Database 19c.

You can create Oracle ASM disk groups on NFS (Network File System) or block devices 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.
- Oracle ASM Filter Driver and Oracle ACFS and Oracle ADVM are not supported with NFS. These features are incompatible because of the nature of the operating system interface for NFS-based storage.



Note:

Oracle ACFS does not support NFS disks in an Oracle ASM disk group, but you can use NFS disks with Oracle ACFS using Oracle ACFS NAS Maximum Availability eXtensions (Oracle ACFS NAS MAX).

- 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.
- When you choose Oracle ASM on NFS, you cannot use Oracle Automatic Storage Management Cluster File System (Oracle ACFS) for storage. This cluster configuration cannot be used as a Oracle Fleet Patching and Provisioning Server.

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.

Related Topics

- [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.

Using Logical Volume Managers with Oracle Grid Infrastructure and Oracle RAC

Oracle Grid Infrastructure and Oracle RAC only support cluster-aware volume managers.

Using Logical 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. If you are

installing on IBM POWER, and you want to use a cluster file system, then you must use the IBM General Parallel File System (GPFS).

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 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 through 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_HOME/dbs/oranfstab`
2. `/etc/oranfstab`
3. `/etc/mtab`

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.

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

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.
- [Configuring HACMP Multinode Disk Heartbeat \(MNDHB\)](#)
Review these topics to configure Multi-node Disk Heartbeat (MNDHB) for Oracle Clusterware.
- [Creating Member Cluster Manifest File for Oracle Member Clusters](#)
Create a Member Cluster Manifest file to specify the Oracle Member Cluster configuration for the Grid Infrastructure Management Repository (GIMR), Grid Naming Service, Oracle ASM storage server, and Oracle Fleet Patching and Provisioning configuration.
- [Creating and Using Oracle ASM Credentials File](#)
Review this information to create an Oracle ASM credentials file.
- [Configuring Oracle Automatic Storage Management Cluster File System](#)
Review this information to configure Oracle ACFS for an Oracle RAC Oracle Database home.

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 various Oracle 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. Plan your Oracle ASM disk groups requirement, based on the cluster configuration you want to deploy. Oracle Domain Services Clusters store Oracle Clusterware files and the Grid Infrastructure Management Repository (GIMR) on separate Oracle ASM disk groups and hence require configuration of two separate Oracle ASM disk groups, one for OCR and voting files and the other for the GIMR.

2. 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 can use either Oracle ASM or a shared file system to store OCR and voting files on Oracle Standalone Cluster deployments, but you must use Oracle ASM to store OCR and voting files on Oracle Domain Services cluster deployments.
- 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.

3. 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).

- Extended redundancy

Extended redundancy disk group has similar features as the flex redundancy disk group. Extended redundancy is available when you configure an Oracle Extended Cluster. Extended redundancy extends Oracle ASM data protection to cover failure of sites by placing enough copies of data in different failure groups of each site. A site is a collection of failure groups. For extended redundancy with three sites, for example, two data sites, and one quorum failure group, the minimum number of disks is seven (three disks each for two data sites and one quorum failure group outside the two data sites). The maximum number of supported sites for extended redundancy is three. In an extended redundancy disk group, each site maintains the user data redundancy as specified by the file group attribute. Each site can host data failure groups and quorum failure groups for a given disk group. For example, if the file group redundancy is specified as 2 or 3, each site has 2 or 3 mirrors respectively, provided there are enough failure groups to accommodate the mirrors. See *About Oracle Extended Clusters* for more information about selecting redundancy levels for extended clusters.

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.

4. 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.

5. 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.

6. 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 and ACFS proxy instances

nodes = Number of nodes in cluster.

clients - Number of database instances for each node.

disks - Number of disks in disk group.

7. 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.

8. 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.
9. 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 various Oracle Cluster deployments.

Total Oracle Clusterware Available Storage Space Required by Oracle Cluster Deployment Type

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, a new GIMR is created only if the source Grid home has a GIMR configured.

Based on the cluster configuration you want to install, the Oracle Clusterware space requirements vary for different redundancy levels. The following tables list the space requirements for each cluster configuration and 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	Oracle Fleet Patching and Provisioning	Total Storage
External	1 GB	28 GB Each node beyond four: 5 GB	1 GB	30 GB
Normal	2 GB	56 GB Each node beyond four: 5 GB	2 GB	60 GB
High/Flex/Extended	3 GB	84 GB Each node beyond four: 5 GB	3 GB	90 GB

- Oracle recommends that you use separate disk groups for Oracle Clusterware files and 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.
- By default, all new Oracle Standalone Cluster deployments are configured with Oracle Fleet Patching and Provisioning for patching that cluster only. This deployment requires a minimal ACFS file system that is automatically configured in the same disk group as the GIMR.

Table 7-2 Minimum Available Space Requirements for Oracle Standalone Cluster Without GIMR Configuration

Redundancy Level	DATA Disk Group	Oracle Fleet Patching and Provisioning	Total Storage
External	1 GB	1 GB	2 GB

Table 7-2 (Cont.) Minimum Available Space Requirements for Oracle Standalone Cluster Without GIMR Configuration

Redundancy Level	DATA Disk Group	Oracle Fleet Patching and Provisioning	Total Storage
Normal	2 GB	2 GB	4 GB
High/Flex/Extended	3 GB	3 GB	6 GB

- Oracle recommends that you use separate disk groups for Oracle Clusterware files and 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.
- By default, all new Oracle Standalone Cluster deployments are configured with Oracle Fleet Patching and Provisioning for patching that cluster only. This deployment requires a minimal ACFS file system that is automatically configured.

Table 7-3 Minimum Available Space Requirements for Oracle Member Cluster with Local ASM

Redundancy Level	DATA Disk Group	Oracle Clusterware Backup Files	Total Storage
External	1 GB	4 GB	5 GB
Normal	2 GB	8 GB	10 GB
High/Flex/Extended	3 GB	12 GB	15 GB

- For Oracle Member Cluster, the storage space for the GIMR is pre-allocated in the centralized GIMR on the Oracle Domain Services Cluster as described in *Table 8–5*.
- Oracle recommends that you use separate disk groups for Oracle Clusterware files and for Oracle Clusterware backup files.

Table 7-4 Minimum Available Space Requirements for Oracle Domain Services Cluster

Redundancy Level	DATA Disk Group	MGMT Disk Group	Trace File Analyzer	Total Storage	Additional Oracle Member Cluster
External	1 GB and 1 GB for each Oracle Member Cluster	140 GB	200 GB	345 GB (excluding Oracle Fleet Patching and Provisioning)	Oracle Fleet Patching and Provisioning: 100 GB GIMR for each Oracle Member Cluster beyond four: 28 GB

Table 7-4 (Cont.) Minimum Available Space Requirements for Oracle Domain Services Cluster

Redundancy Level	DATA Disk Group	MGMT Disk Group	Trace File Analyzer	Total Storage	Additional Oracle Member Cluster
Normal	2 GB and 2 GB for each Oracle Member Cluster	280 GB	400 GB	690 GB (excluding Oracle Fleet Patching and Provisioning)	Oracle Fleet Patching and Provisioning: 200 GB GIMR for each Oracle Member Cluster beyond four: 56 GB
High/Flex/Extended	3 GB and 3 GB for each Oracle Member Cluster	420 GB	600 GB	1035 GB (excluding Oracle Fleet Patching and Provisioning)	Oracle Fleet Patching and Provisioning: 300 GB GIMR for each Oracle Member Cluster beyond four: 84 GB

- By default, the initial space allocation for the GIMR is for the Oracle Domain Services Cluster and four or fewer Oracle Member Clusters. You must add additional storage space for each Oracle Member Cluster beyond four.
- At the time of installation, TFA storage space requirements are evaluated to ensure the growth to the maximum sizing is possible. Only the minimum space is allocated for the ACFS file system, which extends automatically up to the maximum value, as required.
- Oracle recommends that you pre-allocate storage space for largest foreseeable configuration of the Oracle Domain Services Cluster according to the guidelines in the above table.

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 Management Database (MGMTDB) 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 Member Cluster for Database 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, and additional PDB for each member cluster that is added.

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 11g Release 2 (11.2.0.4) 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, Oracle ASM volumes, and Oracle Automatic Storage Management Cluster File System (Oracle ACFS).

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 11g Release 2 (11.2) 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 IBM AIX](#)
Complete these tasks to configure disk devices for use with Oracle Automatic Storage Management (Oracle ASM).
- [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 IBM AIX

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. To ensure that the disks are available, enter the following command:

```
# /usr/sbin/lsdev -Cc disk
```

The output from this command is similar to the following:

```
hdisk0 Available 1A-09-00-8,0 16 Bit LVD SCSI Disk Drive
hdisk1 Available 1A-09-00-9,0 16 Bit LVD SCSI Disk Drive
hdisk2 Available 17-08-L      SSA Logical Disk Drive
```

3. If a disk is not listed as available, then configure the new disks:

```
# /usr/sbin/cfgmgr
```

4. Enter the following command on any node to identify the device names for the physical disks that you want to use:

```
# /usr/sbin/lspv | grep -i none
```

This command displays information similar to the following for each disk that is not configured in a volume group:

```
hdisk2      0000078752249812  None
```

In this example, `hdisk2` is the device name of the disk and `0000078752249812` is the physical volume ID (PVID). The disks that you want to use may have a PVID, but they must not belong to a volume group.

5. Enter commands similar to the following to clear the PVID from each disk device that you want to use:

```
# /usr/sbin/chdev -l hdiskn -a pv=clear
```

6. Enter commands similar to the following to change the owner, group, and permissions on the character file for each disk that you want to add to the disk group:

```
# chown -R grid:asmadmin /dev/rhdiskn  
# chmod 660 /dev/rhdiskn
```

 Note:

If you are using a multi-pathing disk driver with Automatic Storage Management, then ensure that you set the permissions only on the correct logical device name for the disk.

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 Oracle 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 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 `orantstab` File for Direct NFS Client](#)
Direct NFS uses a configuration file, `orantstab`, 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 HCC on Direct NFS Client](#)
Perform the 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 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_mount:
dev = "/NFS_mount"
vfs = nfs
nodename = NFS_server
mount = true
options = bg,soft,intr,rsize=32768,wsize=32768
account = false
```

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:

```
# /usr/sbin/no -o sb_max
```

Oracle recommends that you set the value based on the link speed of your servers. For example:

```
# /usr/sbin/no -p -o sb_max=524288
```

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 `oranstab`. A single `oranstab` 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 HCC on Direct NFS Client

Perform the steps to enable Hybrid Columnar Compression (HCC) on Direct NFS Client.

1. Ensure that SNMP is enabled on the ZFS storage server. For example:

```
# snmpinfo -m get -c public -v -h server_name 1.3.6.1.4.1.42.2.225.1.4.2.0
enterprises.42.2.225.1.4.2.0 =
53:75:6e:20:53:74:6f:72:61:67:65:20:37:34:31:30
```

The above sequence of bytes 53:75:6e: . . . :31:30 is an ascii character representation for 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. If `libnetsnmp.so` is not installed, download and install it from the following URL:

<https://oss.oracle.com/netsnmp/>

**Note:**

On IBM AIX on POWER Systems (64-bit), HCC is not supported for block based devices.

Configuring HACMP Multinode Disk Heartbeat (MNDHB)

Review these topics to configure Multi-node Disk Heartbeat (MNDHB) for Oracle Clusterware.

- [About Using HACMP with Oracle Clusterware](#)
You must define one Multi-node Disk Heartbeat (MNDHB) network for each Oracle Clusterware voting file.
- [Deploying HACMP and MDNDHB for Oracle Clusterware](#)
Complete the following procedure to deploy HACMP and MDNDHP for Oracle Clusterware.

- [Upgrading an Existing Oracle Clusterware and HACMP Installation](#)
Complete the following procedure to upgrade an existing HACMP and Oracle Clusterware installation.

About Using HACMP with Oracle Clusterware

You must define one Multi-node Disk Heartbeat (MNDHB) network for each Oracle Clusterware voting file.

Each MNDHB and voting disk pair must be located on a single hard disk, separate from the other pairs. You must also configure MNDHB so that the node is halted if access is lost to a quorum of the MNDHB networks in the enhanced concurrent volume group.

To reduce the likelihood of a cluster partition, IBM recommends that HACMP is deployed with multiple IP networks and at least one non-IP network. The non-IP networks can be implemented using RS232 or disk heart-beating. For systems using Oracle RAC and HACMP enhanced concurrent resources (enhanced concurrent logical volumes) for database storage, you must configure MNDHB networks.

Install, configure and have HACMP running before installing Oracle Clusterware. For an Oracle RAC configuration, do not use HACMP for IP failovers on the Oracle RAC network interfaces (public, VIP or private). These network interfaces should not be configured to use HACMP IP failover, as Oracle Clusterware manages VIP failovers for Oracle RAC. The Oracle RAC network interfaces are bound to individual nodes and Oracle RAC instances. Problems can occur with Oracle Clusterware if HACMP reconfigures IP addresses over different interfaces, or fails over addresses across nodes. You only can use HACMP for failover of IP address on Oracle RAC nodes if Oracle RAC does not use these addresses.

Deploying HACMP and MDNDHB for Oracle Clusterware

Complete the following procedure to deploy HACMP and MDNDHP for Oracle Clusterware.

Ensure that your versions of HACMP and AIX meet the system requirements as listed in this guide.

1. Start HACMP.
2. Enter the following command to ensure that the HACMP clcomdES daemon is running:

```
# lssrc -s clcomdES
```

If the daemon is not running, then start it using the following command:

```
# startsrc -s clcomdES
```

3. Create HACMP cluster and add the Oracle Clusterware nodes. For example:

```
# smitty cm_add_change_show_an_hacmp_cluster.dialog
* Cluster Name [mycluster]
```

4. Create an HACMP cluster node for each Oracle Clusterware node. For example:

```
# smitty cm_add_a_node_to_the_hacmp_cluster_dialog
* Node Name [mycluster_node1]
Communication Path to Node []
```

5. Create HACMP Ethernet heartbeat networks. The HACMP configuration requires network definitions. Select **NO** for the IP address takeover for these networks, since they are used by Oracle Clusterware.

Create at least two network definitions: one for the Oracle public interface and a second one for the Oracle private (cluster interconnect) network. Additional Ethernet heartbeat networks can be added if desired. For example:

```
# smitty cm_add_a_network_to_the_hacmp_cluster_select
- select ether network
* Network Name [my_network_name]
* Network Type ether
* Netmask [my.network.netmask.here]
* Enable IP Address Takeover via IP Aliases [No]
IP Address Offset for Heart beating over IP Aliases []
```

6. For each of the networks added in the previous step, define all of the IP names for each Oracle Clusterware node associated with that network, including the public, private and VIP names for each Oracle Clusterware node. For example:

```
# smitty cm_add_communication_interfaces_devices.select
- select: Add Pre-defined Communication Interfaces and Devices /
Communication Interfaces / desired network
* IP Label/Address [node_ip_address]
* Network Type ether
* Network Name some_network_name
* Node Name [my_node_name]
Network Interface []
```

7. Create an HACMP resource group for the enhanced concurrent volume group resource with the following options:

```
# smitty config_resource_group.dialog.custom
* Resource Group Name [my_resource_group_name]
* Participating Nodes (Default Node Priority) [mynode1,mynode2,mynode3]
Startup Policy Online On All Available Nodes
Failover Policy Bring Offline (On Error Node Only)
Fallback Policy Never Fallback
```

8. Create an AIX enhanced concurrent volume group (Big VG, or Scalable VG) using either the command `smitty mkvg`, or using command lines. The VG must contain at least one hard disk for each voting disk. You must configure at least three voting disks.

In the following example, where you see default, accept the default response:

```
# smitty _mksvg
VOLUME GROUP name [my_vg_name] PP SIZE in MB
* PHYSICAL VOLUME names [mydisk1,mydisk2,mydisk3]
Force the creation of a volume group? no
Activate volume group AUTOMATICALLY no at system restart?
Volume Group MAJOR NUMBER []
Create VG Concurrent Capable? enhanced concurrent
Max PPs per VG in kilobytes default
Max Logical Volumes default
```

9. Under "Change/Show Resources for a Resource Group (standard)", add the concurrent volume group to the resource group added in the preceding steps.

For example:

```
# smitty cm_change_show_resources_std_resource_group_menu_dmn.select
- select_resource_group_from_step_6
Resource Group Name shared_storage
Participating Nodes (Default Node Priority) mynode1,mynode2,mynode3
Startup Policy Online On All Available Nodes
Fallover Policy Bring Offline (On Error Node Only)
Fallback Policy Never Fallback
Concurrent Volume Groups [enter_VG_from_step_7]
Use forced varyon of volume groups, if necessary false
Application Servers []
```

10. Using the following command, ensure that one MNDHB network is defined for each Oracle Clusterware voting disk. Each MNDHB and voting disk pair must be collocated on a single hard disk, separate from the other pairs. The MNDHB network and Voting Disks exist on shared logical volumes in an enhanced concurrent logical volume managed by HACMP as an enhanced concurrent resource. For each of the hard disks in the VG created in step 8 on which you want to place a voting disk logical volume (LV), create a MNDHB LV.

```
# smitty cl_add_mndhb_lv
- select_resource_group_defined_in_step_6
* Physical Volume name enter F4, then select a hard disk
Logical Volume Name []
Logical Volume Label []
Volume Group name ccvg
Resource Group Name shared_storage
Network Name [n]
```

When you define the LVs for the Oracle Clusterware voting disks, they should be defined on the same disks: one for each disk, as used in this step for the MNDHB LVs.

11. Configure MNDHB so that the node is halted if access is lost to a quorum of the MNDHB networks in the enhanced concurrent volume group. For example:

```
# smitty cl_set_mndhb_response
- select_the_VG_created_in_step_7
On loss of access Halt the node
Optional notification method []
Volume Group ccvg
```

12. Verify and Synchronize HACMP configuration. For example:

```
# smitty cm_initialization_and_standard_config_menu_dmn -
select Verify and Synchronize HACMP Configuration
```

Enter **Yes** if prompted: "Would you like to import shared VG: ccvg, in resource group my_resource_group onto node: mynode to node: racha702 [Yes / No]:"

13. Add the HACMP cluster node IP names to the file /usr/es/sbin/cluster/etc/rhosts.

Upgrading an Existing Oracle Clusterware and HACMP Installation

Complete the following procedure to upgrade an existing HACMP and Oracle Clusterware installation.

1. Back up all databases, and back up the Oracle Cluster Registry (OCR).
2. Shut down on all nodes all Oracle RAC databases, all node applications, and Oracle Clusterware.
3. Enter the following command to disable Oracle Clusterware from starting when nodes are restarted:

```
# crsctl disable crs
```

4. Shut down HACMP on all nodes.
5. Install HACMP APAR IZ01809, following the directions in the README included with that APAR.
6. Determine if the existing voting disk LVs are already on separate hard disks, and if each of these disks have sufficient space (at least 256 MB for the MNDHB LVs). If this is true, then create a MNDHB LV on each of the hard disks. If this is not true, then create new MNDHB LVs and new voting disk LVs, located on separate hard disks using the following command, responding to the sections in *italics* with the appropriate information for your system:

```
# smitty cl_add_mndhb_lv
- Select_resource_group
* Physical Volume name Enter F4, then select disk for the MNDHB and Voting
Disk pair
Logical Volume Name []
Logical Volume Label []
Volume Group name ccvg
Resource Group Name shared_storage
Network Name [net_diskhbmulti_01]
```

7. Verify and Synchronize HACMP configuration.
8. Start HACMP on all nodes.
9. If you added new LVs for voting disks in step 5, then replace each of the existing voting disks with the new ones.
10. Enter the following command to re-enable Oracle Clusterware:

```
# crsctl enable crs
```

11. Start Oracle Clusterware on all nodes, and verify that all resources start correctly.

Creating Member Cluster Manifest File for Oracle Member Clusters

Create a Member Cluster Manifest file to specify the Oracle Member Cluster configuration for the Grid Infrastructure Management Repository (GIMR), Grid Naming Service, Oracle ASM storage server, and Oracle Fleet Patching and Provisioning configuration.

Oracle Member Clusters use Oracle ASM storage from the Oracle Domain Services Cluster. Grid Naming Service (GNS) without zone delegation must be configured so that the GNS virtual IP address (VIP) is available for connection.

1. (Optional) If the Oracle Member Cluster accesses direct or indirect Oracle ASM storage, then, enable access to the disk group. Connect to any Oracle ASM instance as SYSASM user and run the command:

```
ALTER DISKGROUP diskgroup_name SET ATTRIBUTE 'access_control.enabled' = 'true';
```

2. From the Grid home on the Oracle Domain Services Cluster, create the member cluster manifest file:

```
cd Grid_home/bin
./crsctl create member_cluster_configuration member_cluster_name
-file cluster_manifest_file_name -member_type database|application
[-version member_cluster_version
[-domain_services [asm_storage local|direct|indirect][rhpf] [acfs]]
```

member_cluster_name is the name of the member cluster.

-file specifies the full path of the xml file to export the credentials, *-version* is the five digit Member Cluster version, for example, 19.0.0.0.0, if it is different from the Storage Server version. The Storage Server version is used if *-version* is not specified.

In the options for *-domain_services*, specifying *rhpf* generates credentials and configuration for a Oracle Fleet Patching and Provisioning Member Cluster, and *asm_storage* generates credentials and configuration for an Oracle ASM Member Cluster. *direct* if specified, signifies direct storage access, otherwise indirect.

This command creates a member cluster manifest file containing configuration details about Grid Infrastructure Management Repository (GIMR), Storage services, and Oracle Fleet Patching and Provisioning for the Oracle Member Cluster.

3. GNS client data is required if the Oracle Member Cluster uses dynamic networks and the server cluster has GNS with zone delegation. Provide the GNS client data as follows:

- As root or grid user, export the Grid Naming Service (GNS) client data, to the member cluster manifest file created earlier:

```
srvctl export gns -clientdata manifest_file_name -role CLIENT
```

The GNS configuration is appended to the member cluster manifest file.

4. Copy the manifest file to a location on the Oracle Member Cluster, and select the file during the installation and configuration of the Oracle Member Cluster.

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.

Configuring Oracle Automatic Storage Management Cluster File System

Review this information to configure Oracle ACFS for an Oracle RAC Oracle Database home.

**Note:**

The ACFS storage option on IBM AIX is deprecated in Oracle Database 19c. To align the support for the IBM AIX platform, Oracle Automatic Storage Management Cluster File System (ACFS) on IBM AIX (called Oracle Advanced Cluster File System starting with Oracle Database 21c) is deprecated in Oracle Database 19c. For data or files currently stored on an ACFS file system that cannot be stored in the Oracle database, the IBM General Parallel File System (GPFS) can be used.

Oracle ACFS is installed as part of an Oracle Grid Infrastructure 19c installation.

You can also create a General Purpose File System configuration of ACFS using ASMCA.

To configure Oracle ACFS for an Oracle Database home for an Oracle RAC database:

1. Install Oracle Grid Infrastructure for a cluster.
2. Change directory to the Oracle Grid Infrastructure home. For example:

```
$ cd /u01/app/19.0.0/grid
```

3. Ensure that the Oracle Grid Infrastructure installation owner has read and write permissions on the storage mountpoint you want to use. For example, if you want to use the mountpoint `/u02/acfsmounts/`:

```
$ ls -l /u02/acfsmounts
```

4. Start Oracle ASM Configuration Assistant as the grid installation owner. For example:

```
./asmca
```

5. The Configure ASM: ASM Disk Groups page shows you the Oracle ASM disk group you created during installation. Click the **ASM Cluster File Systems** tab.
6. On the ASM Cluster File Systems page, right-click the Data disk, then select **Create ACFS for Database Use**.
7. In the Create ACFS for Database window, enter the following information:
 - **Volume Name:** Enter the name of the database home. The name must be unique in your enterprise. For example: `dbase_01`
 - **Mount Point:** Enter the directory path for the mount point. For example: `/u02/acfsmounts/dbase_01`
Make a note of this mount point for future reference.
 - **Size (GB):** Enter in gigabytes the size you want the database home to be. The default is 12 GB and the minimum recommended size.
 - **Owner Name:** Enter the name of the Oracle Database installation owner you plan to use to install the database. For example: `oracle1`
 - **Owner Group:** Enter the OSDBA group whose members you plan to provide when you install the database. Members of this group are given operating system authentication for the SYSDBA privileges on the database. For example: `dba1`

Select **Automatically run configuration commands** to run ASMCA configuration commands automatically. To use this option, you must provide the root credentials on the ASMCA Settings page.

Click **OK** when you have completed your entries.

8. If you did not select to run configuration commands automatically, then run the script generated by Oracle ASM Configuration Assistant as a privileged user (`root`). On an Oracle Clusterware environment, the script registers the ACFS as a resource managed by Oracle Clusterware. Registering ACFS as a resource helps Oracle Clusterware to mount ACFS automatically in proper order when ACFS is used for an Oracle RAC Oracle Database home.
9. During Oracle RAC installation, ensure that you or the DBA who installs Oracle RAC selects for the Oracle home the mount point you provided in the **Mount Point** field (in the preceding example, `/u02/acfsmounts/dbase_01`).

Related Topics

- *Oracle Automatic Storage Management Administrator's Guide*

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. You can also use Oracle Fleet Patching and Provisioning for subsequent Oracle Grid Infrastructure and Oracle Database deployments.

- [About Image-Based Oracle Grid Infrastructure Installation](#)
Installation and configuration of Oracle Grid Infrastructure software is simplified with image-based installation.
- [Understanding Cluster Configuration Options](#)
Review these topics to understand the cluster configuration options available in Oracle Grid Infrastructure 19c.
- [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 providing cluster configuration information manually, or 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.
- [About Deploying Oracle Grid Infrastructure Using Oracle Fleet Patching and Provisioning](#)
Oracle Fleet Patching and Provisioning (Oracle FPP) is a software lifecycle management method for provisioning and maintaining Oracle homes. Oracle Fleet Patching and Provisioning enables mass deployment and maintenance of standard operating environments for databases, clusters, and user-defined software types.
- [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.

Understanding Cluster Configuration Options

Review these topics to understand the cluster configuration options available in Oracle Grid Infrastructure 19c.

- [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 Cluster Domain and Oracle Domain Services Cluster](#)
An Oracle Cluster Domain is a choice of deployment architecture for new clusters, introduced in Oracle Clusterware 12c Release 2.
- [About Oracle Member Clusters](#)
Oracle Member Clusters use centralized services from the Oracle Domain Services Cluster and can host databases or applications. Oracle Member Clusters can be of two types - Oracle Member Clusters for Oracle Databases or Oracle Member Clusters for applications.
- [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, and includes metadata required for Oracle Fleet Patching and Provisioning.

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 Cluster Domain and Oracle Domain Services Cluster

An Oracle Cluster Domain is a choice of deployment architecture for new clusters, introduced in Oracle Clusterware 12c Release 2.


Oracle Cluster Domain enables you to standardize, centralize, and optimize your Oracle Real Application Clusters (Oracle RAC) deployment for the private database cloud. Multiple cluster configurations are grouped under an Oracle Cluster Domain for management purposes and make use of shared services available within that Oracle Cluster Domain. The cluster configurations within that Oracle Cluster Domain include Oracle Domain Services Cluster and Oracle Member Clusters.

The Oracle Domain Services Cluster provides centralized services to other clusters within the Oracle Cluster Domain. These services include:

- A centralized Grid Infrastructure Management Repository (housing the MGMTDB for each of the clusters within the Oracle Cluster Domain)
- Trace File Analyzer (TFA) services, for targeted diagnostic data collection for Oracle Clusterware and Oracle Database
- Consolidated Oracle ASM storage management service, including the use of Oracle ACFS.
- An optional Oracle Fleet Patching and Provisioning Service to install clusters, and provision, patch, and upgrade Oracle Grid Infrastructure and Oracle Database homes. When you configure the Oracle Domain Services Cluster, you can also choose to configure the Oracle Fleet Patching and Provisioning Server.

An Oracle Domain Services Cluster provides these centralized services to Oracle Member Clusters. Oracle Member Clusters use these services for centralized management and to reduce their local resource usage.

Figure 8-1 Oracle Cluster Domain



About Oracle Member Clusters

Oracle Member Clusters use centralized services from the Oracle Domain Services Cluster and can host databases or applications. Oracle Member Clusters can be of two types - Oracle Member Clusters for Oracle Databases or Oracle Member Clusters for applications.

Oracle Member Clusters do not need direct connectivity to shared disks. Using the shared Oracle ASM service, they can leverage network connectivity to the IO Service or the ACFS Remote Service to access a centrally managed pool of storage. To use shared Oracle ASM services from the Oracle Domain Services Cluster, the member cluster needs connectivity to the Oracle ASM networks of the Oracle Domain Services Cluster.

Oracle Member Clusters cannot provide services to other clusters. For example, you cannot configure and use a member cluster as a GNS server or Oracle Fleet Patching and Provisioning Server.

Oracle Member Cluster for Oracle Databases

An Oracle Member Cluster for Oracle Databases supports Oracle Real Application Clusters (Oracle RAC) or Oracle RAC One Node database instances. This cluster registers with the management repository service and uses the centralized TFA service. It can use additional services as needed. An Oracle Member Cluster for Oracle Databases can be configured with

local Oracle ASM storage management or make use of the consolidated Oracle ASM storage management service offered by the Oracle Domain Services Cluster.

An Oracle Member Cluster for Oracle Database always uses remote Grid Infrastructure Management Repository (GIMR) from its Oracle Domain Services Cluster. For two-node or four-node clusters, hosting the GIMR on a remote cluster reduces the overhead of running an extra infrastructure repository on a cluster.

Oracle Member Cluster for Applications

Oracle Member Cluster for Applications hosts applications other than Oracle Database, as part of an Oracle Cluster Domain. The Oracle Member Cluster requires connectivity to Oracle Cluster Domain Services for centralized management and resource efficiency. The Oracle Member Cluster uses remote Oracle ASM storage, with any required shared storage provided through the Oracle ACFS Remote service. This cluster configuration enables high availability of any software application.

Unlike other cluster configurations that require public and private network interconnects, the Oracle Member Cluster for Application can be configured to use a single public network interface.



Note:

Before running Oracle Universal Installer, you must specify the Oracle Domain Services Cluster configuration details for the Oracle Member Cluster by creating the Member Cluster Manifest file.

Oracle Member Cluster for Oracle Database does not support Oracle Database 12.1 or earlier, where Oracle Member Cluster is configured with Oracle ASM storage as direct or indirect.

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 8-1 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, Oracle Domain Services Cluster, or Oracle Member Cluster.
- [Installing Oracle Standalone Cluster](#)
Complete this procedure to install Oracle Grid Infrastructure software for Oracle Standalone Cluster.
- [Installing Oracle Domain Services Cluster](#)
Complete this procedure to install Oracle Grid Infrastructure software for Oracle Domain Services Cluster.
- [Installing Oracle Member Clusters](#)
Complete this procedure to install Oracle Grid Infrastructure software for Oracle Member Cluster for Oracle Database and Oracle Member Cluster for Applications.

About Oracle Grid Infrastructure Installation

You can install Oracle Grid Infrastructure as Oracle Standalone Cluster, Oracle Domain Services Cluster, or Oracle Member 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. For Oracle Member Cluster installations, ensure that you have created a Member Cluster Manifest File as explained in this guide.

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/19.0.0/grid
chown grid:oinstall /u01/app/19.0.0/grid
cd /u01/app/19.0.0/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/19.0.0/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. Choose whether you want to create a Grid Infrastructure Management Repository for your Oracle Standalone Cluster installation, then click **Next**.

If you choose **Yes** on this window, then the Grid Infrastructure Management Repository Option window appears. Otherwise, the Create ASM Disk Group window appears.

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/19.0.0/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/19.0.0/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 Domain Services Cluster

Complete this procedure to install Oracle Grid Infrastructure software for Oracle Domain Services 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/19.0.0/grid
chown grid:oinstall /u01/app/19.0.0/grid
cd /u01/app/19.0.0/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/19.0.0/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 Domain Services Cluster**, then click **Next**.

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. This example uses the cluster name `mycluster` and the cluster SCAN name of `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 Create ASM Disk Group window appears.

- 8. 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**.

The Grid Infrastructure Management Repository Option window appears

- 9. 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 DATA1.
 - 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.
 - d. Select the **Configure Rapid Home Provisioning Server** option to configure an Oracle Fleet Patching and Provisioning (Oracle FPP) Server as part of the Oracle Domain Services Cluster. Oracle FPP enables you to install clusters, and provision, patch, and upgrade Oracle Grid Infrastructure and Oracle Database homes.

When you have finished providing the information for the disk group, click **Next**.

The Specify ASM Password window appears.

- 10. 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.

- 11. Select the option **Do not use Intelligent Platform Management Interface (IPMI)**, then click **Next**.

The Specify Management Options window appears.

- 12. 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.

You can manage Oracle Grid Infrastructure and Oracle Automatic Storage Management (Oracle ASM) using Oracle Enterprise Manager Cloud Control. To register the Oracle Grid Infrastructure cluster with Oracle Enterprise Manager, ensure that Oracle Management Agent is installed and running on all nodes of the cluster.

The Privileged Operating System Groups window appears.

13. Accept the default operating system group names for Oracle ASM administration and click **Next**.

The Specify Install Location window appears.

14. 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.

15. 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.

16. 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.

17. 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.

18. Review the contents of the Summary window and then click **Install**.

The installer displays a progress indicator enabling you to monitor the installation process.

19. 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/19.0.0/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. If your cluster has three or more nodes, then `root.sh` can be run concurrently on all nodes but the first. Node numbers are assigned according to the order of running `root.sh`. If you want to create a particular node number assignment, then run the root scripts in the order of the node assignments you want to make, and wait for the script to finish running on each node before proceeding to run the script on the next node. However, Oracle system identifier, or SID, for your Oracle RAC databases, do not follow the node numbers.

- 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/19.0.0/grid
[root@node2 grid]# ./root.sh
```

After the `root.sh` script completes, return to the OUI window where the Installer prompted you to run the `orainstRoot.sh` and `root.sh` scripts. Click **OK**.

The software installation monitoring window reappears.

When you run `root.sh` during Oracle Grid Infrastructure installation, the Trace File Analyzer (TFA) Collector is also installed in the directory `grid_home/tfa`.

20. After `root.sh` runs on all the nodes, OUI runs Net Configuration Assistant (`netca`) and Cluster Verification Utility. These programs run without user intervention.
21. During the installation, Oracle Automatic Storage Management Configuration Assistant (`asmca`) configures Oracle ASM for storage.

22. 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 Domain Services Cluster installation is complete, you can install Oracle Member Clusters for Oracle Databases and Oracle Member Clusters for Applications.

Installing Oracle Member Clusters

Complete this procedure to install Oracle Grid Infrastructure software for Oracle Member Cluster for Oracle Database and Oracle Member Cluster for Applications.

You must create a Member Cluster Manifest File as explained in this guide before performing the installation.

Use this procedure to install an Oracle Member Cluster for Oracle Databases or Oracle Member Cluster for Applications.

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/19.0.0/grid
chown grid:oinstall /u01/app/19.0.0/grid
cd /u01/app/19.0.0/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/19.0.0/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 either the **Configure an Oracle Member Cluster for Oracle Databases** or **Configure an Oracle Member Cluster for Applications** option, then click **Next**.

The Cluster Domain Services window appears.

5. Select the Manifest file that contains the configuration details about the management repository and other services for the Oracle Member Cluster.

For Oracle Member Cluster for Oracle Databases, you can also specify the Grid Naming Service and Oracle ASM Storage server details using a Member Cluster Manifest file.

Click **Next**.

6. If you selected to configure an Oracle Member Cluster for applications, then the Configure Virtual Access window appears. Provide a **Cluster Name** and optional **Virtual Host Name**.

The virtual host name serves as a connection address for the Oracle Member Cluster, and to provide service access to the software applications that you want the Oracle Member Cluster to install and run.

Click **Next**.

The Cluster Node Information window appears.

7. 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, Oracle Universal Installer (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.

- 8. Select the usage type for each network interface displayed, then click **Next**.

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**.

Note: For Oracle Member Cluster for Applications configurations, you can choose to have one public network interface, instead of requiring two network interfaces, one private or ASM, and the other public.

Click **Next**. The ASM Client Storage window appears.

- 9. Choose the disk group to store Oracle Cluster Registry (OCR) and voting files for the cluster on the Oracle Domain Services Cluster.

Click **Next**.

The Operating System Groups window appears.

- 10. Accept the default operating system group names for Oracle ASM administration and click **Next**.

The Specify Install Location window appears.

- 11. 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.

- 12. 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.

- 13. 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.

- 14. 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.

15. Review the contents of the Summary window and then click **Install**.

The installer displays a progress indicator enabling you to monitor the installation process.

16. 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 appears. Do *not* click **OK** until you have run 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/19.0.0/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. If your cluster has three or more nodes, then `root.sh` can be run concurrently on all nodes but the first. Node numbers are assigned according to the order of running `root.sh`. If you want to create a particular node number assignment, then run the root scripts in the order of the node assignments you want to make, and wait for the script to finish running on each node before proceeding to run the script on the next node. However, Oracle system identifier, or SID, for your Oracle RAC databases, do not follow the node numbers.

- 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/19.0.0/grid
[root@node2 grid]# ./root.sh
```

After the `root.sh` script completes, return to the OUI window where the Installer prompted you to run the `oraInstRoot.sh` and `root.sh` scripts. Click **OK**.

The software installation monitoring window reappears.

When you run `root.sh` during Oracle Grid Infrastructure installation, the Trace File Analyzer (TFA) Collector is also installed in the directory `grid_home/tfa`.

17. After `root.sh` runs on all the nodes, OUI runs Net Configuration Assistant (`netca`) and Cluster Verification Utility. These programs run without user intervention.
18. During installation of Oracle Member Cluster for Oracle Databases, if the Member Cluster Manifest file does not include configuration details for Oracle ASM, then Oracle Automatic Storage Management Configuration Assistant (`asmca`) configures Oracle ASM for storage.
19. 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, other applications, 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 providing cluster configuration information manually, or 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
node2 node2-vip
.
.
.
```

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] [site-name]
#
```

[illegible]

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.

Use this procedure to install Oracle Grid Infrastructure for a cluster software:

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/19.0.0/grid/gridSetup.sh [-debug] [-silent -responseFile filename]
```

For example:

```
$ cd /u01/app/19.0.0/grid
$ ./gridSetup.sh -responseFile /u01/app/grid/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
```

About Deploying Oracle Grid Infrastructure Using Oracle Fleet Patching and Provisioning

Oracle Fleet Patching and Provisioning (Oracle FPP) is a software lifecycle management method for provisioning and maintaining Oracle homes. Oracle Fleet Patching and Provisioning enables mass deployment and maintenance of standard operating environments for databases, clusters, and user-defined software types.



Note:

Starting with Oracle Grid Infrastructure 19c, the feature formerly known as Rapid Home Provisioning (RHP) is now Oracle Fleet Patching and Provisioning (Oracle FPP).

Oracle Fleet Patching and Provisioning enables you to install clusters, and provision, patch, scale, and upgrade Oracle Grid Infrastructure, Oracle Restart, and Oracle Database homes. The supported versions are 11.2, 12.1, 12.2, 18c, and 19c. You can also provision applications and middleware using Oracle Fleet Patching and Provisioning.

Oracle Fleet Patching and Provisioning is a service in Oracle Grid Infrastructure that you can use in either of the following modes:

- Central Oracle Fleet Patching and Provisioning Server

The Oracle Fleet Patching and Provisioning Server stores and manages standardized images, called gold images. Gold images can be deployed to any number of nodes across the data center. You can create new clusters and databases on the deployed homes and can use them to patch, upgrade, and scale existing installations.

The Oracle Fleet Patching and Provisioning Server can manage the following types of installations:

- Software homes on the cluster hosting the Oracle Fleet Patching and Provisioning Server itself.
- Oracle Fleet Patching and Provisioning Clients running Oracle Grid Infrastructure 12c Release 2 (12.2), 18c, and 19c.
- Installations running Oracle Grid Infrastructure 11g Release 2 (11.2) and 12c Release 1 (12.1).
- Installations running without Oracle Grid Infrastructure.

The Oracle Fleet Patching and Provisioning Server can provision new installations and can manage existing installations without requiring any changes to the existing installations. The Oracle Fleet Patching and Provisioning Server can automatically share gold images among peer servers to support enterprises with geographically distributed data centers.

- Oracle Fleet Patching and Provisioning Client

The Oracle Fleet Patching and Provisioning Client can be managed from the Oracle Fleet Patching and Provisioning Server, or directly by executing commands on the client itself. The Oracle Fleet Patching and Provisioning Client is a service built into the Oracle Grid Infrastructure and is available in Oracle Grid Infrastructure 12c Release 2 (12.2) and later releases. The Oracle Fleet Patching and Provisioning Client can retrieve gold images from the Oracle Fleet Patching and Provisioning Server, upload new images based on the policy, and apply maintenance operations to itself.

Oracle Fleet Patching and Provisioning

Deploying Oracle software using Oracle Fleet Patching and Provisioning has the following advantages:

- Ensures standardization and enables high degrees of automation with gold images and managed lineage of deployed software.
- Minimizes downtime by deploying new homes as images (called gold images) out-of-place, without disrupting active databases or clusters.
- Simplifies maintenance by providing automations which are invoked with a simple, consistent API across database versions and deployment models.
- Reduces maintenance risk with built-in validations and a “dry run” mode to test the operations.
- Enables you to resume or restart the commands in the event of an unforeseen issue, reducing the risk of maintenance operations.
- Minimizes and often eliminates the impact of patching and upgrades, with features that include:
 - Zero-downtime database upgrade with fully automated upgrade, executed entirely within the deployment without requiring any extra nodes or external storage.

- Adaptive management of database sessions and OJVM during rolling patching.
- Options for management of consolidated deployments.
- The deployment and maintenance operations enable customizations to include environment-specific actions into the automated workflow.

Related Topics

- *Oracle Clusterware Administration and Deployment Guide*

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 delays, 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 11g Release 2 (11.2) or later installations, 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.
- [About Changes in Default SGA Permissions for Oracle Database](#)
Starting with Oracle Database 12c Release 2 (12.2.0.1), by default, permissions to read and write to the System Global Area (SGA) are limited to the Oracle software installation owner.
- [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 19c 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 Monthly Recommended Patches (MRPs) patches for your Oracle software after you complete installation.

Downloading Release Update Patches

Download and install Release Updates (RU) and Monthly Recommended Patches (MRPs) patches for your Oracle software after you complete installation.

Oracle provides quarterly updates in the form of Release Updates (RU) and Monthly Recommended Patches (MRPs). 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. In the **sign in to other Oracle Support portals** section, click **Support for Oracle Hardware, Software, and Managed Cloud services** to log in to My Oracle Support.

 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 888.1](#)
- [Patch Delivery Methods for Oracle Database](#)

Recommended Postinstallation Tasks

Oracle recommends that you complete these tasks after installation.

- [Configuring IPMI-based Failure Isolation Using Crsctl](#)
On Oracle Solaris and AIX platforms, where Oracle does not currently support the native IPMI driver, DHCP addressing is not supported and manual configuration is required for IPMI support.
- [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](#)
Install the latest version of Oracle Autonomous Health Framework to perform proactive health checks and collect diagnostics data for the Oracle software stack.
- [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. 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.
- [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.

Configuring IPMI-based Failure Isolation Using Crsctl

On Oracle Solaris and AIX platforms, where Oracle does not currently support the native IPMI driver, DHCP addressing is not supported and manual configuration is required for IPMI support.

Oracle Universal Installer (OUI) will not collect the administrator credentials, so failure isolation must be manually configured, the BMC must be configured with a static IP address, and the address must be manually stored in the OLR.

Configure BMC as described in this guide.

1. If necessary, start Oracle Clusterware using the following command:

```
$ crsctl start crs
```

2. Use the BMC management utility to obtain the BMC's IP address and then use the cluster control utility `crsctl` to store the BMC's IP address in the Oracle Local Registry (OLR) by issuing the `crsctl set css ipmiaddr address` command. For example:

```
$ crsctl set css ipmiaddr 192.168.10.45
```

3. Enter the following `crsctl` command to store the user ID and password for the resident BMC in the OLR, where *youradminacct* is the IPMI administrator user account, and provide the password when prompted:

```
$ crsctl set css ipmiadmin youradminacct  
IPMI BMC Password:
```

This command attempts to validate the credentials you enter by sending them to another cluster node. The command fails if that cluster node is unable to access the local BMC using the credentials.

When you store the IPMI credentials in the OLR, you must have the anonymous user specified explicitly, or a parsing error will be reported.

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

Install the latest version of Oracle Autonomous Health Framework to perform proactive health checks and collect diagnostics data for the Oracle software stack.

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:

- Standalone Oracle Database
- Oracle Grid Infrastructure
- Oracle Real Application Clusters
- Maximum Availability Architecture (MAA) Validation
- Upgrade Readiness Validations
- Oracle GoldenGate

Oracle Autonomous Health Framework is pre-installed with Oracle Database. However, Oracle recommends that you update to the latest version of Oracle Autonomous Health Framework 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>

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/19.0.0/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 Groups 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**.

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. 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.

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=
```

About Changes in Default SGA Permissions for Oracle Database

Starting with Oracle Database 12c Release 2 (12.2.0.1), by default, permissions to read and write to the System Global Area (SGA) are limited to the Oracle software installation owner.

In previous releases, both the Oracle installation owner account and members of the OSDBA group had access to shared memory. The change in Oracle Database 12c Release 2 (12.2) and later releases to restrict access by default to the Oracle installation owner account provides greater security than previous configurations. However, this change may prevent DBAs who do not have access to the Oracle installation owner account from administering the database.

The Oracle Database initialization parameter `ALLOW_GROUP_ACCESS_TO_SGA` determines if the Oracle Database installation owner account (`oracle` in Oracle documentation examples) is the only user who can read and write to the database System Global Area (SGA), or if members of the OSDBA group can read the SGA. In Oracle Database 12c Release 2 (12.2) and later releases, the default value for this parameter is `FALSE`, so that only the Oracle Database installation owner has read and write permissions to the SGA. Group access to the SGA is removed by default. This change affects all Linux and UNIX platforms.

If members of the OSDBA group require read access to the SGA, then you can change the initialization parameter `ALLOW_GROUP_ACCESS_TO_SGA` setting from `FALSE` to `TRUE`. Oracle strongly recommends that you accept the default permissions that limit access to the SGA to the `oracle` user account.

Related Topics

- [Oracle Database Reference](#)

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 19c installations:

- [General Restrictions for Using Earlier Oracle Database Releases](#)
You can use Oracle Database 19c, 18c, Oracle Database 12c releases 1 and 2 and Oracle Database 11g release 2 (11.2.0.3 or later) with Oracle Grid Infrastructure 19c.
- [Configuring Earlier Release Oracle Database on Oracle ACFS](#)
Review this information to configure a 11.2 release Oracle Database on Oracle Automatic Storage Management Cluster File System (Oracle ACFS).
- [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 19c 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 19c, 18c, Oracle Database 12c releases 1 and 2 and Oracle Database 11g release 2 (11.2.0.3 or later) with Oracle Grid Infrastructure 19c.

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.

Oracle Database homes can only be stored on Oracle ASM Cluster File System (Oracle ACFS) if the database version is Oracle Database 11g release 2 or later.

When installing 11.2 databases on an Oracle Flex ASM cluster, the Oracle ASM cardinality must be set to `ALL`.



Note:

If you are installing Oracle Database 11g release 2 with Oracle Grid Infrastructure 12c release 2 (12.2), then before running Oracle Universal Installer (OUI) for Oracle Database, run the following command on the local node only:

```
Grid_home/oui/bin/runInstaller -ignoreSysPrereqs -updateNodeList
ORACLE_HOME=Grid_home "CLUSTER_NODES={comma_separated_list_of_nodes}"
CRS=true LOCAL_NODE=local_node [-cfs]
```

Use the `-cfs` option only if the *Grid_home* is on a shared location.

On IBM AIX on POWER Systems (64-Bit), when installing Oracle RAC database 12c release 1 (12.1.0.2) on Oracle Grid Infrastructure 12c release 2 (12.2), do the following:

1. Install and configure Oracle Grid Infrastructure 12c release 2 (12.2) as described in this guide.
2. Copy the Oracle RAC database 12c release 1 (12.1.0.2) shiphome locally on one of the nodes.
3. Edit the file `database/install/oraparam.ini` and comment the following line:

```
*****
EXT_LIB=../stage/ext/lib
*****
```

4. Save the changes and run the Oracle RAC database 12c release 1 (12.1.0.2) installer.

```
./runInstaller
```

Configuring Earlier Release Oracle Database on Oracle ACFS

Review this information to configure a 11.2 release Oracle Database on Oracle Automatic Storage Management Cluster File System (Oracle ACFS).

1. Install Oracle Grid Infrastructure 19c as described in this guide.
2. Start Oracle ASM Configuration Assistant (ASMCA) as the grid installation owner. For example:

```
./asmca
```

Follow the steps in the configuration wizard to create Oracle ACFS storage for the earlier release Oracle Database home.

3. Install Oracle Database 11g release 2 (11.2) software-only on the Oracle ACFS file system you configured.

4. From the 11.2 Oracle Database home, run Oracle Database Configuration Assistant (DBCA) and create the Oracle RAC Database, using Oracle ASM as storage for the database data files.

```
./dbca
```

5. Modify the Oracle ACFS path dependency:

```
srvctl modify database -d my_112_db -j Oracle_ACFS_path
```

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.

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.

You can 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 19c

Using the Correct LSNRCTL Commands

To administer Oracle Database 19c 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.

Unlock the home using the following procedure:

1. Go to the `/u01/app/19.0.0/grid/crs/install` directory, and unlock the Grid home using the command `rootcrs.sh -unlock`:

```
# cd /u01/app/19.0.0/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/19.0.0/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/19.0.0/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/19.0.0/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*".

Removing Oracle Database Software

These topics describe how to remove Oracle software and configuration files.

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 standalone 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.
- [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.
- [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.
- [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, if you modify the Oracle Grid Infrastructure binaries, then you must relink the binaries on Oracle Linux 8 or earlier releases.
- [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 existing 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.

- [Unconfiguring Oracle Member Cluster](#)
Run this procedure to unconfigure Oracle Member Cluster.

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 19c Oracle home to remove Oracle software from an existing 11.2.0.4 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:

- Dismount Oracle Automatic Storage Management Cluster File System (Oracle ACFS) and disable Oracle Automatic Storage Management Dynamic Volume Manager (Oracle ADVN).

- 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
<code>-silent</code>	Use this flag to run <code>deinstall</code> 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 <code>-silent</code> 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 <code>deinstall</code> with the <code>-checkonly</code> flag. <code>deinstall</code> then discovers information from the Oracle home to deinstall and deconfigure. It generates the response file that you can then use with the <code>-silent</code> option. You can also modify the template file <code>deinstall.rsp.tmpl</code>, located in the <code>\$ORACLE_HOME/deinstall/response</code> directory.
<code>-checkonly</code>	Use this flag to check the status of the Oracle software home configuration. Running <code>deinstall</code> with the <code>-checkonly</code> flag does not remove the Oracle configuration. The <code>-checkonly</code> flag generates a response file that you can then use with the <code>deinstall</code> command and <code>-silent</code> option.
<code>-paramfile</code> <i>complete path of input response file</i>	Use this flag to run <code>deinstall</code> 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 <code>\$ORACLE_HOME/deinstall/response</code>.
<code>-params</code> [<i>name1=value name2=value name3=value ...</i>]	Use this flag with a response file to override one or more values to change in a response file you have created.
<code>-o</code> <i>complete path of directory for saving response files</i>	Use this flag to provide a path other than the default location where the response file (<code>deinstall.rsp.tmpl</code>) is saved. The default location of the response file is <code>\$ORACLE_HOME/deinstall/response</code>.
<code>-tmpdir</code> <i>complete path of temporary directory to use</i>	Use this flag to specify a non-default location where <code>deinstall</code> writes the temporary files for the deinstallation.
<code>-logdir</code> <i>complete path of log directory to use</i>	Use this flag to specify a non-default location where <code>deinstall</code> writes the log files for the deinstallation.

Parameter	Description
<code>-local</code>	Use this flag on a multinode environment to deinstall Oracle software in a cluster. When you run <code>deinstall</code> with this flag, it deconfigures and deinstalls the Oracle software on the local node (the node where <code>deinstall</code> is run). On remote nodes, it deconfigures Oracle software, but does not deinstall the Oracle software.
<code>-skipLocalHomeDeletion</code>	Use this flag in Oracle Grid Infrastructure installations on a multinode environment to deconfigure a local Grid home without deleting the Grid home.
<code>-skipRemoteHomeDeletion</code>	Use this flag in Oracle Grid Infrastructure installations on a multinode environment to deconfigure a remote Grid home without deleting the Grid home.
<code>-help</code>	Use this option to obtain additional information about the command option flags.

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.

Unlock the Oracle Grid Infrastructure installation using the following procedure:

1. Log in as `root`, and change the permission and ownership of the previous release Grid home using the following command syntax, where *oldGH* is the previous release Grid home, *giowner* 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 giowner oldGH
#chown giowner oldGHParent
```

For example:

```
#chmod -R 755 /u01/app/18.0.0/grid
#chown -R grid /u01/app/18.0.0/grid
#chown grid /u01/app/18.0.0
```

2. As the `grid` user, and use the `-checkonly` parameter from previous release Grid home *oldGH/deinstall* directory to check the status of the Oracle software home configuration.

```
oldGH/deinstall/deinstall -checkonly -silent -o /tmp/ -tmpdir /tmp
```

3. Use the `deinstall` command from previous release Grid home `oldGH/deinstall` directory to deinstall the previous release Grid home.

```
oldGH/deinstall/deinstall -silent -tmpdir /tmp/rsp_file
```

4. As the `root` user, run the `rootdeinstall.sh` script from the `/tmp/deinstall<date>/` directory.

```
/tmp/deinstall<date>/rootdeinstall.sh
```

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 19.0.0.0 Oracle home to remove Oracle software from an existing 18.0.0.0 Oracle home.

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/19.0.0/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/19.0.0/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.

In this example, the Oracle Grid Infrastructure home is `/u01/app/19.0.0/grid`

```
$ cd /u01/app/19.0.0/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.

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/19.0.0/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/`, the central Oracle Inventory home (`oraInventory`) is `/u01/app/oraInventory`, the virtual IP addresses (VIP) are `192.0.2.2` and `192.0.2.4`, the local node (the node where you run the deinstallation session from) is `node1`:

```
# Copyright (c) 2005, 2016 Oracle Corporation. All rights reserved.
ORACLE_HOME=/u01/app/12.2.0/grid
CDATA_AUSIZE=4
BIG_CLUSTER=true
ISROLLING=true
LOCAL_NODE=node1
OCR_VD_DISKGROUPS="+DATA1"
MGMTDB_DIAG=/u01/app/grid
OCRID=
MGMTDB_SPFILE="+DATA1/_MGMTDB/PARAMETERFILE/spfile.271.923210081"
ObaseCleanupPtrLoc=/tmp/deinstall2016-10-06_09-36-04AM/utl/orabase_cleanup.lst
CDATA_BACKUP_QUORUM_GROUPS=
ASM_CREDENTIALS=
MGMTDB_NODE_LIST=node1,node2
EXTENDED_CLUSTER=false
LISTENER_USERNAME=cuser
local=false
inventory_loc=/u01/app/oraInventory
ORACLE_HOME=/u01/app/12.2.0/grid
ASM_HOME=/u01/app/grid
ASM_DISK_GROUPS="+DATA1"
HUB_NODE_VIPS=AUTO,AUTO
PING_TARGETS=
ORA_DBA_GROUP=oinstall
ASM_DISCOVERY_STRING=/dev/rds/*
CDATA_DISKS=/dev/rds/c0t600144F0C4A01A3F000056E6A12A0022d0s3
MinimumSupportedVersion=11.2.0.1.0
NEW_HOST_NAME_LIST=
ORACLE_HOME_VERSION=12.2.0.1.0
PRIVATE_NAME_LIST=
MGMTDB_DB_UNIQUE_NAME=_mgmtdb
ASM_DISKSTRING=/dev/rds/*,AFD:*
CDATA_QUORUM_GROUPS=
CRS_HOME=true
ODA_CONFIG=
JLIBDIR=/u01/app/jlib
CRFHOME="/u01/app/"
USER_IGNORED_PREREQ=true
MGMTDB_ORACLE_BASE=/u01/app/grid/
```



```

DROP_MGMTDB=true
RHP_CONF=false
OCRLOC=
GNS_TYPE=local
CRS_STORAGE_OPTION=1
CDATA_SITES=
GIMR_CONFIG=local
CDATA_BACKUP_SIZE=0
GPNPGCONFIGDIR=$ORACLE_HOME
MGMTDB_IN_HOME=true
CDATA_DISK_GROUP=+DATA2
LANGUAGE_ID=AMERICAN_AMERICA.AL32UTF8
CDATA_BACKUP_FAILURE_GROUPS=
CRS_NODEVIPS='AUTO/255.255.254.0/net0,AUTO/255.255.254.0/net0'
ORACLE_OWNER=cuser
GNS_ALLOW_NET_LIST=
silent=true
INSTALL_NODE=node1.example.com
ORACLE_HOME_VERSION_VALID=true
inst_group=oinstall
LOGDIR=/tmp/deinstall12016-10-06_09-36-04AM/logs/
EXTENDED_CLUSTER_SITES=
CDATA_REDUNDANCY=EXTERNAL
CDATA_BACKUP_DISK_GROUP=+DATA2
APPLICATION_VIP=
HUB_NODE_LIST=node1,node2
NODE_NAME_LIST=node1,node2
GNS_DENY_ITF_LIST=
ORA_CRS_HOME=/u01/app/12.2.0/grid/
JREDIR=/u01/app/12.2.0/grid/jdk/jre/
ASM_LOCAL_SID=+ASM1
ORACLE_BASE=/u01/app/
GNS_CONF=true
CLUSTER_CLASS=DOMAINSERVICES
ORACLE_BINARY_OK=true
CDATA_BACKUP_REDUNDANCY=EXTERNAL
CDATA_FAILURE_GROUPS=
ASM_CONFIG=near
OCR_LOCATIONS=
ASM_ORACLE_BASE=/u01/app/12.2.0/
OLRLOC=
GIMR_CREDENTIALS=
GPNPCONFIGDIR=$ORACLE_HOME
ORA_ASM_GROUP=asmadmin
GNS_CREDENTIALS=
CDATA_BACKUP_AUSIZE=4
GNS_DENY_NET_LIST=
OLD_CRS_HOME=
NEW_NODE_NAME_LIST=
GNS_DOMAIN_LIST=node1.example.com
ASM_UPGRADE=false
NETCA_LISTENERS_REGISTERED_WITH_CRS=LISTENER
CDATA_BACKUP_DISKS=/dev/rdisk/
ASMCA_ARGS=
CLUSTER_GUID=
CLUSTER_NODES=node1,node2

```

```

MGMTDB_NODE=node2
ASM_DIAGNOSTIC_DEST=/u01/app/
NEW_PRIVATE_NAME_LIST=
AFD_LABELS_NO_DG=
AFD_CONFIGURED=true
CLSCFG_MISSCOUNT=
MGMT_DB=true
SCAN_PORT=1521
ASM_DROP_DISKGROUPS=true
OPC_NAT_ADDRESS=
CLUSTER_TYPE=DB
NETWORKS="net0"/IP_Address:public,"net1"/IP_Address:asm,"net1"/
IP_Address:cluster_interconnect
OCR_VOTINGDISK_IN_ASM=true
HUB_SIZE=32
CDATA_BACKUP_SITES=
CDATA_SIZE=0
REUSEDG=false
MGMTDB_DATAFILE=
ASM_IN_HOME=true
HOME_TYPE=CRS
MGMTDB_SID="-MGMTDB"
GNS_ADDR_LIST=mycluster-gns.example.com
CLUSTER_NAME=node1-cluster
AFD_CONF=true
MGMTDB_PWDFILE=
OPC_CLUSTER_TYPE=
VOTING_DISKS=
SILENT=false
VNDR_CLUSTER=false
TZ=localtime
GPNP_PA=
DC_HOME=/tmp/deinstall2016-10-06_09-36-04AM/logs/
CSS_LEASEDURATION=400
REMOTE_NODES=node2
ASM_SPFILE=
NEW_NODEVIPS='n1-vip/255.255.252.0/eth0,n2-vip/255.255.252.0/eth0'
SCAN_NAME=node1-cluster-scan.node1-cluster.com
RIM_NODE_LIST=
INVENTORY_LOCATION=/u01/app/oraInventory

```

**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"/192.0.2.1\:public,"eth1"/10.0.0.1\:cluster_interconnect
VIP1_IP=192.0.2.2
```

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
```

Write down the configuration information for the server.

3. Stop all of the databases, services, and listeners that you discovered in step 1.
4. If present, unmount all Oracle Automatic Storage Management Cluster File System (Oracle ACFS) file systems.
5. As the `root` user, unconfigure the Oracle Grid Infrastructure installation for a standalone server (Oracle Restart).

```
# cd /u01/app/19.0.0/grid/crs/install
# roothas.sh -deconfig -force
```

6. As the `grid` user, update the Oracle central inventory (`oraInventory`).

```
$ /u01/app/19.0.0/grid/oui/bin/runInstaller -nowait -waitforcompletion
-ignoreSysPrereqs -updateNodeList -silent -local CRS=false
ORACLE_HOME=Grid_home
"CLUSTER_NODES=node_list" -doNotUpdateNodeList"
```

7. Open the `/etc/oratab` file and remove the entry corresponding to `oracle_restart_home`, in the following format:

```
+ASM:oracle_restart_home:N
```

8. 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.
9. 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.

- 10. As the Oracle Grid Infrastructure installation owner, create and stage the response file for this installation as described in *Recording Response Files*.
- 11. Complete the installation in the silent mode using the following command:

```
$ Grid_home/gridSetup.sh -silent -responseFile $ORACLE_HOME/GI.rsp
```

- 12. Run `root.sh`.
- 13. Mount the Oracle ASM disk group used by Oracle Restart.
- 14. If you used Oracle ACFS with Oracle Restart, then:
 - a. Start Oracle ASM Configuration Assistant (ASMCA). Run the `volenable` command to enable all Oracle Restart disk group volumes.
 - b. Mount all Oracle ACFS file systems manually.
- 15. Add back Oracle Clusterware services to the Oracle Clusterware home, using the information you wrote down in step 1, including adding back Oracle ACFS resources. 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
```

- 16. 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
```

17. 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, if you modify the Oracle Grid Infrastructure binaries, then you must relink the binaries on Oracle Linux 8 or earlier releases.



Note:

You do not need to relink the Oracle Clusterware and Oracle ASM binaries when you:

- Perform a fresh install of Oracle Grid Infrastructure on Oracle Linux 9 or later releases.
- Upgrade the operating system to Oracle Linux 9 or later releases.
- Apply any operating system patch to Oracle Linux 9 and later releases.



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. If present, unmount all Oracle Automatic Storage Management Cluster File System (Oracle ACFS) file systems.

As root user:

```
# 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 all
```

As root user 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.

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 existing 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/19.0.0/grid` and the new Grid home is `/u01/app/19c/grid`.

```
# mkdir /u01/app/19c/grid
# cp -pR /u01/app/19.0.0/grid /u01/app/19c/grid
```

2. Change the owner and group for the directory and files in the new Grid home.

```
# chown -R grid:oinstall /u01/app/19c/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/19c/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/19c/grid/root.sh
```

5. As the `grid` user, enable Oracle RAC, dNFS, and `ioracle` from the new Grid home.

```
$ export ORACLE_HOME=/u01/app/19c/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. In addition, dismount Oracle Automatic Storage Management Cluster File System (Oracle ACFS) and disable Oracle Automatic Storage Management Dynamic Volume Manager (Oracle ADVM) volumes.



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.

Unconfiguring Oracle Member Cluster

Run this procedure to unconfigure Oracle Member Cluster.

1. Run the deinstall tool to unconfigure Oracle Member Cluster:

```
Grid_home/deinstall/deinstall.sh
```

2. Complete the deinstallation by running the root script on all the nodes when prompted.

```
# rootcrs.sh -deconfig
```

3. Delete the Member Cluster Manifest File for the Oracle Member Cluster and stored on the Oracle Domain Services Cluster:

```
crsctl delete member_cluster_configuration member_cluster_name
```

Related Topics

- *Oracle Clusterware Administration and Deployment Guide*

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. Starting with Oracle ASM 11g Release 2 (11.2), Oracle ASM upgrades can be rolling upgrades.

You can also use Oracle Fleet Patching and Provisioning to upgrade Oracle Grid Infrastructure for a cluster.

- [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](#)
When you upgrade to Oracle Grid Infrastructure 19c, you upgrade to an Oracle Flex Cluster configuration.
- [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.
- [About Upgrading Oracle Grid Infrastructure Using Oracle Fleet Patching and Provisioning](#)
Oracle Fleet Patching and Provisioning (Oracle FPP) is a software lifecycle management method for provisioning and patching Oracle homes.
- [Applying Patches to Oracle Grid Infrastructure](#)
After you have upgraded Oracle Grid Infrastructure 19c, you can install individual software patches by downloading them from My Oracle Support.
- [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](#)
- [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 restore 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.

Related Topics

- [About Upgrading Oracle Grid Infrastructure Using Oracle Fleet Patching and Provisioning](#)
Oracle Fleet Patching and Provisioning (Oracle FPP) is a software lifecycle management method for provisioning and patching Oracle homes.

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 19c to Oracle Grid Infrastructure 18c, Oracle Grid Infrastructure 12c Release 2 (12.2), Oracle Grid Infrastructure 12c Release 1 (12.1), and Oracle Grid Infrastructure 11g Release 2 (11.2). 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

When you upgrade to Oracle Grid Infrastructure 19c, you upgrade to an Oracle Flex Cluster configuration.

Supported upgrade paths for Oracle Grid Infrastructure for this release are:

- Oracle Grid Infrastructure upgrade from 11g Release 2 (11.2.0.4) to Oracle Grid Infrastructure 19c.
- Oracle Grid Infrastructure upgrade from 12c Release 1 (12.1.0.2) to Oracle Grid Infrastructure 19c.
- Oracle Grid Infrastructure upgrade from 12c Release 2 (12.2) to Oracle Grid Infrastructure 19c.
- Oracle Grid Infrastructure upgrade from 18c release to Oracle Grid Infrastructure 19c.

Upgrade options from Oracle Grid Infrastructure 11g Release 2 (11.2.0.4), Oracle Grid Infrastructure 12c Release 1 (12.1.0.2), Oracle Grid Infrastructure 12c Release 2 (12.2), and Oracle Grid Infrastructure 18c to Oracle Grid Infrastructure 19c 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 19c.

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 19c upgrade.
- Oracle ASM and Oracle Clusterware both run in the Oracle Grid Infrastructure home.
- When you upgrade Oracle Grid Infrastructure, 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 19c, 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 19c 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 19c cluster, your release must be greater than or equal to Oracle Grid Infrastructure 11g Release 2 (11.2.0.4).

 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, then you must migrate them to an Oracle ASM disk group, certified NAS device, or a shared file system before upgrading to Oracle Grid Infrastructure 19c.
- If you want to upgrade Oracle Grid Infrastructure releases before Oracle Grid Infrastructure 11g Release 2 (11.2), where the OCR and voting files are on raw or block devices, then you must upgrade to Oracle Grid Infrastructure 12c Release 1 (12.1.0.2), and move the Oracle Cluster Registry (OCR) and voting files to an Oracle ASM disk group or a shared file system, before upgrading to Oracle Grid Infrastructure 19c.
- If you have Oracle Automatic Storage Management Cluster File System (Oracle ACFS) file systems on Oracle Grid Infrastructure 11g Release 2 (11.2.0.1), you upgrade Oracle Grid Infrastructure to any later release, and you take advantage of Redundant Interconnect Usage and add one or more additional private interfaces to the private network, then you must restart the Oracle ASM instance on each upgraded cluster member node.

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 19c.
- 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 19c 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 11-1 Upgrade Checklist for Oracle Grid Infrastructure Installation

Check	Task
Review Upgrade Guide for deprecation and desupport information that may affect upgrade planning.	<i>Oracle Database Upgrade Guide</i>
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 11-1 (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 <code>ocrcheck</code> 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, a new GIMR is created only if the source Grid home has a GIMR configured. Allocate additional storage space as described in <i>Oracle Clusterware Storage Space Requirements</i> . 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 1 (12.1), Oracle Grid Infrastructure 12c Release 2 (12.2), or Oracle Grid Infrastructure 18c to Oracle Grid Infrastructure 19c, move the Oracle ASM password file from file system to Oracle ASM, before proceeding with the upgrade using the following ASMCMD command: <pre>ASMCMD [+] > pwcopy --asm current_location_of_ASM_password_file_in_OS_directory +target_disk_group_name/orapwASM</pre> When upgrading from Oracle Grid Infrastructure 11g Release 2 (11.2) to Oracle Grid Infrastructure 19c, move the Oracle ASM password file from file system to Oracle ASM, after the upgrade.
 Note: Set <code>compatible.asm</code> to at least 12.1.0.2, 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.
Unset Environment variables	As the user performing the upgrade, unset the environment variables <code>\$ORACLE_HOME</code> and <code>\$ORACLE_SID</code> . Check that the <code>\$ORA_CRS_HOME</code> environment variable is not set. Do not use <code>\$ORA_CRS_HOME</code> as an environment variable, except under explicit direction from Oracle Support. Refer to the <i>Checks to Complete Before Upgrading Oracle Grid Infrastructure</i> for a complete list of environment variables to unset.
Check system upgrade readiness with dry-run upgrade	Run Oracle Grid Infrastructure installation wizard, <code>gridSetup.sh</code> , in dry-run upgrade mode to perform system readiness checks for Oracle Grid Infrastructure upgrades.
Oracle ORAchK Upgrade Readiness Assessment	Download and run the Oracle ORAchK 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.
HugePages memory allocation	Allocate memory to HugePages large enough for the System Global Areas (SGA) of all databases planned to run on the cluster, and to accommodate the System Global Area for the Grid Infrastructure Management Repository.

Table 11-1 (Cont.) Upgrade Checklist for Oracle Grid Infrastructure Installation

Check	Task
Remove encryption of Oracle ACFS File Systems Before Upgrading from Oracle Grid Infrastructure 11g Release 2 (11.2.0.4)	To avoid data corruption, ensure that encryption of Oracle ACFS file systems is removed before upgrading from Oracle Grid Infrastructure 11g Release 2 (11.2.0.4). If you do not want to remove Oracle ACFS file system encryption, see My Oracle Support note 2147979.1 for an alternate: https://support.oracle.com/rs?type=doc&id=2147979.1

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 `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. Check to 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 11.2.0.3, 11.2.0.4, 12.1.0.1, 12.1.0.2, 12.2, 18c, and 19c. 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*
- [My Oracle Support Note 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.

- 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.

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]
```

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: 19.0.0.0.0.
- `-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.

```
$ /u01/app/19.0.0/grid/runcluvfy.sh stage -pre crsinst -upgrade -rolling -src_crshome  
/u01/app/18.0.0/grid -dest_crshome /u01/app/19.0.0/grid -dest_version  
19.0.0.0.0 -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 19c 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.

- As `grid` user, download the Oracle Grid Infrastructure image file and extract the file to the Grid home.

For example:

```
$ mkdir -p /u01/app/19.0.0/grid
$ chown grid:oinstall /u01/app/19.0.0/grid
$ cd /u01/app/19.0.0/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.

- Start the Oracle Grid Infrastructure installation wizard in dry-run upgrade mode.

```
$ /u01/app/19.0.0/grid/gridSetup.sh -dryRunForUpgrade
```

- Select **Upgrade Oracle Grid Infrastructure** option to perform dry-run upgrade for Oracle Grid Infrastructure (Oracle Clusterware and Oracle ASM).
- 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.
- Run `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 `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/19.0.0/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.

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 `grid` user, download the Oracle Grid Infrastructure image files and extract the files to the Grid home.

For example:

```
mkdir -p /u01/app/19.0.0/grid
chown grid:oinstall /u01/app/19.0.0/grid
cd /u01/app/19.0.0/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/19.0.0/grid/gridSetup.sh
```

3. Select the following configuration option:

- **Upgrade Oracle Grid Infrastructure:** Select this option to upgrade Oracle Grid Infrastructure (Oracle Clusterware and Oracle ASM).

 Note:

Oracle Clusterware must always be the later release, so you cannot upgrade Oracle ASM to a release that is more recent than Oracle Clusterware.

4. On the Node Selection page, select all nodes.
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 `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 `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 19c, 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.

You can confirm that the upgrade is incomplete by entering the command `crsctl query crs activeversion`.

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/19.0.0/grid/rootupgrade -force
```

This command forces the upgrade to complete. Verify that the upgrade has completed by using the command `crsctl query crs activeversion`. The active release should be the upgrade release.

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.

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 19c 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 19c *Grid_home* directory. For example:

```
$ cd /u01/app/19.0.0/grid/
```

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
```

About Upgrading Oracle Grid Infrastructure Using Oracle Fleet Patching and Provisioning

Oracle Fleet Patching and Provisioning (Oracle FPP) is a software lifecycle management method for provisioning and patching Oracle homes.

Oracle Fleet Patching and Provisioning enables you to install clusters, and provision, patch, and upgrade Oracle Grid Infrastructure, Oracle Database, and Oracle Restart homes. The supported versions are 11.2, 12.1, 12.2, 18c, and 19c. You can also provision applications and middleware using Oracle FPP. A single cluster, known as the Oracle Fleet Patching and Provisioning Server, stores and manages standardized images, called gold images, which can be provisioned to any number of nodes. You can install Oracle Grid Infrastructure cluster configurations such as Oracle Standalone Clusters, Oracle Member Clusters, and Oracle Member Cluster for Applications. After deployment, you can expand and contract clusters and Oracle RAC Databases.

You can provision Oracle Grid Infrastructure on a remote set of nodes in a cloud computing environment from a single cluster where you store templates of Oracle homes as images (called gold images) of Oracle software, such as databases, middleware, and applications.

Oracle Fleet Patching and Provisioning

Deploying Oracle software using Oracle FPP has the following advantages:

- Ensures standardization and enables high degrees of automation with gold images and managed lineage of deployed software.
- Supports change management. With standardized Oracle homes, an administrator has better control of the hosted Oracle software and can easily manage the mass deployment and maintenance of the software through a single location for change management.
- Minimizes downtime during patching and upgrades, eases rollbacks, and makes provisioning for large systems easier and more efficient.
- Reduces the cumulative time to patch software images, since a single Oracle home may be used for many database instances.

See Also:

Oracle Clusterware Administration and Deployment Guide for information about setting up the Oracle Fleet Patching and Provisioning Server and Client, creating and using gold images for provisioning and patching Oracle Grid Infrastructure, Oracle Database, and Oracle Restart homes.

Applying Patches to Oracle Grid Infrastructure

After you have upgraded Oracle Grid Infrastructure 19c, 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 OPatch 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.
- [Patching and Switching Oracle Grid Infrastructure Homes](#)
Perform an out-of-place Oracle Restart patching by switching from the current Oracle Grid Infrastructure home to a patched Oracle Grid Infrastructure home.

See Also:

Oracle Clusterware Administration and Deployment Guide for more information about patching Oracle Grid Infrastructure using Oracle Fleet Patching and Provisioning.

About Individual Oracle Grid Infrastructure Patches

Download Oracle ASM individual (one-off) patch and apply it to Oracle Grid Infrastructure using the OPatch 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 OPatch 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 OPatch 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 applied to the Oracle Grid Infrastructure software release, including Oracle ASM. The release is the release number, in the format of major, minor, and patch set release number. For example, with the release number 19.1.0.1, the major release is 19, the minor release is 1, and 0.0 is the patch set number. With one-off patches, the major and minor release remains the same, though the patch levels change each time you apply or roll back an interim patch.

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](#)
Starting with Oracle Grid Infrastructure 18c, 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

Starting with Oracle Grid Infrastructure 18c, 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/19.0.0/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.

Related Topics

- *Oracle Grid Infrastructure Upgrading Oracle Restart for Linux and Unix-Based Operating Systems*

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/patch 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. Check if CRS is in the `ROLLING PATCH` mode.

```
$ crsctl query crs activeversion -f
```

If CRS is not in the `ROLLING PATCH` mode, then set CRS in `ROLLING PATCH` mode from a node where CRS is online.

```
$ crsctl start rollingpatch
```

9. Start the Oracle Clusterware stack on the first node.

```
# $ORACLE_HOME/bin/crsctl start crs
```

10. Update Oracle Cluster Registry (OCR) to the patch level on the first node.

```
# $ORACLE_HOME/bin/clscfg -patch
```

11. Repeat steps 3 through 10 on all other cluster nodes.

You can patch all other cluster nodes in parallel.

12. Transitions Oracle Clusterware out of the `ROLLING PATCH` mode.

```
$ crsctl stop rollingpatch
```

Patching and Switching Oracle Grid Infrastructure Homes

Perform an out-of-place Oracle Restart patching by switching from the current Oracle Grid Infrastructure home to a patched Oracle Grid Infrastructure home.

1. Download the 19.3 Oracle Grid Infrastructure base release image files.

<https://www.oracle.com/database/technologies/oracle-database-software-downloads.html#19c>

2. As the `grid` user, extract the downloaded image files into a new Oracle Grid Infrastructure home directory.

```
$ mkdir -p /u01/app/oracle/product/19.17.0/grid
$ chown grid:oinstall /u01/app/oracle/product/19.17.0/grid
$ cd /u01/app/oracle/product/19.17.0/grid
$ unzip -q download_location/grid.zip
```

Here:

- `/u01/app/oracle/product/19.17.0/grid` is the new Grid home.
- `/u01/app/oracle/product/19.16.0/grid` is the old Grid home.

3. As the `grid` user, download and install the latest version of the OPatch utility in the new Grid home.

<https://updates.oracle.com/download/6880880.html>

```
$ mv /u01/app/oracle/product/19.17.0/grid/OPatch /u01/app/oracle/product/19.17.0/grid/bak_OPatch
$ unzip latest_Opatch.zip -d /u01/app/oracle/product/19.17.0/grid/
```

4. Download the Oracle Database RU version that you want to apply from [My Oracle Support](#). In this example, Oracle Database 19.17 RU.

For more information, see, [Downloading Release Update Patches](#)

5. Start the Oracle Grid Infrastructure installer to perform a software-only Oracle Restart installation. You can apply the optional `-applyRU` or `-applyOneOff` flags to apply Release Updates (RUs) during the installation.

```
$ /u01/app/oracle/product/19.17.0/grid/gridSetup.sh [-applyRU
patch_directory_location]
[-applyOneOffs comma_separated_list_of_patch_directory_locations]
```

6. Follow the steps in the configuration wizard to complete the Oracle Grid Infrastructure installation.
7. As the `root` user, run the following command to prepare the new home for the out-of-place patching:

```
# /u01/app/oracle/product/19.17.0/grid/crs/install/rootcrs.sh -prepatch -
dstcrshome
/u01/app/oracle/product/19.17.0/grid
```

This command does not shut down any services.

8. Run the following command to switch to the new Oracle Grid Infrastructure home and perform the out of place patching:

```
# /u01/app/oracle/product/19.17.0/grid/crs/install/rootcrs.sh -postpatch -
dstcrshome
/u01/app/oracle/product/19.17.0/grid
```

This command shuts down the old Oracle Grid Infrastructure home and starts resources from the new Oracle Grid Infrastructure home. All Oracle Grid Infrastructure services start running from the new Grid home.

9. Update the Oracle central inventory (oraInventory).

```
$ /u01/app/oracle/product/19.17.0/grid/oui/bin/runInstaller -
updateNodeList ORACLE_HOME=/u01/app/oracle/product/19.17.0/grid CRS=TRUE
$ /u01/app/oracle/product/19.16.0/grid/oui/bin/runInstaller -
updateNodeList ORACLE_HOME=/u01/app/oracle/product/19.16.0/grid CRS=FALSE
```

10. If the patching fails, then perform the following steps to rollback the patch:

- a. As the `root` user, run the `prepatch` script.

```
# /u01/app/19.17.0/grid/crs/install/rootcrs.sh -prepatch -dstcrshome
Old_GI_Home -rollback
```

- b. As the `root` user, run the `postpatch` script.

```
# /u01/app/19.17.0/grid/crs/install/rootcrs.sh -postpatch -dstcrshome
Old_GI_Home -rollback
```

11. If you have successfully switched to the new Grid home on all nodes and want to switch back to the old Grid home, then perform the following steps:

- a. As the `root` user, run the `prepatch` script.

```
# Old_GI_Home/crs/install/rootcrs.sh -prepatch -dstcrshome Old_GI_Home
```


- b. As the `grid` user, run the `postpatch` script.

```
# Old_GI_Home/crs/install/rootcrs.sh -postpatch -dstcrshome Old_GI_Home
```

- c. Update the Oracle central inventory (`oraInventory`).

```
$ /u01/app/oracle/product/19.16.0/grid/oui/bin/runInstaller -
updateNodeList ORACLE_HOME=/u01/app/oracle/product/19.16.0/grid CRS=TRUE
$ /u01/app/oracle/product/19.17.0/grid/oui/bin/runInstaller -
updateNodeList ORACLE_HOME=/u01/app/oracle/product/19.17.0/grid
CRS=FALSE
```

Updating Oracle Enterprise Manager Cloud Control Target Parameters

After upgrading Oracle Grid Infrastructure, upgrade the Enterprise Manager Cloud Control target.

Because Oracle Grid Infrastructure 19c 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 1-3 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.

Unlock the Oracle Grid Infrastructure installation using the following procedure:

1. Log in as `root`, and change the permission and ownership of the previous release Grid home using the following command syntax, where *oldGH* is the previous release Grid home, *giowner* 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 giowner oldGH
#chown giowner oldGHParent
```

For example:

```
#chmod -R 755 /u01/app/18.0.0/grid
#chown -R grid /u01/app/18.0.0/grid
#chown grid /u01/app/18.0.0
```

2. As the `grid` user, and use the `-checkonly` parameter from previous release Grid home *oldGH/deinstall* directory to check the status of the Oracle software home configuration.

```
oldGH/deinstall/deinstall -checkonly -silent -o /tmp/ -tmpdir /tmp
```

3. Use the `deinstall` command from previous release Grid home *oldGH/deinstall* directory to deinstall the previous release Grid home.

```
oldGH/deinstall/deinstall -silent -tmpdir /tmp/rsp_file
```

4. As the `root` user, run the `rootdeinstall.sh` script from the `/tmp/deinstall<date>/` directory.

```
/tmp/deinstall<date>/rootdeinstall.sh
```

 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 19.0.0.0.0 Oracle home to remove Oracle software from an existing 18.0.0.0.0 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 restore Oracle Clusterware to the previous release.

Downgrading Oracle Clusterware restores the Oracle Clusterware configuration to the state it was in before the Oracle Grid Infrastructure 19c upgrade. Any configuration changes you performed during or after the Oracle Grid Infrastructure 19c 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:

- Starting with Oracle Grid Infrastructure 12c Release 2 (12.2), 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 19c to earlier releases.
- [Restrictions for Oracle Grid Infrastructure Downgrades](#)
Review the following information for restrictions and changes for downgrading Oracle Grid Infrastructure installations.
- [Downgrading Oracle Clusterware to 18c](#)
Downgrade procedures for downgrading all Oracle Grid Infrastructure cluster types to 18c after successful or failed upgrade.
- [Downgrading Oracle Clusterware to 12c Release 2 \(12.2\)](#)
Downgrade procedures for downgrading all Oracle Grid Infrastructure cluster types to 12c Release 2 (12.2) after successful or failed upgrade.
- [Downgrading to Oracle Grid Infrastructure 12c Release 1 \(12.1\)](#)
Use this procedure to downgrade to Oracle Grid Infrastructure 12c Release 1 (12.1).
- [Downgrading to Oracle Grid Infrastructure 11g Release 2 \(11.2\)](#)
Use this procedure to downgrade to Oracle Grid Infrastructure 11g Release 2 (11.2).
- [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.

Options for Oracle Grid Infrastructure Downgrades

You can downgrade Oracle Grid Infrastructure 19c to earlier releases.

Downgrade options include the following earlier releases:

- Oracle Grid Infrastructure downgrade to Oracle Grid Infrastructure 18c.
- Oracle Grid Infrastructure downgrade to Oracle Grid Infrastructure 12c Release 2 (12.2).
- Oracle Grid Infrastructure downgrade to Oracle Grid Infrastructure 12c Release 1 (12.1).
- Oracle Grid Infrastructure downgrade to Oracle Grid Infrastructure 11g Release 2 (11.2).
Because all cluster configurations in Oracle Grid Infrastructure 19c are Oracle Flex Clusters, when you downgrade to Oracle Grid Infrastructure 11g Release 2 (11.2), you downgrade from an Oracle Flex cluster configuration to a Standard cluster configuration.



Note:

When you downgrade Oracle Grid Infrastructure to an earlier release, for example from Oracle Grid Infrastructure 19c to Oracle Grid Infrastructure 18c, 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.

- When downgrading Oracle Member Cluster to 12c Release 2 (12.2), you must start Oracle Clusterware on last downgraded node first and then on other nodes.
- When you downgrade from an Oracle Grid Infrastructure 12c Release 2 (12.2) or later releases to Oracle Grid Infrastructure 11g Release 2 (11.2), you downgrade from an Oracle Flex cluster configuration to a Standard cluster configuration since all cluster configurations in releases earlier than Oracle Grid Infrastructure 12c are Standard cluster configurations.
- You can only downgrade to the Oracle Grid Infrastructure release you upgraded from. For example, if you upgraded from Oracle Grid Infrastructure 11g Release 2 (11.2) to Oracle Grid Infrastructure 19c, you can only downgrade to Oracle Grid Infrastructure 11g Release 2 (11.2).

Downgrading Oracle Clusterware to 18c

Downgrade procedures for downgrading all Oracle Grid Infrastructure cluster types to 18c after successful or failed upgrade.

- [Downgrading Oracle Standalone Cluster to 18c](#)
Use this procedure to downgrade Oracle Standalone Cluster to Oracle Grid Infrastructure 18c after a successful upgrade.
- [Downgrading Oracle Domain Services Cluster to 18c](#)
Use this procedure to downgrade Oracle Domain Services Cluster to Oracle Grid Infrastructure 18c after a successful upgrade.
- [Downgrading Oracle Member Cluster to 18c](#)
Use this procedure to downgrade Oracle Member Cluster to Oracle Grid Infrastructure 18c after a successful upgrade.
- [Downgrading Oracle Grid Infrastructure to 18c 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 Standalone Cluster to 18c

Use this procedure to downgrade Oracle Standalone Cluster to Oracle Grid Infrastructure 18c after a successful upgrade.

1. As `grid` user, delete the Oracle Grid Infrastructure 19c Management Database:

```
$ $ORACLE_HOME/bin/dbca -silent -deleteDatabase -sourceDB -MGMTDB
```

2. As `root` user, use the command syntax `rootcrs.sh -downgrade` from 19c 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. As `root` user, downgrade the last node after you downgrade all other nodes:

```
# $ORACLE_HOME/crs/install/rootcrs.sh -downgrade
```

4. As `grid` user, remove Oracle Grid Infrastructure 19c 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/19.0.0/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/19.0.0/grid "CLUSTER_NODES=node1,node2,node3" -
doNotUpdateNodeList
```

Add the flag `-cfs` if the Grid home is a shared home.

5. As `grid` user, set Oracle Grid Infrastructure 18c 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=/u01/app/18.0.0/grid
"CLUSTER_NODES=node1,node2,node3"
```

6. As `root` user, start the 18c Oracle Clusterware stack on all nodes.

```
# crsctl start crs
```

7. As `grid` user, from any Oracle Grid Infrastructure 18c node, remove the MGMTDB resource as follows:

```
$ $ORACLE_HOME/bin/srvctl remove mgmtdb -f
```

8. As `grid` user, run DBCA in the silent mode from the 18c Grid home and create the Management Database container database (CDB) as follows:

```
$ $ORACLE_HOME/bin/dbca -silent -createDatabase -createAsContainerDatabase true
-templateName MGMTSeed Database.dbc -sid -MGMTDB -gdbName _mgmtdb
-storageType ASM -diskGroupName ASM_DG_NAME
-datafileJarLocation /u01/app/18.0.0/grid/assistants/dbca/templates
-characterSet AL32UTF8 -autoGeneratePasswords -skipUserTemplateCheck
```

9. Configure the Management Database by running the Configuration Assistant from the location `$ORACLE_HOME/bin/mgmtca -local`.

10. As `grid` user, run the `post_gimr_ugdg.pl` script from 19c Grid home:

```
$ $ORACLE_HOME/crs/install/post_gimr_ugdg.pl -downgrade -clusterType SC -
destHome /u01/app/19.0.0/grid
-lowerVersion 18.0.0.0.0 -oraBase /u01/app/grid2
```

Where:

SC is the type of the cluster as Oracle Standalone Cluster. The value for `-clusterType` can be SC for Oracle Standalone Cluster, DSC for Oracle Domain Services Cluster, or MC for Oracle Member Cluster.

`/u01/app/19.0.0/grid` is the Oracle home for Oracle Grid Infrastructure 19c.

`18.0.0.0.0` is the version of Oracle Grid Infrastructure to which you are downgrading.

`/u01/app/grid2` is the Oracle base for Oracle Grid Infrastructure 19c

Downgrading Oracle Domain Services Cluster to 18c

Use this procedure to downgrade Oracle Domain Services Cluster to Oracle Grid Infrastructure 18c after a successful upgrade.

1. As `grid` user, downgrade the Management Database to Oracle Grid Infrastructure 18c:
 - a. Manually copy the most recent time zone files from 19c Grid home to 18c Grid home, where `timezlg_number` is the name of the most recent `timzlg` file and `timezone_number` is the name of the most recent `timezone` file:

```
$ cp $ORACLE_HOME/oracore/zoneinfo/timezlg_number.dat /u01/app/18.0.0/
grid/oracore/zoneinfo/timezlg_number.dat
$ cp $ORACLE_HOME/oracore/zoneinfo/timezone_number.dat /u01/app/18.0.0/
grid/oracore/zoneinfo/timezone_number.dat
```

- b. Downgrade application schema using the following command syntax from 19c Grid home:

```
$ $ORACLE_HOME/bin/mgmtua downgrade -local -oldOracleHome /u01/app/
18.0.0/grid -skipSystemSchemaDowngrade
```


c. Disable and stop MGMTDB resource from 19c Grid home:

```
$ cd $ORACLE_HOME/bin
$ ./srvctl disable mgmtdb
$ ./srvctl stop mgmtdb
```

d. Downgrade system schema using the following procedure:

i. Set ORACLE_SID environment variable for 19c Grid home:

```
$ export ORACLE_SID=-MGMTDB
$ cd $ORACLE_HOME/bin
```

ii. Start CDB and all PDBs in downgrade mode:

```
$ ./sqlplus / as sysdba
SQL> startup downgrade
SQL> alter pluggable database all open downgrade;
SQL> exit
```

iii. Downgrade 19c Management Database using the following command syntax, where /u01/app/grid2 is the Oracle base for Oracle Grid Infrastructure 19c:

```
$ $ORACLE_HOME/perl/bin/perl $ORACLE_HOME/rdbms/admin/catcon.pl -d
/u01/app/grid2 -e -l /u01/app/grid2/cfgtoollogs/mgmtua -b
mgmtdowngrade -r
$ORACLE_HOME/rdbms/admin/catdwgrd.sql
```

iv. Set ORACLE_HOME and ORACLE_SID environment variables for 18c Grid home:

```
$ export ORACLE_HOME=/u01/app/18.0.0/grid/
$ export ORACLE_SID=-MGMTDB
$ cd $ORACLE_HOME/bin
```

v. Start CDB and all PDBs in upgrade mode:

```
$ ./sqlplus / as sysdba
SQL> shutdown immediate
SQL> startup upgrade
SQL> alter pluggable database all open upgrade;
SQL> exit
```

vi. Run catrelod script for 18c Management Database using the following command syntax, where /u01/app/grid is the Oracle base for Oracle Grid Infrastructure 18c:

```
$ $ORACLE_HOME/perl/bin/perl $ORACLE_HOME/rdbms/admin/catcon.pl -d
/u01/app/grid -e -l /u01/app/grid/cfgtoollogs/mgmtua -b
mgmtdowngrade
$ORACLE_HOME/rdbms/admin/catrelod.sql
```

- vii. Recompile all invalid objects after downgrade using the following command syntax from 18c Grid home:

```
$ $ORACLE_HOME/perl/bin/perl $ORACLE_HOME/rdbms/admin/catcon.pl -d
/u01/app/grid -e -l /u01/app/grid/cfgtoollogs/mgmtua -b
mgmtdowngrade
$ORACLE_HOME/rdbms/admin/utlpr.sql
```

- e. Manually stop the Management Database:

```
$ ./sqlplus / as sysdba
SQL> shutdown immediate
SQL> exit
```

2. As `root` user, use the following command syntax `rootcrs.sh -downgrade` from 19c 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. As `root` user, downgrade the last node after you downgrade all other nodes:

```
# $ORACLE_HOME/crs/install/rootcrs.sh -downgrade
```

4. As `grid` user, remove Oracle Grid Infrastructure 19c 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/19.0.0/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/19.0.0/grid
"CLUSTER_NODES=node1,node2,node3"
-doNotUpdateNodeList
```

Add the flag `-cfs` if the Grid home is a shared home.

5. As `root` user, start the 18c Oracle Clusterware stack on all nodes.

```
# crsctl start crs
```

6. As `grid` user, set Oracle Grid Infrastructure 18c 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=/u01/app/18.0.0/grid
"CLUSTER_NODES=node1,node2,node3"
```

7. As `grid` user, downgrade CHA models from any node where the Grid Infrastructure stack is running from 18c Grid home and Management Database and `ochad` are up:

```
$ ./srvctl stop cha
$ ./chactl import model -file /u01/app/18.0.0/grid/cha/model/os_gold.svm -
name DEFAULT_CLUSTER
$ ./chactl import model -file /u01/app/18.0.0/grid/cha/model/db_gold.svm -
name DEFAULT_DB
$ ./srvctl start cha
```

In the example above, `DEFAULT_CLUSTER` and `DEFAULT_DB` are function names that you must pass as values.

Downgrading Oracle Member Cluster to 18c

Use this procedure to downgrade Oracle Member Cluster to Oracle Grid Infrastructure 18c after a successful upgrade.

1. As `grid` user, use the command syntax `mgmtua downgrade` from 19c Grid home to downgrade Oracle Member Cluster where `oldOracleHome` is 18c Grid home and `version` is the five digit release number:

```
$ $ORACLE_HOME/bin/mgmtua downgrade -member -oldOracleHome /u01/app/18.0.0/
grid -toVersion version
```

2. As `root` user, use the command syntax `rootcrs.sh -downgrade` from 19c 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. As `root` user, downgrade the last node after you downgrade all other nodes:

```
# $ORACLE_HOME/crs/install/rootcrs.sh -downgrade
```

4. As `grid` user, remove Oracle Grid Infrastructure 19c 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/19.0.0/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/19.0.0/grid
"CLUSTER_NODES=node1,node2,node3"
-doNotUpdateNodeList
```

Add the flag `-cfs` if the Grid home is a shared home.

5. As `root` user, start the 18c Oracle Clusterware stack on all nodes.

```
# crsctl start crs
```

Note:

You must start Oracle Clusterware on last downgraded node first, and then on other nodes.

6. As `grid` user, set Oracle Grid Infrastructure 18c 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=/u01/app/18.0.0/grid
"CLUSTER_NODES=node1,node2,node3"
```

7. As `grid` user, downgrade CHA models from any node where the Grid Infrastructure stack is running from 18c Grid home and Management Database and `ochad` are up:

```
$ ./srvctl stop cha
$ ./chactl import model -file /u01/app/18.0.0/grid/cha/model/os_gold.svm -
name DEFAULT_CLUSTER
$ ./chactl import model -file /u01/app/18.0.0/grid/cha/model/db_gold.svm -
name DEFAULT_DB
$ ./srvctl start cha
```

In the example above, `DEFAULT_CLUSTER` and `DEFAULT_DB` are function names that you must pass as values.

Downgrading Oracle Grid Infrastructure to 18c 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/19.0.0/grid/gridSetup.sh -silent -downgrade [-nodes]
nodes_to_be_downgraded
[-oldHome] previous_release_grid_home_to_downgrade_to
[-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 Clusterware to 12c Release 2 (12.2)

Downgrade procedures for downgrading all Oracle Grid Infrastructure cluster types to 12c Release 2 (12.2) after successful or failed upgrade.

- [Downgrading Oracle Standalone Cluster to 12c Release 2 \(12.2\)](#)
Use this procedure to downgrade Oracle Standalone Cluster to Oracle Grid Infrastructure 12c Release 2 (12.2) after a successful upgrade.
- [Downgrading Oracle Domain Services Cluster to 12c Release 2 \(12.2\)](#)
Use this procedure to downgrade Oracle Domain Services Cluster to Oracle Grid Infrastructure 12c Release 2 (12.2) after a successful upgrade.
- [Downgrading Oracle Member Cluster to 12c Release 2 \(12.2\)](#)
Use this procedure to downgrade Oracle Member Cluster to Oracle Grid Infrastructure 12c Release 2 (12.2) after a successful upgrade.
- [Downgrading Oracle Grid Infrastructure to 12c Release 2 \(12.2\) 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 Standalone Cluster to 12c Release 2 (12.2)

Use this procedure to downgrade Oracle Standalone Cluster to Oracle Grid Infrastructure 12c Release 2 (12.2) after a successful upgrade.

1. As `grid` user, delete the Oracle Grid Infrastructure 19c Management Database:

```
$ $ORACLE_HOME/bin/dbca -silent -deleteDatabase -sourceDB -MGMTDB
```

2. As `root` user, use the command syntax `rootcrs.sh -downgrade` from 19c 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. As `root` user, downgrade the last node after you downgrade all other nodes:

```
# $ORACLE_HOME/crs/install/rootcrs.sh -downgrade
```

4. As `grid` user, remove Oracle Grid Infrastructure 19c 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/19.0.0/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/19.0.0/grid "CLUSTER_NODES=node1,node2,node3" -
doNotUpdateNodeList
```

Add the flag `-cfs` if the Grid home is a shared home.

5. As `grid` user, set Oracle Grid Infrastructure 12c Release 2 (12.2) 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=/u01/app/12.2.0/grid
"CLUSTER_NODES=node1,node2,node3"
```

6. As `root` user, start the 12c Release 2 (12.2) Oracle Clusterware stack on all nodes.

```
# crsctl start crs
```

7. As `grid` user, from any Oracle Grid Infrastructure 12c Release 2 (12.2) node, remove the MGMTDB resource as follows:

```
$ $ORACLE_HOME/bin/srvctl remove mgmtldb -f
```

8. As `grid` user, run DBCA in the silent mode from the 12.2.0.1 Grid home and create the Management Database container database (CDB) as follows:

```
$ $ORACLE_HOME/bin/dbca -silent -createDatabase -createAsContainerDatabase
true
-templateName MGMTSeed_Database.dbc -sid -MGMTDB -gdbName _mgmtdb
-storageType ASM -diskGroupName ASM_DG_NAME
-datafileJarLocation /u01/app/12.2.0/grid/assistants/dbca/templates
-characterSet AL32UTF8 -autoGeneratePasswords -skipUserTemplateCheck
```

9. Configure the Management Database by running the Configuration Assistant from the location `$ORACLE_HOME/bin/mgmtca -local`.

10. As `grid` user, run the `post_gimr_ugdg.pl` script from 19c Grid home:

```
$ $ORACLE_HOME/crs/install/post_gimr_ugdg.pl -downgrade -clusterType SC -
destHome /u01/app/19.0.0/grid
-lowerVersion 12.2.0.1.0 -oraBase /u01/app/grid2
```

Where:

SC is the type of the cluster as Oracle Standalone Cluster. The value for `-clusterType` can be SC for Oracle Standalone Cluster, DSC for Oracle Domain Services Cluster, or MC for Oracle Member Cluster.

`/u01/app/19.0.0/grid` is the Oracle home for Oracle Grid Infrastructure 19c.

12.2.0.1.0 is the version of Oracle Grid Infrastructure to which you are downgrading.

`/u01/app/grid2` is the Oracle base for Oracle Grid Infrastructure 19c

Downgrading Oracle Domain Services Cluster to 12c Release 2 (12.2)

Use this procedure to downgrade Oracle Domain Services Cluster to Oracle Grid Infrastructure 12c Release 2 (12.2) after a successful upgrade.

1. As `grid` user, downgrade the Management Database to Oracle Grid Infrastructure 12c Release 2 (12.2):
 - a. Manually copy the most recent time zone files from 19c Grid home to 12c Release 2 (12.2) Grid home, where `timezlrgr_number` is the name of the most recent `timzlrgr` file and `timezone_number` is the name of the most recent `timezone` file:

```
$ cp $ORACLE_HOME/oracore/zoneinfo/timezlrgr_number.dat /u01/app/12.2.0/
grid/oracore/zoneinfo/timezlrgr_number.dat
$ cp $ORACLE_HOME/oracore/zoneinfo/timezone_number.dat /u01/app/12.2.0/
grid/oracore/zoneinfo/timezone_number.dat
```

- b. Downgrade application schema using the following command syntax from 19c Grid home:

```
$ $ORACLE_HOME/bin/mgmtua downgrade -local -oldOracleHome /u01/app/
12.2.0/grid -skipSystemSchemaDowngrade
```

c. Disable and stop MGMTDB resource from 19c Grid home:

```
$ cd $ORACLE_HOME/bin
$ ./srvctl disable mgmtdb
$ ./srvctl stop mgmtdb
```

d. Downgrade system schema using the following procedure:

i. Set ORACLE_SID environment variable for 19c Grid home:

```
$ export ORACLE_SID=-MGMTDB
$ cd $ORACLE_HOME/bin
```

ii. Start CDB and all PDBs in downgrade mode:

```
$ ./sqlplus / as sysdba
SQL> startup downgrade
SQL> alter pluggable database all open downgrade;
SQL> exit
```

iii. Downgrade 19c Management Database using the following command syntax, where /u01/app/grid2 is the Oracle base for Oracle Grid Infrastructure 19c:

```
$ $ORACLE_HOME/perl/bin/perl $ORACLE_HOME/rdbms/admin/catcon.pl -d
/u01/app/grid2 -e -l /u01/app/grid2/cfgtoollogs/mgmtua -b
mgmtdowngrade -r
$ORACLE_HOME/rdbms/admin/catdwgrd.sql
```

iv. Set ORACLE_HOME and ORACLE_SID environment variables for 12c Release 2 (12.2) Grid home:

```
$ export ORACLE_HOME=/u01/app/12.2.0/grid/
$ export ORACLE_SID=-MGMTDB
$ cd $ORACLE_HOME/bin
```

v. Start CDB and all PDBs in upgrade mode:

```
$ ./sqlplus / as sysdba
SQL> shutdown immediate
SQL> startup upgrade
SQL> alter pluggable database all open upgrade;
SQL> exit
```

vi. Run catrelod script for 12c Release 2 (12.2) Management Database using the following command syntax, where /u01/app/grid is the Oracle base for Oracle Grid Infrastructure 12c Release 2 (12.2):

```
$ $ORACLE_HOME/perl/bin/perl $ORACLE_HOME/rdbms/admin/catcon.pl -d
/u01/app/grid -e -l /u01/app/grid/cfgtoollogs/mgmtua -b
mgmtdowngrade
$ORACLE_HOME/rdbms/admin/catrelod.sql
```


- vii. Recompile all invalid objects after downgrade using the following command syntax from 12c Release 2 (12.2) Grid home:

```
$ $ORACLE_HOME/perl/bin/perl $ORACLE_HOME/rdbms/admin/catcon.pl -d
/u01/app/grid -e -l /u01/app/grid/cfgtoollogs/mgmtua -b
mgmtdowngrade
$ORACLE_HOME/rdbms/admin/utlrlp.sql
```

- e. Manually stop the Management Database:

```
$ ./sqlplus / as sysdba
SQL> shutdown immediate
SQL> exit
```

2. As **root** user, use the following command syntax `rootcrs.sh -downgrade` from 19c 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. As **root** user, downgrade the last node after you downgrade all other nodes:

```
# $ORACLE_HOME/crs/install/rootcrs.sh -downgrade
```

4. As **grid** user, remove Oracle Grid Infrastructure 19c 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/19.0.0/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/19.0.0/grid
"CLUSTER_NODES=node1,node2,node3"
-doNotUpdateNodeList
```

Add the flag `-cfs` if the Grid home is a shared home.

5. As **root** user, start the 12c Release 2 (12.2) Oracle Clusterware stack on all nodes.

```
# crsctl start crs
```

6. As **grid** user, set Oracle Grid Infrastructure 12c Release 2 (12.2) 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=/u01/app/12.2.0/grid
"CLUSTER_NODES=node1,node2,node3"
```

7. As `grid` user, downgrade CHA models from any node where the Grid Infrastructure stack is running from 12c Release 2 (12.2) Grid home and Management Database and `ochad` are up:

```
$ ./srvctl stop cha
$ ./chactl import model -file /u01/app/12.2.0/grid/cha/model/os_gold.svm -
name DEFAULT_CLUSTER
$ ./chactl import model -file /u01/app/12.2.0/grid/cha/model/db_gold.svm -
name DEFAULT_DB
$ ./srvctl start cha
```

In the example above, `DEFAULT_CLUSTER` and `DEFAULT_DB` are function names that you must pass as values.

Related Topics

- [Downgrading to Oracle Grid Infrastructure 12c Release 1 \(12.1\)](#)

Downgrading Oracle Member Cluster to 12c Release 2 (12.2)

Use this procedure to downgrade Oracle Member Cluster to Oracle Grid Infrastructure 12c Release 2 (12.2) after a successful upgrade.

1. As `grid` user, use the command syntax `mgmtua downgrade` from 19c Grid home to downgrade Oracle Member Cluster where `oldOracleHome` is 12c Release 2 (12.2) Grid home and `version` is the five digit release number:

```
$ $ORACLE_HOME/bin/mgmtua downgrade -member -oldOracleHome /u01/app/12.2.0/
grid -toVersion version
```

2. As `root` user, use the command syntax `rootcrs.sh -downgrade` from 19c 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. As `root` user, downgrade the last node after you downgrade all other nodes:

```
# $ORACLE_HOME/crs/install/rootcrs.sh -downgrade
```

4. As `grid` user, remove Oracle Grid Infrastructure 19c 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/19.0.0/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/19.0.0/grid
"CLUSTER_NODES=node1,node2,node3"
-doNotUpdateNodeList
```

Add the flag `-cfs` if the Grid home is a shared home.

5. As `root` user, start the 12c Release 2 (12.2) Oracle Clusterware stack on all nodes.

```
# crsctl start crs
```

 Note:

You must start Oracle Clusterware on last downgraded node first, and then on other nodes.

6. As `grid` user, set Oracle Grid Infrastructure 12c Release 2 (12.2) 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=/u01/app/12.2.0/grid
"CLUSTER_NODES=node1,node2,node3"
```

7. As `grid` user, downgrade CHA models from any node where the Grid Infrastructure stack is running from 12c Release 2 (12.2) Grid home and Management Database and `ochad` are up:

```
$ ./srvctl stop cha
$ ./chactl import model -file /u01/app/12.2.0/grid/cha/model/os_gold.svm -
name DEFAULT_CLUSTER
$ ./chactl import model -file /u01/app/12.2.0/grid/cha/model/db_gold.svm -
name DEFAULT_DB
$ ./srvctl start cha
```

In the example above, `DEFAULT_CLUSTER` and `DEFAULT_DB` are function names that you must pass as values.

Related Topics

- [Downgrading Oracle Domain Services Cluster to 12c Release 2 \(12.2\)](#)
Use this procedure to downgrade Oracle Domain Services Cluster to Oracle Grid Infrastructure 12c Release 2 (12.2) after a successful upgrade.
- [Downgrading to Oracle Grid Infrastructure 12c Release 1 \(12.1\)](#)
Use this procedure to downgrade to Oracle Grid Infrastructure 12c Release 1 (12.1).

Downgrading Oracle Grid Infrastructure to 12c Release 2 (12.2) 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/19.0.0/grid/gridSetup.sh -silent -downgrade [-nodes]
nodes_to_be_downgraded
[-oldHome] previous_release_grid_home_to_downgrade_to
[-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.

Related Topics

- [Downgrading Oracle Domain Services Cluster to 12c Release 2 \(12.2\)](#)

Downgrading to Oracle Grid Infrastructure 12c Release 1 (12.1)

Use this procedure to downgrade to Oracle Grid Infrastructure 12c Release 1 (12.1).

1. Delete the Oracle Grid Infrastructure 19c Management Database:

```
$ dbca -silent -deleteDatabase -sourceDB -MGMTDB
```

2. Use the command syntax `rootcrs.sh -downgrade` to downgrade Oracle Grid Infrastructure on all nodes, in any sequence. For example:

```
# /u01/app/19.0.0/grid/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:

```
# /u01/app/19.0.0/grid/crs/install/rootcrs.sh -downgrade
```

4. Remove Oracle Grid Infrastructure 19c 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/19.0.0/grid` is the location of the new (upgraded) Grid home:

```
$ cd /u01/app/19.0.0/grid/oui/bin
./runInstaller -nowait -waitforcompletion -ignoreSysPrereqs
-updateNodeList -silent CRS=false
ORACLE_HOME=/u01/app/19.0.0/grid
"CLUSTER_NODES=node1,node2,node3"
-doNotUpdateNodeList
```

Add the flag `-cfs` if the Grid home is a shared home.

5. Set Oracle Grid Infrastructure 12c Release 1 (12.1) 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 /u01/app/12.1.0/grid/oui/bin
$ ./runInstaller -nowait -waitforcompletion -ignoreSysPrereqs
-updateNodeList -silent CRS=true
ORACLE_HOME=/u01/app/12.1.0/grid
"CLUSTER_NODES=node1,node2,node3"
```

6. Start the 12.1 Oracle Clusterware stack on all nodes.

```
# crsctl start crs
```

7. On any node, remove the MGMTDB resource as follows:

```
$ 121_Grid_home/bin/srvctl remove mgmtadb
```

8. If you are downgrading to Oracle Grid Infrastructure 12c Release 1 (12.1.0.2), run the following commands to configure the Grid Infrastructure Management Database:

- a. Run DBCA in the silent mode from the 12.1.0.2 Oracle home and create the Management Database container database (CDB) as follows:

```
12102_Grid_home/bin/dbca -silent -createDatabase -
createAsContainerDatabase true
-templateName MGMTSeed_Database.dbc -sid -MGMTDB -gdbName _mgmtadb
-storageType ASM -diskGroupName ASM_DG_NAME
-datafileJarLocation 12102_Grid_home/assistants/dbca/templates
-characterSet AL32UTF8 -autoGeneratePasswords -skipUserTemplateCheck
```

- b. Run DBCA in the silent mode from the 12.1.0.2 Oracle home and create the Management Database pluggable database (PDB) as follows:

```
12102_Grid_home/bin/dbca -silent -createPluggableDatabase -sourceDB
-MGMTDB -pdbName cluster_name -createPDBFrom RMANBACKUP
-PDBBackUpfile 12102_Grid_home/assistants/dbca/templates/
mgmtseed_pdb.dfb
-PDBMetadataFile 12102_Grid_home/assistants/dbca/templates/
mgmtseed_pdb.xml
-createAsClone true -internalSkipGIHomeCheck
```

9. If you are downgrading to Oracle Grid Infrastructure 12c Release 1 (12.1.0.1), run DBCA in the silent mode from the 12.1.0.1 Oracle home and create the Management Database as follows:

```
12101_Grid_home/bin/dbca -silent -createDatabase
-templateName MGMTSeed_Database.dbc -sid -MGMTDB -gdbName _mgmtdb
-storageType ASM -diskGroupName ASM_DG_NAME
-datafileJarLocation 12101_Grid_home/assistants/dbca/templates
-characterSet AL32UTF8 -autoGeneratePasswords
```

10. Configure the Management Database by running the Configuration Assistant from the location `121_Grid_home/bin/mgmtca`.

Downgrading to Oracle Grid Infrastructure 11g Release 2 (11.2)

Use this procedure to downgrade to Oracle Grid Infrastructure 11g Release 2 (11.2) .

1. Delete the Oracle Grid Infrastructure 19c Management Database:

```
$ dbca -silent -deleteDatabase -sourceDB -MGMTDB
```

2. Use the command syntax `Grid_home/crs/install/rootcrs.sh -downgrade` to stop the Oracle Grid Infrastructure 19c resources, and to shut down the stack. 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:

```
# /u01/app/19.0.0/grid/crs/install/rootcrs.sh -downgrade
```

4. Follow these steps to remove Oracle Grid Infrastructure 19c 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/19.0.0/grid` is the location of the new (upgraded) Grid home:

```
$ cd /u01/app/19.0.0/grid/oui/bin
$ ./runInstaller -nowait -waitforcompletion -ignoreSysPrereqs -
updateNodeList
```

```
-silent CRS=false ORACLE_HOME=/u01/app/19.0.0/grid
"CLUSTER_NODES=node1,node2,node3" -doNotUpdateNodeList
```

Add the `-cfs` option if the Grid home is a shared home.

5. Follow these steps to set the Oracle Grid Infrastructure 11g Release 2 (11.2) 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 the `ORACLE_HOME` is the location of the home directory from the earlier Oracle Clusterware installation.

```
$ cd /u01/app/11.2.0/grid/oui/bin
$ ./runInstaller -nowait -waitforcompletion -ignoreSysPrereqs -
updateNodeList
-silent CRS=true ORACLE_HOME=/u01/app/11.2.0/grid
```

Add the `-cfs` option if the Grid home is a shared home.

6. Start the Oracle Clusterware stack manually from the earlier release Oracle Clusterware home using the command `crsctl start crs`. For example, where the earlier release home is `/u01/app/11.2.0/grid`, use the following command on each node:

```
/u01/app/11.2.0/grid/bin/crsctl start crs
```

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 required to downgrade the cluster 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 node that you are downgrading.

```
crsctl stop crs
```

2. If you are upgrading from 11.2.0.4, then apply the latest available patches on all nodes in the cluster. If the pre-upgrade version is 12.1.0.2 or later, then patch is not required.
 - a. On all other nodes except the first node, where the earlier release Grid Infrastructure stack is running, apply the latest patch using the `opatchauto` procedure.
 - b. On the first node where the earlier release Grid Infrastructure stack is stopped, apply the latest patch using the `opatch apply` procedure.

For the list of latest available patches, see My Oracle Support at the following link:

<https://support.oracle.com/>

- i. Unlock the Grid Infrastructure home from the earlier release:

```
rootcrs.pl -unlock -crshome pre-upgrade-grid-home
```

pre-upgrade-grid-home is the previous release Grid home.

- ii. Apply the patch:

```
opatch apply -local -oh pre-upgrade-grid-home
```

- iii. Relock the Grid home from the earlier release:

```
rootcrs.pl -lock
```

- c. From any other node where the Grid Infrastructure stack from the earlier release is running, unset the Oracle ASM rolling migration mode as explained in step 2.

3. Run the following command on all cluster nodes on which Oracle Grid Infrastructure is running, but not on the node that you are downgrading.

```
clscfg -nodedowngrade -h hostname
```

hostname is the name of the node that you are downgrading.

4. From the later release Grid home, run the following command on all cluster nodes, which are successfully upgraded, 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 all the nodes from the earlier release Grid home.

```
crsctl start crs
```

Note:

You can downgrade the successfully upgraded cluster nodes in any sequence.

6. On the last node where the Grid Infrastructure stack from the earlier release is running, unset the Oracle ASM rolling migration mode as follows:
 - Log in as `grid` user, and run the following command as SYSASM user on the Oracle ASM instance:

```
SQL> ALTER SYSTEM STOP ROLLING MIGRATION;
```

Related Topics

- [Downgrading Oracle Grid Infrastructure to 12c Release 2 \(12.2\) when Upgrade Fails](#)

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 Upgrades on Remote Nodes](#)
Review this information to continue incomplete upgrade on remote nodes.
- [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 /u01/app/19.0.0/grid
[root@node1]# ./rootupgrade.sh
```

3. 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]$ /u01/app/19.0.0/grid/gridSetup.sh -executeConfigTools -
responseFile
/u01/app/19.0.0/grid/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 Upgrades on Remote Nodes

Review this information to continue incomplete upgrade on remote nodes.

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 root automation is being used, click **Retry** on the OUI instance on the first node.
3. If root automation is not being used, log into the affected node as `root`. Change directory to the Grid home, and run the `rootupgrade.sh` script on that node. For example:

```
[root@node6]# cd /u01/app/19.0.0/grid
[root@node6]# ./rootupgrade.sh
```

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 /u01/app/19.0.0/grid
[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]$ /u01/app/19.0.0/grid/gridSetup.sh -executeConfigTools -
responseFile /u01/app/19.0.0/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 /u01/app/19.0.0/grid
[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 19c as described in this chapter.

1. As `root` user, log in to the first node, and run the 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.

For example:

```
rootcrs.sh -converttoextended -first -sites newyork,newjersey,conn -site  
newyork
```

2. As `root` user, on all other nodes, run the following command:

```
rootcrs.sh -converttoextended -site node_site
```

node_site is the node containing the site.

For example:

```
rootcrs.sh -converttoextended -site newjersey
```

3. Delete the default site after the associated nodes and storage are migrated.

```
crsctl delete cluster site site_name
```

For example:

```
[root@node4]#crsctl delete cluster site mycluster
```

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](#)
Passwordless SSH configuration is a mandatory installation requirement. 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 Shell Limits and System Configuration Parameters for IBM AIX](#)
These topics explain how to configure shell limits and system configuration parameters manually if you cannot complete them using the fixup scripts.
- [Setting UDP and TCP Kernel Parameters Manually](#)
If you do not use a Fixup script or CVU to set ephemeral ports, then use NDD to ensure that the kernel TCP/IP ephemeral port range is broad enough to provide enough ephemeral ports for the anticipated server workload.

Configuring SSH Manually on All Cluster Nodes

Passwordless SSH configuration is a mandatory installation requirement. 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.

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.



Note:

The supported version of SSH for Linux distributions is OpenSSH.

- [Checking Existing SSH Configuration on the System](#)
To determine if SSH is running, enter the following command.
- [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.

Checking Existing SSH Configuration on the System

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. In the home directory of the installation software owner (`grid`, `oracle`), use the command `ls -al` to ensure that the `.ssh` directory is owned and writable only by the user.

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.

Configuring SSH on Cluster Nodes

You must configure SSH separately for each Oracle software installation owner that you intend to use for installation.

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.

To configure SSH, complete the following:

- [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 SFTP (Secure FTP) to copy the `authorized_keys` file to the `oracle` user `.ssh` directory on a remote node. The following example is on a node called `node2`, with the Oracle Grid Infrastructure owner `grid`, where the `grid` user path is `/home/grid`:

```
[grid@node1 .ssh]$ sftp 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]$ sftp 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 `sftp` 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 (without domain name and in lowercase), are the public host names 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 `sftp` 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 Shell Limits and System Configuration Parameters for IBM AIX

These topics explain how to configure shell limits and system configuration parameters manually if you cannot complete them using the fixup scripts.



Note:

The parameter and shell limit values shown in this section are recommended values only. For production database systems, Oracle recommends that you tune these values to optimize the performance of the system. See your operating system documentation for more information about tuning kernel parameters.

- [Configuring System Configuration Parameters for IBM AIX](#)
Use these instructions to verify and set the system configuration values manually if they are different from the minimum recommended value.
- [Configuring Shell Limits on IBM AIX](#)
For IBM AIX on POWER Systems (64-bit), it is the `ulimit` settings that determine process memory related resource limits. Use these instructions to display and change the shell limit values if they are different from the recommended value.
- [Checking Asynchronous Input Output Processes](#)
On IBM AIX 6 and IBM AIX 7, the Asynchronous Input Output (AIO) device drivers are enabled by default. For both IBM AIX 6 and IBM AIX 7, increase the number of `aio_server` processes from the default value. The recommended value for `aio_maxreqs` is 64k (65536). Confirm this value for both IBM AIX 6 and IBM AIX 7.

Configuring System Configuration Parameters for IBM AIX

Use these instructions to verify and set the system configuration values manually if they are different from the minimum recommended value.

If you cannot use the Fixup scripts, then verify that the kernel parameters shown in the following table are set to values greater than or equal to the minimum value shown. If the current value for any parameter is greater than the value listed in this table, then the Fixup scripts do not change the value of that parameter.

Table A-1 Minimum System Configuration Parameter Values

Parameter	Minimum Value
<code>maxuproc</code> s	16384
<code>ncargs</code>	128

To verify and set the values manually:

1. To verify that the maximum number of processes allowed per user is set to 16384 or greater:



Note:

For production systems, this value should be at least 128 plus the sum of the `PROCESSES` and `PARALLEL_MAX_SERVERS` initialization parameters for each database running on the system.

- a. Enter the following command:


```
# smit chgsys
```
 - b. Verify that the value shown for `Maximum number of PROCESSES allowed per user` is greater than or equal to 16384.
If necessary, edit the existing value.
 - c. When you have finished making changes, press `F10` to exit.
2. To verify that long commands can be run from the shell, perform the following steps:

 Note:

Oracle recommends that you set the `ncargs` system attribute to a value greater than or equal to 128. The `ncargs` attribute determines the maximum number of values that can be passed as command line arguments.

- a. Enter the following command:


```
# smit chgsys
```
- b. Verify that the value shown for `ARG/ENV list size in 4K byte blocks` is greater than or equal to 128.
If necessary, edit the existing value.
- c. When you have finished making changes, press `F10` to exit.

Configuring Shell Limits on IBM AIX

For IBM AIX on POWER Systems (64-bit), it is the `ulimit` settings that determine process memory related resource limits. Use these instructions to display and change the shell limit values if they are different from the recommended value.

Set shell limits for the Oracle Grid Infrastructure installation owner and for `root`. Verify that `unlimited` is set for both accounts either by using the `smit` utility or by editing the `/etc/security/limits` file. The root user requires these settings because the `crs` daemon (`crsd`) runs as `root`.

Verify that the shell limits displayed in the following table are set to the values shown:

Table A-2 Recommended Shell Limit Values

Shell Limit (As Shown in smit)	Recommended Value
Soft File Descriptors	at least 1024
Hard File Descriptors	at least 65536
Number of processes (Soft)	at least 2047
Number of processes (Hard)	at least 16384
Soft STACK size	at least 10240 KB
Hard STACK size	at least 10240 KB; at most 32768 KB
Soft FILE size	unlimited

Table A-2 (Cont.) Recommended Shell Limit Values

Shell Limit (As Shown in smit)	Recommended Value
Soft CPU time	unlimited
Soft DATA segment	unlimited
Soft Real Memory size	unlimited



Note:

This is the default value.

To display the current value specified for these shell limits, and to change them if necessary

1. Enter the following command:

```
# smit chuser
```

2. In the **User NAME** field, enter the user name of the Oracle software owner, for example `oracle`.
3. Scroll down the list and verify the values shown for the limits listed in the previous table.
If necessary, edit the existing value. To edit the values, you can use the `smit` utility. However, to set the value of Soft Real Memory size, you must edit the file `/etc/security/limits`. If you have permissions to run the `smit` utility, then you automatically have the permissions to edit the limits file.
4. When you have finished making changes, press `F10` to exit.

Checking Asynchronous Input Output Processes

On IBM AIX 6 and IBM AIX 7, the Asynchronous Input Output (AIO) device drivers are enabled by default. For both IBM AIX 6 and IBM AIX 7, increase the number of `aio_server` processes from the default value. The recommended value for `aio_maxreqs` is 64k (65536). Confirm this value for both IBM AIX 6 and IBM AIX 7.

Confirm the `aio_maxreqs` value using the following procedure:

```
# ioo -o aio_maxreqs
aio_maxreqs = 65536
```

When performing an asynchronous I/O to a file system, each asynchronous I/O operation is tied to an asynchronous I/O server. Thus, the number of asynchronous I/O servers limits the number of concurrent asynchronous I/O operations in the system.

The initial number of servers that are started during a system restart is determined by the `aio_minservers` parameter. As concurrent asynchronous I/O operations occur, additional asynchronous I/O servers are started, up to a maximum of the value set in the `aio_maxservers` parameter.

In general, to set the number of asynchronous I/O servers, complete the following procedure:

1. Adjust the initial value of `aio_maxservers` to 10 times the number of logical disks divided by the number of CPUs that are to be used concurrently but no more than 80.
2. Monitor the performance effects on the system during periods of high I/O activity. If all AIO server processes are started, then increase the `aio_maxservers` value. Also, continue to monitor the system performance during peak I/O activity to determine if there was a benefit from the additional AIO servers. Too many asynchronous I/O servers increase memory and processor overload of additional processes, but this disadvantage is small.

See your operating system vendor documentation for information about tuning AIO parameters.

To monitor the number of AIO server processes that have started, enter the following:

```
# ps -ek|grep -v grep|grep -v posix_aioserver|grep -c aioserver
```

Setting UDP and TCP Kernel Parameters Manually

If you do not use a Fixup script or CVU to set ephemeral ports, then use NDD to ensure that the kernel TCP/IP ephemeral port range is broad enough to provide enough ephemeral ports for the anticipated server workload.

Ensure that the lower range is set to at least 9000 or higher, to avoid Well Known ports, and to avoid ports in the Registered Ports range commonly used by Oracle and other server ports. Set the port range high enough to avoid reserved ports for any applications you may intend to use. If the lower value of the range you have is greater than 9000, and the range is large enough for your anticipated workload, then you can ignore OUI warnings regarding the ephemeral port range.

1. Use the following commands to check your current range for ephemeral ports:

```
# /usr/sbin/no -a | fgrep ephemeral
tcp_ephemeral_low = 32768
tcp_ephemeral_high = 65500
udp_ephemeral_low = 32768
udp_ephemeral_high = 65500
```

In the preceding example, the TCP and UDP ephemeral ports are set to the default range (32768-65536).

2. If you expect your workload to require a high number of ephemeral ports, such as high node counts or heavy use of Parallel Query, then update the UDP and TCP ephemeral port range to a broader range. For example:

```
# /usr/sbin/no -p -o tcp_ephemeral_low=9000 -o tcp_ephemeral_high=65500
# /usr/sbin/no -p -o udp_ephemeral_low=9000 -o udp_ephemeral_high=65500
```

B

Response Files

Review the following topics to install and configure Oracle Database, Oracle Grid Infrastructure, and other Oracle products using response files.

- [How Response Files Work](#)
Response files can assist you with installing an Oracle product multiple times on multiple computers.
- [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](#)
Review this information to use response files.
- [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 Configuration Assistants Using Response Files](#)
You can run configuration assistants in response file or silent mode to configure and start Oracle software after it is installed on the system. To run configuration assistants in response file or silent mode, you must copy and edit a 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.
- [Postinstallation Configuration Using the ConfigToolAllCommands Script](#)
You can create and run a response file configuration after installing Oracle software. The `configToolAllCommands` script requires users to create a second response file, of a different format than the one used for installing the product.

How Response Files Work

Response files can assist you with installing an Oracle product multiple times on multiple computers.

When you start Oracle Universal Installer (OUI), you can use a response file to automate the installation and configuration of Oracle software, either fully or partially. OUI uses the values contained in the response file to provide answers to some or all installation prompts.

Typically, the installer runs in interactive mode, which means that it prompts you to provide information in graphical user interface (GUI) screens. When you use response files to provide this information, you run the installer from a command prompt using either of the following modes:

- Silent mode

If you include responses for all of the prompts in the response file and specify the `-silent` option when starting the installer, then it runs in silent mode. During a silent mode installation, the installer does not display any screens. Instead, it displays progress information in the terminal that you used to start it.

- **Response file mode**

If you include responses for some or all of the prompts in the response file and omit the `-silent` option, then the installer runs in response file mode. During a response file mode installation, the installer displays all the screens, screens for which you specify information in the response file, and also screens for which you did not specify the required information in the response file.

You define the settings for a silent or response file installation by entering values for the variables listed in the response file. For example, to specify the Oracle home name, provide the Oracle home path for the `ORACLE_HOME` environment variable:

```
ORACLE_HOME=/u01/app/oracle/product/19.0.0/dbhome_1
```

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 <code>at</code>. • 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

Review this information to use response files.

Use the following general steps 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. This method is useful for Advanced or software-only 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. This method 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.

To record a 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/19.0.0/grid/gridSetup.sh -help
```

The help information appears in a window after some time.

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/19.0.0/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 Configuration Assistants Using Response Files

You can run configuration assistants in response file or silent mode to configure and start Oracle software after it is installed on the system. To run configuration assistants in response file or silent mode, you must copy and edit a response file template.

 Note:

If you copied the software to a hard disk, then the response file template is located in the `/response` directory.

- [Running Oracle DBCA Using Response Files](#)
You can run Oracle Database Configuration Assistant (Oracle DBCA) in response file mode to configure and start an Oracle database on the system.
- [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.

Running Oracle DBCA Using Response Files

You can run Oracle Database Configuration Assistant (Oracle DBCA) in response file mode to configure and start an Oracle database on the system.

To run Oracle DBCA in response file mode, you must copy and edit a response file template. Oracle provides a response file template named `dbca.rsp` in the `ORACLE_HOME/assistants/dbca` directory. To run Oracle DBCA in response file mode, you must use the `-responseFile` flag in combination with the `-silent` flag. You must also use a graphical display and set the `DISPLAY` environment variable.

To run Oracle DBCA in response file mode:

1. Copy the `dbca.rsp` response file template from the response file directory to a directory on your system:

```
$ cp /directory_path/assistants/dbca/dbca.rsp local_directory
```

In this example, `directory_path` is the path of the directory where you have copied the installation binaries.

As an alternative to editing the response file template, you can also create a database by specifying all required information as command line options when you run Oracle DBCA. For information about the list of options supported, enter the following command:

```
$ $ORACLE_HOME/bin/dbca -help
```

2. Open the response file in a text editor:

```
$ vi /local_dir/dbca.rsp
```

3. Follow the instructions in the file to edit the file.

 Note:

Oracle DBCA 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. To run Oracle DBCA in response file mode, set the `DISPLAY` environment variable.
6. Use the following command syntax to run Oracle DBCA in silent or response file mode using a response file:

```
$ORACLE_HOME/bin/dbca [-silent] -createDatabase -responseFile /local_dir/  
dbca.rsp
```

In this example:

- `-silent` option indicates that Oracle DBCA runs in silent mode.
- `local_dir` is the full path of the directory where you copied the `dbca.rsp` response file template.

During configuration, Oracle DBCA displays a window that contains the status messages and a progress bar.

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](#)
Starting with Oracle Database 12c release 2 (12.2), 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

Starting with Oracle Database 12c release 2 (12.2), 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.ASMNMPPassword=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, 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 `ASMSNMP` variable. If you are not using Oracle ASM, then leave the value for this password variable blank.
- If you create a multitenant container database (CDB) with one or more pluggable databases (PDBs), then you must specify a password for the `PDBADMIN` 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`. For example:

For Oracle Database (`oracle` user)

```
oracle.install.asm.SYSASMPassword=password
oracle.install.config.emAdminPassword=password
```

For Oracle Grid Infrastructure (`grid` user)

```
grid.install.asm.SYSASMPassword=password
grid.install.config.emAdminPassword=password
```

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/19.0.0/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/inventory/response/`:

```
$ ./runInstaller -executeConfigTools -responseFile $ORACLE_HOME/inventory/
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/19.0.0/grid/  
response/grid_2016-01-09_01-03-36PM.rsp [-silent]
```

For Oracle Database:

```
$ ./runInstaller -executeConfigTools -responseFile /u01/app/oracle/product/  
19.0.0/dbhome_1/inventory/response/db_2016-01-09_01-03-36PM.rsp [-silent]
```

Postinstallation Configuration Using the ConfigToolAllCommands Script

You can create and run a response file configuration after installing Oracle software. The `configToolAllCommands` script requires users to create a second response file, of a different format than the one used for installing the product.

Starting with Oracle Database 12c Release 2 (12.2), the `configToolAllCommands` script is deprecated and may be desupported in a future release.

- [About the Postinstallation Configuration File](#)
When you run a silent or response file installation, you provide information about your servers in a response file that you otherwise provide manually during a graphical user interface installation.
- [Creating a Password Response File](#)
You can create a password response file and use it with configuration assistants to perform silent installation.
- [Running Postinstallation Configuration Using a Password Response File](#)
Complete this procedure to run configuration assistants with the `configToolAllCommands` script.

Related Topics

- [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.

About the Postinstallation Configuration File

When you run a silent or response file installation, you provide information about your servers in a response file that you otherwise provide manually during a graphical user interface installation.

However, the response file does not contain passwords for user accounts that configuration assistants require after software installation is complete. The configuration assistants are started with a script called `configToolAllCommands`. You can run this script in response file mode by using a password response file. The script uses the passwords to run the configuration tools in succession to complete configuration.

If you keep the password file to use for clone installations, then Oracle strongly recommends that you store the password file in a secure location. In addition, if you have to stop an installation to fix an error, then you can run the configuration assistants using `configToolAllCommands` and a password response file.

The `configToolAllCommands` password response file has the following syntax options:

- `oracle.crs` for Oracle Grid Infrastructure components or `oracle.server` for Oracle Database components that the configuration assistants configure
- `variable_name` is the name of the configuration file variable
- `value` is the desired value to use for configuration.

The command syntax is as follows:

```
internal_component_name|variable_name=value
```

For example:

```
oracle.crs|S_ASMPASSWORD=PassWord
```

The database configuration assistants require the SYS, SYSTEM, and DBSNMP passwords for use with Oracle DBCA. You may need to specify the following additional passwords, depending on your system configuration:

- If the database is using Oracle Automatic Storage Management (Oracle ASM) for storage, then you must specify a password for the `S_ASMSNMPPASSWORD` variable. If you are not using Oracle ASM, then leave the value for this password variable blank.
- If you create a multitenant container database (CDB) with one or more pluggable databases (PDBs), then you must specify a password for the `S_PDBADMINPASSWORD` variable. If you are not using Oracle ASM, then leave the value for this password variable blank.

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.

Creating a Password Response File

You can create a password response file and use it with configuration assistants to perform silent installation.

Perform the following steps to create a password response file:

1. Create a response file that has a name of the format `filename.properties`, for example:

```
$ touch pwdrsp.properties
```

2. Open the file with a text editor, and cut and paste the sample password file contents, as shown in the examples, modifying as needed.
3. Change permissions to secure the password response file. For example:

```
$ ls -al pwdrsp.properties
-rw----- 1 oracle oinstall 0 Apr 30 17:30 pwdrsp.properties
```

Example B-4 Password response file for Oracle Grid Infrastructure (grid user)

```
grid.crs|S_ASMPASSWORD=password
grid.crs|S_OMSPASSWORD=password
grid.crs|S_BMCPASSWORD=password
grid.crs|S_ASMMONITORPASSWORD=password
```

If you do not have a BMC card, or you do not want to enable IPMI, then leave the `S_BMCPASSWORD` input field blank.

Example B-5 Password response file for Oracle Grid Infrastructure for a Standalone Server (oracle user)

```
oracle.crs|S_ASMPASSWORD=password
oracle.crs|S_OMSPASSWORD=password
oracle.crs|S_ASMMONITORPASSWORD=password
```

Example B-6 Password response file for Oracle Database (oracle user)

This example provides a template for a password response file to use with the database configuration assistants.

```
oracle.server|S_SYSPASSWORD=password
oracle.server|S_SYSPASSWORD=password
oracle.server|S_EMADMINPASSWORD=password
oracle.server|S_DBSNMPPASSWORD=password
oracle.server|S_ASMSNMPPASSWORD=password
oracle.server|S_PDBADMINPASSWORD=password
```

If you do not want to enable Oracle Enterprise Manager for management, then leave those password fields blank.

Running Postinstallation Configuration Using a Password Response File

Complete this procedure to run configuration assistants with the `configToolAllCommands` script.

1. Create a password response file as described in *Creating a Password File*.
2. Change directory to `$ORACLE_HOME/cfgtoollogs`.
3. Run the configuration script using the following syntax:

```
configToolAllCommands RESPONSE_FILE=/path/name.properties
```

For example:

```
$ ./configToolAllCommands RESPONSE_FILE=/home/oracle/pwdrsp.properties
```

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.
- 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.
- 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 `/etc/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/19.0.0/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/19.0.0/dbhome_1/, owned by oracle:oinstall
/u01/app/grid, owned by grid:oinstall
/u01/app/19.0.0/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/19.0.0/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
<code>/u01/app/ oracle</code>	Oracle Database Oracle base, where the Oracle Database software installation owner name is <code>oracle</code> . The Oracle Database binary home is located underneath the Oracle base path.
<code>/u01/app/grid</code>	Oracle Grid Infrastructure Oracle base, where the Oracle Grid Infrastructure software installation owner name is <code>grid</code> . 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 <code>root</code> 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/sl/product/v/type_[n]`. The following table describes the variables used in this syntax:

Variable	Description
<i>pm</i>	A mount point name.
<i>s</i>	A standard directory name.

Variable	Description
<i>u</i>	The name of the owner of the directory.
<i>v</i>	The version of the software.
<i>type</i>	The type of installation. For example: Database (<i>dbhome</i>), Client (<i>client</i>), or Oracle Grid Infrastructure (<i>grid</i>)
<i>n</i>	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 (<i>dbhome_1</i> , <i>dbhome_2</i>)

For example, the following path is typical for the first installation of Oracle Database on this system:

```
/u01/app/oracle/product/19.0.0/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

Table C-2 (Cont.) Optimal Flexible Architecture Hierarchical File Path Examples

Directory	Description
/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 <code>oracle</code> . 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 <code>grid</code> . 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
/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

Table C-2 (Cont.) Optimal Flexible Architecture Hierarchical File Path Examples

Directory	Description
/u01/app/oracle/ fast_recovery_ a/db_sales	Recovery files for database named "sales"
/u01/app/oracle/ fast_recovery_ a/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/19.0.0/ dbhome_1	Oracle home directory for Oracle Database 1, owned by Oracle Database installation owner account <code>oracle</code>
/u01/app/oracle/ product/19.0.0/ dbhome_2	Oracle home directory for Oracle Database 2, owned by Oracle Database installation owner account <code>oracle</code>
/u01/app/oracle2/ product/19.0.0/ dbhome_2	Oracle home directory for Oracle Database 2, owned by Oracle Database installation owner account <code>oracle2</code>
/u01/app/oracle/ product/19.0.0/ grid	Oracle home directory for Oracle Grid Infrastructure for a standalone server, owned by Oracle Database and Oracle Grid Infrastructure installation owner <code>oracle</code> .

Table C-2 (Cont.) Optimal Flexible Architecture Hierarchical File Path Examples

Directory	Description
/u01/app/19.0.0/ grid	Oracle home directory for Oracle Grid Infrastructure for a cluster (Grid home), owned by user <code>grid</code> before installation, and owned by <code>root</code> after installation.
/u01/app/oracle/ product/19.0.0/ client_1	Oracle home directory for Oracle Database Client 1, owned by Oracle Database installation owner account <code>oracle</code>

Index

Symbols

/tmp directory, [2-2](#)

/usr/sbin/no, [7-17](#)

A

adding Oracle ASM listener, [11-26](#)

aio_maxreqs, [A-8](#)

aioserver, [A-8](#)

apply patches during install

 apply patches during upgrade, [11-20](#)

ASM_DISKSTRING, [7-11](#)

ASMCA

 Used to create disk groups for older Oracle
 Database releases on Oracle ASM,
 [9-10](#)

asmdba groups

 creating, [5-12](#)

asmoper group

 creating, [5-13](#)

ASMSNMP, [1-3](#)

Asynchronous Input Output Processes, [A-8](#)

Automatic Diagnostic Repository (ADR), [C-1](#)

Automatic Storage Management Cluster File
System

 See Oracle ACFS.

B

backupdba group

 creating, [5-14](#)

Bash shell

 default user startup file, [5-24](#)

bash_profile file, [5-24](#)

batch upgrade, [11-14](#)

binaries

 relinking, [9-11](#)

binary files

 supported storage options for, [6-1](#)

BMC

 configuring, [5-30](#)

BMC interface

 preinstallation tasks, [5-28](#)

Bourne shell

 default user startup file, [5-24](#)

C

C shell

 default user startup file, [5-24](#)

central inventory, [1-5](#), [C-5](#)

 See also Oracle inventory directory

 See also OINSTALL directory

changing shell limit values, [A-7](#)

checkdir error, [9-11](#), [11-4](#)

Checking Asynchronous Input Output Processes,
[A-8](#)

checklists, [1-1](#)

client-server configurations, [C-2](#)

clients

 and upgrades, [4-4](#)

 connecting to SCAN, [4-4](#)

 using SCAN, [4-4](#)

cluster configuration

 Oracle Domain Services Cluster, [8-3](#)

 Oracle Extended Clusters, [8-5](#)

 Oracle Member Clusters, [8-4](#)

 Oracle Standalone Clusters, [8-3](#)

cluster file system

 storage option for data files, [6-5](#)

cluster name, [1-3](#)

 requirements for, [1-3](#)

cluster nodes

 private network node interfaces, [1-3](#)

 private node names, [4-4](#)

 public network node names and addresses,
 [1-3](#)

 virtual node names, [1-3](#), [4-3](#)

Cluster Time Synchronization Service, [3-19](#)

CLUSTER_INTERCONNECTS parameter, [4-3](#)

clusterware

 requirements for third party clusterware, [1-3](#)

Clusterware Installation, [8-6](#)

commands

 asmca, [7-27](#)

 asmcmd, [7-10](#)

 crsctl, [11-4](#)

 gridSetup.sh, [8-24](#)

 ipmitool, [5-30](#)

 root.sh, [9-4](#)

 rootcrs.pl

 and deconfig option, [10-14](#)

commands (*continued*)
 rootcrs.sh, [9-11](#)
 rootupgrade.sh, [11-4](#)
 runcluvfy.sh, [8-24](#)
 srvctl, [11-4](#)
 umask, [5-24](#)
 unset, [11-8](#)
 useradd, [5-16](#)
 cron jobs, [1-9](#)
 ctsdd, [3-19](#)
 custom database
 failure groups for Oracle ASM, [7-2](#)
 requirements when using Oracle ASM, [7-2](#)

D

data files
 storage options, [6-5](#)
 supported storage options for, [6-1](#)
 data loss
 minimizing with Oracle ASM, [7-2](#), [7-13](#)
 Database Configuration Assistant
 running in silent mode, [B-7](#)
 databases
 Oracle ASM requirements, [7-2](#)
 DB_RECOVERY_FILE_DEST, [9-5](#)
 DB_RECOVERY_FILE_DEST_SIZE, [9-5](#)
 dba group
 creating, [5-13](#)
 description, [5-9](#)
 SYSDBA privilege, [5-9](#)
 dba groups
 creating, [5-14](#), [5-15](#)
 DBCA
 no longer used for Oracle ASM disk group
 administration, [9-10](#)
 dbca.rsp file, [B-3](#)
 default file mode creation mask
 setting, [5-24](#)
 deinstall, [10-1–10-3](#)
 Oracle Member Cluster, [10-15](#)
 See also removing Oracle software
 deinstall command, [10-2](#)
 deinstallation, [10-2](#)
 examples, [10-6](#)
 Deinstallation tool
 Restriction for Oracle Flex Clusters and -
 lastnode flag, [10-14](#)
 df command, [5-24](#)
 dgdba group
 creating, [5-14](#)
 DHCP
 and GNS, [4-10](#)
 diagnostic data, [C-1](#)
 Direct NFS
 disabling, [7-20](#)

Direct NFS (*continued*)
 enabling, [7-20](#)
 oranfstab file, [7-18](#)
 directory
 creating separate data file directories, [7-15](#)
 disk group
 Oracle ASM, [7-2](#)
 recommendations for Oracle ASM disk
 groups, [7-2](#)
 disk groups
 checking, [7-10](#)
 recommendations for, [7-2](#)
 disk space
 requirements for preconfigured database in
 Oracle ASM, [7-2](#)
 disks
 selecting for use with Oracle ASM, [7-11](#)
 display variable, [1-5](#)
 downgrade, [11-28](#)
 downgrade after failed installation, [11-37](#), [11-44](#),
 [11-47](#)
 downgrade after failed upgrade, [11-37](#), [11-44](#),
 [11-47](#)
 downgrade to 12c Release 2 (12.2), [11-37](#)
 downgrade to 18c, [11-30](#)
 downgrades, [11-30](#), [11-37](#), [11-44](#), [11-47](#)
 downgrades restrictions, [11-30](#)
 downgrading
 Oracle Grid Infrastructure, [11-32](#), [11-35](#),
 [11-39](#), [11-42](#), [11-44](#)
 Oracle Standalone Cluster, [11-30](#), [11-37](#)
 to 12.1, [11-44](#)
 to 12.2, [11-37](#), [11-39](#), [11-42](#)
 to 18c, [11-30](#), [11-32](#), [11-35](#)
 downgrading,
 Downgrade Member Cluster, [11-35](#), [11-42](#)
 Oracle Grid Infrastructure, [11-32](#), [11-35](#),
 [11-39](#), [11-42](#), [11-44](#)
 to 12.2, [11-37](#), [11-39](#), [11-42](#)
 to 18c, [11-30](#), [11-32](#), [11-35](#)
 dry-run upgrade
 check upgrade readiness
 Clusterware upgrade, [11-12](#)
 Oracle Grid Infrastructure preupgrade check,
 [11-12](#)

E

enabling IOCP, [3-12](#)
 enterprise.rsp file, [B-3](#)
 environment
 configuring for Oracle user, [5-23](#)
 environment variables
 ORACLE_BASE, [5-24](#)
 ORACLE_HOME, [5-24](#), [11-8](#)
 ORACLE_SID, [5-24](#), [11-8](#)

environment variables (*continued*)
 removing from shell startup file, [5-24](#)
 SHELL, [5-24](#)
 TEMP and TMPDIR, [5-24](#)
 errors
 X11 forwarding, [5-27](#), [A-5](#)
 errors using Opatch, [11-4](#)
 errors using OPatch, [9-11](#)
 Exadata
 relinking binaries example for, [9-11](#)
 examples
 Oracle ASM failure groups, [7-2](#)
 executeConfigTools, [B-11](#)

F

failed install, [11-49](#)
 failed upgrade, [11-49](#)
 failure group
 characteristics of Oracle ASM failure group,
 [7-2](#), [7-13](#)
 examples of Oracle ASM failure groups, [7-2](#)
 Oracle ASM, [7-2](#)
 fast recovery area, [9-5](#)
 filepath, [C-5](#)
 Grid home
 filepath, [C-5](#)
 fencing
 and IPMI, [5-28](#)
 file mode creation mask
 setting, [5-24](#)
 file system
 storage option for data files, [6-5](#)
 files
 bash_profile, [5-24](#)
 dbca.rsp, [B-3](#)
 editing shell startup file, [5-24](#)
 enterprise.rsp, [B-3](#)
 login, [5-24](#)
 profile, [5-24](#)
 response files, [B-3](#)
 filesets, [3-5](#)
 free disk space, [2-2](#)

G

GIMR, [7-9](#)
 globalization, [1-9](#)
 GNS
 about, [4-11](#)
 configuration example, [4-21](#)
 configuring, [4-10](#)
 GNS client clusters
 and GNS client data file, [4-13](#)
 GNS client data file required for installation,
 [4-12](#)

GNS client clusters (*continued*)
 name resolution for, [4-12](#)
 GNS client data file
 how to create, [4-13](#)
 GNS virtual IP address, [1-3](#)
 grid home
 unlocking, [9-11](#)
 grid infrastructure management repository, [8-3](#)
 Grid Infrastructure Management Repository, [7-9](#)
 global, [7-9](#)
 local, [7-9](#)
 Grid user
 creating, [5-16](#)
 gridSetup script, [8-7](#), [8-12](#), [8-17](#)
 groups
 creating an Oracle Inventory Group, [5-3](#)
 creating the asmdba group, [5-12](#)
 creating the asmoper group, [5-13](#)
 creating the backupdba group, [5-14](#)
 creating the dba group, [5-13](#)
 creating the dgdba group, [5-14](#)
 creating the kmdba group, [5-14](#)
 creating the Oracle ASM group, [5-12](#)
 creating the racdba group, [5-15](#)
 OINSTALL group, [1-2](#)
 OSBACKUPDBA (backupdba), [5-9](#)
 OSDBA (dba), [5-9](#)
 OSDBA group (dba), [5-9](#)
 OSDGDBA (dgdba), [5-9](#)
 OSKMDBA (kmdba), [5-9](#)
 OSOPER (oper), [5-9](#)
 OSOPER group (oper), [5-9](#)

H

hardware requirements
 display, [1-1](#)
 IPMI, [1-1](#)
 local storage for Oracle homes, [1-1](#)
 network, [1-1](#)
 RAM, [1-1](#)
 tmp, [1-1](#)
 highly available IP addresses (HAIP), [4-7](#)
 host names
 legal host names, [1-3](#)
 hugepages, [1-2](#)

I

I/O Completion Ports, [3-12](#)
 image
 install, [8-2](#)
 image-based installation of Oracle Grid
 Infrastructure, [8-7](#), [8-12](#), [8-17](#)
 inaccessible nodes
 upgrading, [11-18](#)

incomplete installations, [11-51](#)
 init.ora
 and SGA permissions, [9-7](#)
 installation
 cloning a Grid infrastructure installation to
 other nodes, [8-25](#)
 response files, [B-3](#)
 preparing, [B-3](#), [B-4](#)
 templates, [B-3](#)
 silent mode, [B-5](#)
 installation planning, [1-1](#)
 installation types
 and Oracle ASM, [7-2](#)
 installer screen
 Specify NFS Locations for ASM Disk Groups,
 [6-7](#)
 Storage Option Information, [6-7](#)
 installer screens
 Cluster Node Information, [4-21](#)
 Grid Plug and Play Information, [4-11](#), [4-20](#),
 [4-21](#)
 Network Interface Usage, [4-19](#)
 Node Selection screen, [11-13](#), [11-15](#)
 installing Oracle Member Cluster, [8-17](#)
 interconnect, [1-3](#)
 interconnects
 single interface, [4-7](#)
 interfaces, [1-3](#)
 requirements for private interconnect, [4-3](#)
 IOCP, [3-12](#)
 IPMI
 addresses not configurable by GNS, [5-29](#)
 preinstallation tasks, [5-28](#)
 IPv4 requirements, [4-2](#), [4-9](#)
 IPv6 requirements, [4-2](#), [4-9](#)

J

JDK requirements, [3-5](#)

K

kernel parameters
 tcp and udp, [A-9](#)
 kernel parameters configuration, [A-5](#)
 kmdba group
 creating, [5-14](#)
 Korn shell
 default user startup file, [5-24](#)

L

legal host names, [1-3](#)
 licensing, [1-9](#)
 login file, [5-24](#)
 LoginGraceTime parameter, [3-15](#)

LVM

recommendations for Oracle ASM, [7-2](#)

M

Management Database, [7-9](#)
 management repository service, [8-3](#)
 manifest file, [7-26](#)
 mask
 setting default file mode creation mask, [5-24](#)
 mixed binaries, [3-5](#)
 mode
 setting default file mode creation mask, [5-24](#)
 Multiple Oracle Homes Support
 advantages, [C-2](#)
 multiversioning, [C-2](#)

N

Net Configuration Assistant (NetCA)
 response files, [B-8](#)
 running at command prompt, [B-8](#)
 netca.rsp file, [B-3](#)
 network requirements, [1-3](#)
 network tuning parameters
 configuring, [3-16](#)
 recommended values, [3-18](#)
 network, minimum requirements, [1-1](#)
 networks
 configuring interfaces, [4-23](#)
 for Oracle Flex Clusters, [4-18](#), [4-19](#)
 IP protocol requirements for, [4-2](#), [4-9](#)
 manual address configuration example, [4-22](#)
 Oracle Flex ASM, [1-3](#)
 NFS
 and data files, [6-9](#)
 buffer size requirements, [7-17](#)
 for data files, [6-9](#)
 NFS mounts
 Direct NFS Client
 requirements, [6-9](#)
 mtab, [6-9](#)
 orantstab, [6-9](#)
 noninteractive mode
 See response file mode

O

OCR
 See Oracle Cluster Registry
 OFA, [C-1](#)
 See also Optimal Flexible Architecture
 oifcfg, [4-3](#)
 OINSTALL directory, [C-5](#)
 oinstall group
 creating, [5-3](#)

- OINSTALL group, [1-5](#)
 - See also Oracle Inventory directory
- Opatch, [11-4](#)
- OPatch, [9-11](#)
- oper group
 - description, [5-9](#)
- operating system
 - different on cluster members, [3-5](#)
 - requirements, [3-5](#)
- operating system privileges groups, [1-5](#)
- operating system requirements, [1-2](#)
- Optimal Flexible Architecture, [C-1](#)
 - about, [C-1](#)
- Oracle ACFS
 - Installing Oracle RAC binaries not supported
 - on Oracle Flex Cluster, [6-4](#)
 - restrictions for Oracle Restart, [6-4](#)
 - supported IBM AIX versions, [6-3](#)
- Oracle ADVM
 - supported IBM AIX versions, [6-3](#)
- Oracle ASM
 - characteristics of failure groups, [7-2](#), [7-13](#)
 - configuring disk devices, [7-14](#)
 - disk groups, [7-2](#)
 - failure groups, [7-2](#)
 - examples, [7-2](#)
 - identifying, [7-2](#)
 - guidelines, [6-6](#)
 - installing, [8-7](#), [8-12](#), [8-17](#)
 - recommendations for disk groups, [7-2](#)
 - space required for preconfigured database, [7-2](#)
- Oracle ASM Filter Driver (Oracle ASMF), [8-17](#)
- Oracle ASM group
 - creating, [5-12](#)
- Oracle ASM on NFS, [6-7](#)
- Oracle ASM password file, [6-6](#)
- Oracle Automatic Storage Management
 - part of Oracle Grid Infrastructure installation, [7-10](#)
- Oracle base, [C-1](#), [C-5](#)
- Oracle Base, [5-5](#)
- Oracle base directory, [5-7](#)
- Oracle Cluster Registry, [1-7](#)
 - configuration of, [1-7](#)
- Oracle Clusterware Files
 - NFS to Oracle ASM, [11-9](#)
- Oracle Database
 - data file storage options, [6-5](#)
 - requirements with Oracle ASM, [7-2](#)
- Oracle Database Configuration Assistant, [B-7](#)
 - response file, [B-3](#)
- Oracle DBCA, [B-7](#)
- Oracle Disk Manager
 - and Direct NFS, [7-20](#)
- Oracle Domain Services Cluster, [8-3](#), [8-4](#)
- Oracle Enterprise Manager, [11-26](#)
- Oracle Extended Cluster, [8-7](#), [11-52](#)
 - converting to, [11-52](#)
- Oracle Extended Clusters, [8-5](#)
- Oracle Fleet Patching and Provisioning, [8-26](#), [11-18](#)
- Oracle Fleet Patching and Provisioning Server, [8-3](#)
- Oracle Flex ASM
 - and Oracle ASM clients, [1-3](#)
 - networks, [1-3](#)
- Oracle Flex Clusters
 - about, [4-18](#)
 - and Oracle Flex ASM, [4-19](#)
 - and Oracle Flex ASM cluster, [4-18](#)
- Oracle Restart
 - restrictions for, [6-4](#)
 - restrictions for Oracle ACFS, [6-4](#)
- Oracle FPP, [8-26](#), [11-18](#)
- Oracle Grid Infrastructure
 - restrictions for Oracle ACFS, [6-4](#)
- Oracle Grid Infrastructure upgrading, [11-18](#)
- Oracle home
 - ASCII path restriction for, [1-2](#)
 - file path, [C-5](#)
 - Grid home
 - filepath, [C-5](#)
 - naming conventions, [C-4](#)
- Oracle home directory, [5-7](#)
- Oracle Inventory, [1-5](#)
 - identifying existing, [5-2](#)
- Oracle Inventory Directory
 - OINSTALL group, [C-2](#)
- Oracle IO Server, [4-19](#)
- Oracle Layered File System, [8-26](#), [11-18](#)
- Oracle Member Clusters, [8-4](#)
 - for applications, [8-4](#)
 - for databases, [8-4](#)
- Oracle Net Configuration Assistant
 - response file, [B-3](#)
- Oracle Optimal Flexible Architecture
 - See Optimal Flexible Architecture
- Oracle ORAch
 - and Upgrade Readiness Assessment, [1-9](#)
- Oracle Software Owner user
 - creating, [5-5](#), [5-16](#)
- Oracle Software Owner users
 - configuring environment for, [5-23](#)
 - determining default shell, [5-24](#)
- Oracle Standalone Cluster, [8-7](#)
 - storage option, [6-7](#)
- Oracle Standalone Clusters, [8-3](#)
- Oracle Universal Installer
 - response files
 - list of, [B-3](#)
- Oracle Upgrade Companion, [3-2](#)

- oracle user, [1-5](#)
 - creating, [5-5](#)
- Oracle user
 - configuring environment for, [5-23](#)
 - determining default shell, [5-24](#)
 - modifying, [5-17](#)
- ORACLE_BASE environment variable
 - removing from shell startup file, [5-24](#)
- ORACLE_HOME environment variable
 - removing from shell startup file, [5-24](#)
- ORACLE_SID environment variable
 - removing from shell startup file, [5-24](#)
- oraInventory, [C-5](#)
- oranfstab configuration file, [7-18](#)
- oranfstab file, [7-20](#)
- OSASM group
 - creating, [5-12](#)
- OSBACKUPDBA group
 - creating, [5-14](#)
- OSBACKUPDBA group (backupdba), [5-9](#)
- OSDBA, [1-5](#)
- OSDBA for ASM
 - creating for Oracle Grid Infrastructure, [5-12](#)
- OSDBA groups
 - creating, [5-13](#)
 - creating for Oracle Grid Infrastructure, [5-13](#)
 - description for database, [5-9](#)
 - SYSDBA privilege, [5-9](#)
- OSDGDBA group
 - creating, [5-14](#)
- OSDGDBA group (dgdbs), [5-9](#)
- OSKMDBA group
 - creating, [5-14](#)
- OSKMDBA group (kmdbs), [5-9](#)
- OSOPER group
 - creating, [5-13](#)
- OSOPER groups
 - description for database, [5-9](#)
 - SYSOPER privilege, [5-9](#)
- OSRACDBA group
 - creating, [5-15](#)
- out-of-place patching, [11-23](#)

P

- parameter file
 - and permissions to read and write the SGA, [9-7](#)
- partition
 - using with Oracle ASM, [7-2](#)
- patch updates, [9-2](#)
- physical RAM, [2-2](#)
- policy-managed databases
 - and SCAN, [4-4](#)

- postinstallation
 - recommended tasks
 - root.sh script, backing up, [9-4](#)
- postinstallation -executeConfigTools option, [B-9](#)
- postinstallation configToolAllCommands script, [B-12](#)
- preconfigured database
 - Oracle ASM disk space requirements, [7-2](#)
 - requirements when using Oracle ASM, [7-2](#)
- preinstallation: rootpre.sh, [3-20](#)
- primary host name, [1-3](#)
- profile file, [5-24](#)
- proxy realm, [1-9](#)
- public node name
 - and primary host name, [1-3](#)

Q

- quorum disks, [6-7](#)
 - NFS mount, [6-7](#)

R

- racdba group
 - creating, [5-15](#)
- RAID
 - recommended Oracle ASM redundancy level, [7-2](#)
- recommendations
 - client access to the cluster, [9-6](#)
 - number of IP address for SCAN resolution, [4-4](#)
 - private network, [4-3](#)
- redundancy level
 - and space requirements for preconfigured database, [7-2](#)
- Redundant Interconnect Usage
 - IPv4 requirement, [4-7](#)
- registering resources, [11-26](#)
- release update revisions, [9-2](#)
- release updates, [9-2](#)
- releases
 - multiple, [C-2](#)
- relinking Oracle Grid Infrastructure home binaries, [9-11](#), [10-13](#)
- removing Oracle software, [10-1](#), [10-2](#)
 - examples, [10-6](#)
- requiremenets
 - interconnects, [4-7](#)
- requirements, [7-2](#)
- response file, [B-7](#)
- response file installation
 - preparing, [B-3](#)
 - response files
 - templates, [B-3](#)
 - silent mode, [B-5](#)

response file mode, [B-1](#)
 about, [B-1](#)
 reasons for using, [B-2](#)
 See also response files, silent mode

response files, [B-1](#), [B-7](#)
 about, [B-1](#)
 creating with template, [B-3](#)
 dbca.rsp, [B-3](#)
 enterprise.rsp, [B-3](#)
 general procedure, [B-2](#)
 Net Configuration Assistant, [B-8](#)
 netca.rsp, [B-3](#)
 passing values at command line, [B-1](#)
 specifying with Oracle Universal Installer, [B-5](#)
 See also silent mode.

root user
 logging in as, [2-1](#)

root.sh script
 backing up, [9-4](#)

rootcrs.pl
 restriction for Oracle Flex Cluster
 deinstallation, [10-14](#)

rootcrs.sh, [10-2](#)

roothas.sh, [10-2](#)

rootpre.sh, [3-20](#)

running gridSetup.sh, [11-13](#), [11-15](#)

running multiple Oracle releases, [C-2](#)

running rootpre.sh, [3-20](#)

S

sb_max, [7-17](#)

SCAN
 required for clients of policy-managed
 databases, [4-4](#)
 shared, [4-6](#)
 understanding, [4-4](#)

SCAN address, [1-3](#)

SCAN addresses, [4-4](#)

SCAN listeners, [4-4](#)

SCANS, [1-3](#), [4-14](#)
 client access, [9-6](#)
 configuring, [1-3](#)
 description, [9-6](#)

shared SCAN, [4-6](#)

shell
 determining default shell for Oracle user, [5-24](#)

SHELL environment variable
 checking value of, [5-24](#)

shell limits, [A-5](#)
 changing, [A-7](#)
 displaying, [A-7](#)

shell startup file
 editing, [5-24](#)
 removing environment variables, [5-24](#)

silent mode
 about, [B-1](#)
 reasons for using, [B-2](#)

silent mode installation, [B-5](#)

single client access names
 See SCAN addresses

smit command, [5-16](#)

software requirements, [3-5](#)

space requirements, [7-6](#)

ssh
 and X11 Forwarding, [5-27](#)
 configuring, [A-1](#)

Standard cluster
 upgrades result in, [11-4](#)

standard operating environment, [8-26](#)

startup file
 for shell, [5-24](#)

stty
 suppressing to prevent installation errors,
 [5-28](#)

swap space, [2-2](#)
 allocation, [1-2](#)

SYSBACKUPDBA system privileges, [5-9](#)

SYSDBA privilege
 associated group, [5-9](#)

SYSDBGDBA system privileges, [5-9](#)

SYSKMDBA system privileges, [5-9](#)

SYSOPER privilege
 associated group, [5-9](#)

system architecture, [2-2](#)

system configuration parameters, [A-5](#)
 verifying, [A-6](#)

system global area
 permissions to read and write, [9-7](#)

system privileges
 SYSBACKUPDBA, [5-9](#)
 SYSDBGDBA, [5-9](#)
 SYSKMDBA, [5-9](#)

system requirements, [1-1](#)

T

TEMP environment variable
 commands
 env, [5-24](#)
 env command, [5-24](#)
 environment
 checking settings, [5-24](#)
 setting, [5-24](#)
 umask, [5-24](#)
 umask command, [5-24](#)

terminal output commands
 suppressing for Oracle installation owner
 accounts, [5-28](#)

TMPDIR environment variable
 setting, [5-24](#)

troubleshooting

- cron jobs and installation, [1-9](#)
- disk space errors, [1-2](#)
- environment path errors, [1-2](#)
- garbage strings in script inputs found in log files, [5-28](#)
- inventory corruption, [5-17](#)
- root.sh errors, [10-14](#)
- ssh, [A-2](#), [A-5](#)
- ssh errors, [5-28](#)
- stty errors, [5-28](#)
- unconfiguring Oracle Clusterware to fix causes of root.sh errors, [10-14](#)
- unset environment variables, [1-2](#)
- user equivalency, [A-2](#), [A-5](#)

Troubleshooting

- DBCA does not recognize Oracle ASM disk size and fails to create disk groups, [9-10](#)

typographic conventions, [xiv](#)

U

umask command, [5-24](#)

unconfiguring Oracle Clusterware, [10-14](#)

uninstall

- See removing Oracle software

UNIX commands

- smit, [5-16](#)
- xhost, [2-1](#)

UNIX workstation

- installing from, [2-1](#)

unreachable nodes

- upgrading, [11-18](#)

upgrade, [3-3](#), [11-18](#)

- Oracle Automatic Storage Management, [7-10](#)
- running gridSetup.sh, [11-13](#), [11-15](#)

upgrade (*continued*)

upgrade tasks, [11-26](#)

upgrades, [11-4](#)

- and SCAN, [4-4](#)
- best practices, [3-2](#)
- restrictions for, [11-4](#)
- unsettling environment variables for, [11-8](#)

upgrading

- and Oracle ORAchk Upgrade Readiness Assessment, [1-9](#)
- inaccessible nodes, [11-18](#)
- options, [3-3](#)

user process parameters

- configuring, [3-16](#)

useradd command, [5-16](#)

users

- creating the oracle user, [5-5](#)

V

vendor clusterware

- and cluster names for Oracle Grid Infrastructure, [1-3](#)

verifying system parameter values, [A-6](#)

Virtual Memory Manager, [3-14](#)

voting files

- configuration of, [1-7](#)

X

X Window System

- enabling remote hosts, [2-1](#)

X11 forwarding errors, [5-27](#), [A-5](#)

xhost command, [2-1](#)

xtitle

- suppressing to prevent installation errors, [5-28](#)